



INFOCOMP 2015

The Fifth International Conference on Advanced Communications and
Computation

ISBN: 978-1-61208-416-9

EMPIRICAL MODELING 2015

The International Symposium on Empirical Modeling

June 21 - 26, 2015

Brussels, Belgium

INFOCOMP 2015 Editors

Claus-Peter Rückemann, Westfälische Wilhelms-Universität Münster / Leibniz
Universität Hannover / North-German Supercomputing Alliance, Germany
Malgorzata Pankowska, University of Economics, Katowice, Poland
Ian Flood, University of Florida, USA

INFOCOMP 2015

Foreword

The Fifth International Conference on Advanced Communications and Computation (INFOCOMP 2015), held between June 21-26, 2015, in Brussels, Belgium, continued a series of events dedicated to advanced communications and computing aspects, covering academic and industrial achievements and visions.

The diversity of semantics of data, context gathering and processing led to complex mechanisms for applications requiring special communication and computation support in terms of volume of data, processing speed, context variety, etc. The new computation paradigms and communications technologies are now driven by the needs for fast processing and requirements from data-intensive applications and domain-oriented applications (medicine, geo-informatics, climatology, remote learning, education, large scale digital libraries, social networks, etc.). Mobility, ubiquity, multicast, multi-access networks, data centers, cloud computing are now forming the spectrum of de facto approaches in response to the diversity of user demands and applications. In parallel, measurements control and management (self-management) of such environments evolved to deal with new complex situations.

INFOCOMP 2015 also featured the following Symposium

- EMPIRICAL MODELING 2015: The International Symposium on Empirical Modeling

We take here the opportunity to warmly thank all the members of the INFOCOMP 2015 Technical Program Committee, as well as the numerous reviewers. The creation of such a broad and high quality conference program would not have been possible without their involvement. We also kindly thank all the authors who dedicated much of their time and efforts to contribute to INFOCOMP 2015. We truly believe that, thanks to all these efforts, the final conference program consisted of top quality contributions.

Also, this event could not have been a reality without the support of many individuals, organizations, and sponsors. We are grateful to the members of the INFOCOMP 2015 organizing committee for their help in handling the logistics and for their work to make this professional meeting a success.

We hope that INFOCOMP 2015 was a successful international forum for the exchange of ideas and results between academia and industry and for the promotion of progress in the areas of communications and computations.

We are convinced that the participants found the event useful and communications very open. We hope that Brussels, Belgium, provided a pleasant environment during the conference and everyone saved some time to enjoy the charm of the city.

INFOCOMP 2015 Chairs:

INFOCOMP General Chair

Claus-Peter Rückemann, Westfälische Wilhelms-Universität Münster / Leibniz Universität Hannover / North-German Supercomputing Alliance, Germany

INFOCOMP Advisory Chairs

Hans-Joachim Bungartz, Technische Universität München (TUM) - Garching, Germany
Petre Dini, Concordia University - Montreal, Canada / China Space Agency Center - Beijing, China
Sik Lee, Supercomputing Center / Korea Institute of Science and Technology Information (KISTI), Korea
Subhash Saini, NASA, USA
Manfred Krafczyk, Institute for Computational Modeling in Civil Engineering (iRMB) - Braunschweig, Germany
Ian Flood, University of Florida, USA

INFOCOMP Academia Chairs

Alexander Knapp, Universität Augsburg, Germany
Malgorzata Pankowska, University of Economics, Katowice, Poland

INFOCOMP Research Institute Liaison Chairs

Kei Davis, Los Alamos National Laboratory, USA
Edgar A. Leon, Lawrence Livermore National Laboratory, USA
Ivor Spence, School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Northern Ireland, Research for High Performance and Distributed Computing / Queen's University Belfast, UK

INFOCOMP Industry Chairs

Alfred Geiger, T-Systems Solutions for Research GmbH, Germany
Hans-Günther Müller, Cray, Germany

INFOCOMP Special Area Chairs on Large Scale and Fast Computation

José Gracia, High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany
Björn Hagemeier, Juelich Supercomputing Centre, Forschungszentrum Juelich GmbH, Germany
Walter Lioen, SURFsara, Netherlands
Lutz Schubert, Institute of Information Resource Management, University of Ulm, Germany

INFOCOMP Special Area Chairs on Networks and Communications

Noelia Correia, University of the Algarve, Portugal
Wolfgang Hommel, Leibniz Supercomputing Centre - Munich, Germany

INFOCOMP Special Area Chairs on Advanced Applications

Bernhard Bandow, Max Planck Institute for Solar System Research (MPS), Göttingen, Germany
Diglio A. Simoni, RTI International - Research Triangle Park, USA

INFOCOMP Special Area Chairs on Evaluation Context

Huong Ha, University of Newcastle, Australia (Singapore campus)
Philipp Kremer, German Aerospace Center (DLR), Institute of Robotics and Mechatronics, Oberpfaffenhofen, Germany

INFOCOMP Special Area Chairs on Biometry

Ulrich Norbisch, Nazarbayev University, Kazakhstan / BIOMETRY.com, Switzerland

INFOCOMP 2015

COMMITTEE

INFOCOMP General Chair

Claus-Peter Rückemann, Westfälische Wilhelms-Universität Münster / Leibniz Universität Hannover / North-German Supercomputing Alliance, Germany

INFOCOMP Advisory Chairs

Hans-Joachim Bungartz, Technische Universität München (TUM) - Garching, Germany
Petre Dini, Concordia University - Montreal, Canada / China Space Agency Center - Beijing, China
Sik Lee, Supercomputing Center / Korea Institute of Science and Technology Information (KISTI), Korea
Subhash Saini, NASA, USA
Manfred Krafczyk, Institute for Computational Modeling in Civil Engineering (iRMB) - Braunschweig, Germany

INFOCOMP Academia Chairs

Alexander Knapp, Universität Augsburg, Germany
Malgorzata Pankowska, University of Economics, Katowice, Poland

INFOCOMP Research Institute Liaison Chairs

Kei Davis, Los Alamos National Laboratory, USA
Edgar A. Leon, Lawrence Livermore National Laboratory, USA
Ivor Spence, School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Northern Ireland, Research for High Performance and Distributed Computing / Queen's University Belfast, UK

INFOCOMP Industry Chairs

Alfred Geiger, T-Systems Solutions for Research GmbH, Germany
Hans-Günther Müller, Cray, Germany

INFOCOMP Special Area Chairs on Large Scale and Fast Computation

José Gracia, High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany
Björn Hagemeier, Juelich Supercomputing Centre, Forschungszentrum Juelich GmbH, Germany
Walter Lioen, SURFsara, Netherlands
Lutz Schubert, Institute of Information Resource Management, University of Ulm, Germany

INFOCOMP Special Area Chairs on Networks and Communications

Noelia Correia, University of the Algarve, Portugal
Wolfgang Hommel, Leibniz Supercomputing Centre - Munich, Germany

INFOCOMP Special Area Chairs on Advanced Applications

Bernhard Bandow, Max Planck Institute for Solar System Research (MPS), Göttingen, Germany
Diglio A. Simoni, RTI International - Research Triangle Park, USA

INFOCOMP Special Area Chairs on Evaluation Context

Huong Ha, University of Newcastle, Australia (Singapore campus)
Philipp Kremer, German Aerospace Center (DLR), Institute of Robotics and Mechatronics,
Oberpfaffenhofen, Germany

INFOCOMP Special Area Chairs on Biometry

Ulrich Norbistrath, Nazarbayev University, Kazakhstan / BIOMETRY.com, Switzerland

INFOCOMP 2015 Technical Program Committee

Wassim Abu Abed, Institute for Computational Modelling in Civil Engineering - Technische Universität Braunschweig, Germany
Ajith Abraham, Machine Intelligence Research Labs (MIR Labs), USA
Mehmet Aksit, University of Twente, Netherlands
Ali I. Al Mussa, King Abdulaziz City for Science and Technology (KACST), Saudi Arabia
Angelos-Christos Anadiotis, School of Electrical and Computer Engineering (SECE), National Technical University of Athens (NTUA), Greece
Daniel Andresen, Kansas State University, USA
Bernadetta Kwintiana Ane, Institute of Computer-aided Product Development Systems, Universität Stuttgart, Germany
Douglas Archibald, University of Ottawa, Canada
John Ashley, NVIDIA Corporation, USA
Hazeline U. Asuncion, University of Washington, Bothell, USA
Simon Reay Atkinson, The University of Sydney, Australia
Bernhard Bandow, Max Planck Institute for Solar System Research (MPS), Göttingen, Germany
Enobong Basse, Auckland University of Technology, New Zealand
Khalid Belhajjame, Université Paris Dauphine, France
Belgacem Ben Youssef, King Saud University Riyadh, KSA / Simon Fraser University Vancouver, British Columbia, Canada
Martin Berzins, University of Utah, USA
Rupak Biswas, NASA Ames Research Center, USA
Rim Bouhouch, National Engineering School of Tunis, Tunisia
Steffen Brinkmann, High Performance Computing Centre Stuttgart (HLRS), Germany
Elzbieta Bukowska, Poznan University of Economics, Poland
Hans-Joachim Bungartz, Technische Universität München (TUM) - Garching, Germany
Diletta Romana Cacciagrano, Computer Science Division, School of Science and Technology, University of Camerino, Italy
Xiao-Chuan Cai, University of Colorado Boulder, USA
Elena Camossi, European Commission, Joint Research Centre, Institute for the Protection and Security of the Citizen (IPSC) - Ispra, Italy

Antonio Martí Campoy, Universitat Politècnica de València, Spain
Laura Carrington, University of California, San Diego/ San Diego Supercomputer Center, USA
Emre Celebi, Louisiana State University in Shreveport, USA
Hsi-Ya Chang, National Center for High-Performance Computing (NCHC), Taiwan
Jian Chang, Bournemouth University, UK
Chien-Hsing Chou, Tamkang University, Taiwan
Jerry Chi-Yuan Chou, National Tsing Hua University, Taiwan
Sung-Bae Cho, Yonsei University, Seoul, Korea
Noelia Correia, University of Algarve, Portugal
Kei Davis, Los Alamos National Laboratory / Computer, Computational, and Statistical Sciences Division, USA
Sergio De Agostino, Sapienza University-Rome, Italy
Flávio de Oliveira Silva, Universidade Federal de Uberlândia, MG, Brasil
Vieri del Bianco, Università dell'Insubria, Italia
Šandor Dembitz, University of Zagreb, Faculty of Electrical Engineering and Computing, Croatia
Amine Dhraief, University of Kairouan, Tunisia
Beniamino Di Martino, Dipartimento di Ingegneria dell'Informazione, Seconda Università di Napoli, Italy
Zhong-Hui Duan, University of Akron, USA
Truong Vinh Truong Duy, JAIST / University of Tokyo, Japan
Jürgen Falkner, Fraunhofer-Institut für Arbeitswirtschaft und Organisation IAO, Stuttgart, Germany
Mohammed Farfour, Energy & Resources Eng. Dept., Chonnam National University, Gwangju, South Korea
Lars Fischer, Research Group for IT Security, University of Siegen, Germany
Ian Flood, Rinker School, College of Design, Construction and Planning, University of Florida, USA
María Consuelo Franky, Department of Computer Science, Pontificia Universidad Javeriana, Bogotá, Colombia
Sara de Freitas, Curtin University, Australia
Dariusz Frejlichowski, West Pomeranian University of Technology, Poland
Munehiro Fukuda, Division of Computing and Software Systems, School of Science, Technology Engineering, and Math, University of Washington, Bothell, USA
Marcelo Garcia, Cray Inc., Germany
Marta Gatiús, Technical University of Catalonia, Spain
José Gracia, High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany
Alfred Geiger, T-Systems Solutions for Research GmbH, Germany
Andy Georgi, Dresden University of Technology, Germany
Birgit Frida Stefanie Gersbeck-Schierholz, Leibniz Universität Hannover/Certification Authority University of Hannover (UH-CA), Germany
Franca Giannini, Consiglio Nazionale delle Ricerche - Genova, Italy
Fabio Gomes de Andrade, Federal Institute of Science, Education and Technology of Paraíba, Brazil
Carina González, University of La Laguna, Spain
Conceicao Granja, Tromsø Telemedicine Laboratory - Norwegian Centre for Telemedicine, University Hospital of North Norway, Tromsø, Norway
Richard Gunstone, Bournemouth University, UK
Huong Ha, University of Newcastle, Australia (Singapore campus)
Björn Hagemeyer, Juelich Supercomputing Centre, Forschungszentrum Juelich GmbH, Germany
Malgorzata Hanzl, Technical University of Lodz, Poland
Abbas Hijazi, Lebanese University, Lebanon
Daniel Holmes, University of Edinburgh, UK

Wolfgang Hommel, Leibniz Supercomputing Centre - Munich, Germany
Tzyy-Leng Horng, Feng Chia University, Taiwan
Kuo-Chan Huang, National Taichung University of Education, Taiwan
Friedrich Hülsmann, Gottfried Wilhelm Leibniz Bibliothek - Hannover, Germany
Chih-Cheng Hung, Southern Polytechnic State University, USA
Udo Inden, Cologne University of Applied Sciences, Research Centre for Applications of Intelligent Systems (CAIS), Germany
Lucjan Jacak, Wroclaw University of Technology - Institute of Physics, Poland
Oleg Jakushkin, Saint Petersburg State University, Russia
Jinyuan Jia, Tongji University, China
Kai Jiang, Shanghai Supercomputer Center, China
Seifedine Kadry, American University of the Middle East, Kuwait
David Kaeli, Northeastern University, USA
Tae-Wook Kang, KICT (Korea Institute of Construction Technology), Korea
Izabela Karsznia, University of Warsaw, Department of Cartography, Warsaw, Poland
Christos Kartsaklis, Oak Ridge National Laboratory (ORNL), USA
Stavros Kassinos, University of Cyprus, Cyprus
Takahiro Kawamura, Toshiba Corporation, Japan
Abdelmajid Khelil, Huawei Research, Germany
Jinoh Kim, Texas A&M University-Commerce, USA
Marie Kim, Electronics and Telecommunications Research Institute (ETRI), Republic of Korea
Alexander Kipp, Robert Bosch GmbH., Germany
Christos Kloukinas, City University London, UK
Alexander Knapp, University of Augsburg, Germany
Manfred Krafczyk, Institute for Computational Modeling in Civil Engineering (iRMB), Braunschweig, Germany
Philipp Kremer, German Aerospace Center (DLR) / Institute of Robotics and Mechatronics - Oberpfaffenhofen, Germany
Herbert Kuchen, Westfälische Wilhelms-Universität Münster, Institut für Wirtschaftsinformatik, Praktische Informatik in der Wirtschaft, Münster, Germany
Michael Lang, Los Alamos National Lab, USA
Robert S. Laramée, Swansea University, UK
Sik Lee, Supercomputing Center / Korea Institute of Science and Technology Information (KISTI), Korea
Paulo Leitão, Polytechnic Institute of Braganca, Portugal
Edgar A. Leon, Lawrence Livermore National Laboratory, USA
Walter Lion, SURFsara, Netherlands
Yanting Li, Kyushu Institute of Technology, Iizuka, Japan
Fotis Liarokapis, Coventry University, UK
Georgios Lioudakis, National Technical University of Athens (NTUA), Greece
Maciej Liśkiewicz, Universität zu Lübeck, Germany
Frank Löffler, Louisiana State University, USA
Darrell Long, University of California, USA
Joan Lu, Informatics, University of Huddersfield, UK
Richard Lucas, University of Canberra and Charles Sturt University, Australia
Tony Maciejewski, Colorado State University, USA
Janna L. Maltseva, Lavrentyev Institute of Hydrodynamics, Novosibirsk, Russia
Dirk Malzahn, Dirk Malzahn Ltd. / HfH University, Germany
Suresh Marru, Indiana University, USA

Nikolaos Matsatsinis, Technical University of Crete, Greece
Lois Curfman McInnes, Mathematics and Computer Science Division - Argonne National Laboratory, USA
Artis Mednis, Institute of Electronics and Computer Science, Latvia
Igor Melatti, Sapienza Università di Roma, Rome, Italy
Despina Meridou, School of Electrical and Computer Engineering (SECE), National Technical University of Athens (NTUA), Greece
Jelena Mirkovic, Center for Shared Decision Making and Collaborative Care Research, Oslo University Hospital, Norway
Hans-Günther Müller, Cray, Germany
Marian Mureşan, Faculty of Mathematics and Computer Science, Babes-Bolyai University, Cluj-Napoca, Romania
Syed Naqvi, Birmingham City University, UK
Lena Noack, Royal Observatory of Belgium, Brussels, Belgium
António Nogueira, University of Aveiro / Instituto de Telecomunicações, Portugal
Ulrich Norbistrath, Nazarbayev University, Kazakhstan / BIOMETRY.com, Switzerland
Krzysztof Okarma, West Pomeranian University of Technology, Poland
Aida Omerovic, SINTEF, Norway
Dhabaleswar K. (DK) Panda, Ohio State University, USA
Malgorzata Pankowska, University of Economics - Katowice, Poland
Maria Eleftheria Papadopoulou, National Technical University of Athens (NTUA), Greece
Giuseppe Patané CNR-IMATI, Italy
Raffaella Pavani, Department of Mathematics, Politecnico di Milano, Italy
Guo Peiqing, Shanghai Supercomputer Center, China
Christian Percebois, University of Toulouse, France
Ana-Catalina Plesa, German Aerospace Center, Institute of Planetary Research, Planetary Physics, Berlin, Germany
Matthias Pocs, Stellar Security Technology Law Research, Germany
Daniela Pöhn, Leibniz Supercomputing Centre, Munich, Germany
Mario Porrman, Center of Excellence Cognitive Interaction Technology - Bielefeld University, Germany
Thomas E. Potok, Computational Data Analytics Group, Oak Ridge National Laboratory, USA
Giovanni Puglisi, University of Catania, Italy
Iouldouz Raguimov, York University, Canada
Mohamed A. Rashed, Department of Geophysics, Faculty of Earth Sciences, King Abdulaziz University, Jeddah, Saudi Arabia
Andreas Rausch, Technische Universität Clausthal, Germany
Ustijana Rechkoska Shikoska, University for Information Science & Technology "St. Paul the Apostle" - Ohrid, Republic of Macedonia
Yenumula B. Reddy, Department of Computer Science, Grambling State University, USA
Shangping Ren, Illinois Institute of Technology - Chicago, USA
Theresa-Marie Rhyne, Visualization Consultant, Durham, USA
Sebastian Ritterbusch, Engineering Mathematics and Computing Lab (EMCL), Karlsruhe Institute of Technology (KIT), Germany
Alessandro Rizzi, Università degli Studi di Milano, Italy
Ivan Rodero, Rutgers University - Piscataway, USA
Dieter Roller, University of Stuttgart, Director Institute of Computer-aided Product Development Systems - Stuttgart, Germany
Claus-Peter Rückemann, WWU Münster / Leibniz Universität Hannover / HLRN, Germany
H. Birali Runesha, University of Chicago, USA

Subhash Saini, NASA, USA
Lutz Schubert, Institute of Information Resource Management, University of Ulm, Germany
Marla Scheweppe, Rochester Institute of Technology, USA
Isabel Schwerdtfeger, IBM, Germany
Yaroslav Sergeev, University of Calabria, Italy
Damián Serrano, University of Grenoble - LIG, France
Gyuzel Shakhmametova, Ufa State Aviation Technical University, Russia
Diglio A. Simoni, RTI International - Research Triangle Park, USA
Theodore E. Simos, King Saud University & University of Peloponnese - Tripolis, Greece
Happy Sithole, Center for High Performance Computing - Cape Town, South Africa
Marc Snir, University of Illinois at Urbana Champaign, USA
Marcin Sokół, Gdansk University of Technology, Poland
Terje Sovoll, Tromsø Telemedicine Laboratory - Norwegian Centre for Telemedicine, University Hospital of North Norway, Tromsø, Norway
Ivor Spence, School of Electronics, Electrical Engineering and Computer Science, Research for High Performance and Distributed Computing, Queen's University Belfast, UK
Rolf Sperber, Embrace, HPC-Network Consulting, Germany
Monika Steinberg, University of Applied Sciences and Arts Hanover, Germany
Mark Stillwell, Cranfield University, UK
Christian Straube, MNM-Team, Institut für Informatik, Ludwig-Maximilians-Universität München (LMU), Germany
Mu-Chun Su, National Central University, Taiwan
Gerson Sunyé, University of Nantes, France
Zdenek Sustr, CESNET, Czech Republic
Seyedamir Tavakoli Taba, Complex Civil Systems Research Group and Project Management Programme, The University of Sydney, Australia
Rahim Tafazolli, CCSR Director, University of Surrey - Guildford, UK
Zhiqi Tao, Intel Corporation, USA
Dominique Thiebaut, Smith College, USA
Vrizlynn Thing, Institute for Infocomm Research, Singapore
Yuan Tian, Oak Ridge National Laboratory, USA
Francesco Tiezzi, IMT Institute for Advanced Studies Lucca, Italy
Daniel Thalmann, Institute for Media Innovation (IMI) - Nanyang Technological University, Singapore
Katya Toneva, International Community School and Middlesex University, London, UK
Rafael P. Torchelsen, Universidade Federal da Fronteira Sul, Brazil
Nicola Tosi, Department of Planetary Geodesy, Technical University Berlin, Germany
Bernard Traversat, Oracle, USA
Dan Tulpan, Information and Communications Technologies - National Research Council Canada / University of Moncton / University of New Brunswick / Atlantic Cancer Research Institute, Canada
Ravi Vadapalli, Texas Tech University, USA
Marten Van Dijk, University of Connecticut, USA
Lisette Van Gemert-Pijnen, University of Twente and University of Groningen, University Medica Centre, The Netherlands
Ana Lucia Varbanescu, University of Amsterdam, Netherlands
Domitila Violeta Velasco Mansilla, Hydrogeology Group (GHS) Institute of Environmental Assessment and Water Research (IDAEA-CSIC), Barcelona, Spain
Claire Vishik, Intel Corporation, UK
Krzysztof Walczak, Poznan University of Economics, Poland

Edward Walker, Whitworth University, USA
Iris Weber, Institut für Planetologie, Westfälische Wilhelms-Universität Münster, Germany
Jacqueline Whalley, Auckland University of Technology, New Zealand
Stefan Wild, Argonne National Laboratory, USA
Wojciech Wiza, Poznan University of Economics, Poland
Michael Wrinn, Intel - Corporate Research division, USA
Dongrong Xu, Columbia University, USA
Qimin Yang, Engineering Department, Harvey Mudd College, Claremont, CA, USA
Yi Yang, NEC Laboratories America, USA
Hongchuan Yu, Bournemouth University, UK
May Yuan, University of Texas at Dallas, USA
Peter Zaspel, University of Bonn, Germany
Yu-Xiang Zhao, National Quemoy University, Taiwan

Symposium Chairs

Ian Flood, University of Florida, USA
Claus-Peter Rückemann, Westfälische Wilhelms-Universität Münster / Leibniz Universität Hannover /
North-German Supercomputing Alliance, Germany

Program Committee Members

Alnoor Allidina, IBI-MAAK Inc., Canada
Turgay Aytac, Prescience Technologies Inc., USA
Catherine Cleophas, RWTH Aachen University, Germany
Madeline M. Diep, Fraunhofer Center for Experimental Software Engineering, USA
Ian Flood, University of Florida, USA
Rachel Harrison, Oxford Brookes University, UK
Hadi Hemmati, University of Manitoba, Canada
Frank Herrmann, OTH Regensburg - Technical University of Applied Sciences, Germany
Mauro Iacono, Seconda Università degli Studi di Napoli, Italy
Raymond R.R. Issa, Rinker School/University of Florida, USA
SangHyun Lee, Department of Civil and Environmental Engineering/University of Michigan, USA
Ulf Lotzmann, University of Koblenz-Landau, Germany
Roderick Melnik, Wilfrid Laurier University, Canada
Daniel Méndez Fernández, Technische Universität München, Germany
Akbar Siami Namin, Texas Tech University, USA
Federica Paci, University of Southampton, UK
Claus-Peter Rückemann, Westfälische Wilhelms-Universität Münster / Leibniz Universität Hannover /
North-German Supercomputing Alliance, Germany
Ravi S. Srinivasan, Rinker School/University of Florida, USA
Stefan Wagner, University of Stuttgart, Germany
Dietmar Winkler, Vienna University of Technology, Austria
Yimin Zhu, Department of Construction Management/Louisiana State University, USA

Copyright Information

For your reference, this is the text governing the copyright release for material published by IARIA.

The copyright release is a transfer of publication rights, which allows IARIA and its partners to drive the dissemination of the published material. This allows IARIA to give articles increased visibility via distribution, inclusion in libraries, and arrangements for submission to indexes.

I, the undersigned, declare that the article is original, and that I represent the authors of this article in the copyright release matters. If this work has been done as work-for-hire, I have obtained all necessary clearances to execute a copyright release. I hereby irrevocably transfer exclusive copyright for this material to IARIA. I give IARIA permission to reproduce the work in any media format such as, but not limited to, print, digital, or electronic. I give IARIA permission to distribute the materials without restriction to any institutions or individuals. I give IARIA permission to submit the work for inclusion in article repositories as IARIA sees fit.

I, the undersigned, declare that to the best of my knowledge, the article does not contain libelous or otherwise unlawful contents or invading the right of privacy or infringing on a proprietary right.

Following the copyright release, any circulated version of the article must bear the copyright notice and any header and footer information that IARIA applies to the published article.

IARIA grants royalty-free permission to the authors to disseminate the work, under the above provisions, for any academic, commercial, or industrial use. IARIA grants royalty-free permission to any individuals or institutions to make the article available electronically, online, or in print.

IARIA acknowledges that rights to any algorithm, process, procedure, apparatus, or articles of manufacture remain with the authors and their employers.

I, the undersigned, understand that IARIA will not be liable, in contract, tort (including, without limitation, negligence), pre-contract or other representations (other than fraudulent misrepresentations) or otherwise in connection with the publication of my work.

Exception to the above is made for work-for-hire performed while employed by the government. In that case, copyright to the material remains with the said government. The rightful owners (authors and government entity) grant unlimited and unrestricted permission to IARIA, IARIA's contractors, and IARIA's partners to further distribute the work.

Table of Contents

Advancements in Tape Outpace Disk Technology <i>Isabel Schwerdtfeger</i>	1
An Adaptive Load-balancer for Task-scheduling in FastFlow <i>Md Moniruzzaman, Kamran Idrees, Michael Rossbory, and Jose Gracia</i>	6
Adaption of the n-way Dissemination Algorithm for GASPI Split-Phase Allreduce <i>Vanessa End, Ramin Yahyapour, Christian Simmendinger, and Thomas Alrutz</i>	13
Gauss-Jordan Matrix Inversion Speed-Up using GPUs with the Consideration of Power Consumption <i>Mahmoud Shirazi, Mehdi Kargahi, and Farshad Khunjush</i>	20
Towards Data Persistency for Fault-tolerance Using MPI Semantics <i>Jose Gracia, Muhammad Wahaj Sethi, Nico Struckmann, and Rainer Keller</i>	26
QoS versus QoE for Database Tuning <i>Malgorzata Pankowska and Jaroslaw Kurpanik</i>	30
Integrated Tool Support in the Context of Security and Network Management Convergence <i>Felix von Eye, David Schmitz, and Wolfgang Hommel</i>	36
Design and Implementation of a Scalable SDN-OF Controller Cluster <i>Min Luo, Quancai Li, Bo Man, Ke Lin, Xiaorong Wu, Chenji Li, Shen Lu, and Wu Chou</i>	43
Management Information Systems in Health Sector: Evidence of Mandatory Use <i>Ioannis Mitropoulos</i>	51
Designing Security Policies for Complex SCADA Systems Protection <i>Djamel Khadraoui and Christophe Feltus</i>	54
Impact of the Entering Time on Collective Performance <i>Christoph Niethammer, Dmitry Khabi, Huan Zhou, Vladimir Marjanovic, and Jose Gracica</i>	60
Remote Control and Monitoring of Smart Home Facilities via Smartphone with Wi-Fly <i>Omar Ghabar and Joan Lu</i>	66
Adaptive Noise Reduction in Ultrasonic Images <i>Somkait Udomhunsakul</i>	74
Writer Identification Method using Inter and Intra-Stroke Information	80

Jungpil Shin and Zhaofeng Liu

Habitability, Interior and Crust: A new Code for Modeling the Thermal Evolution of Planets and Moons 84
Lena Noack, Attilio Rivoldini, and Tim Van Hoolst

Creation of Objects and Concordances for Knowledge Processing and Advanced Computing 91
Claus-Peter Ruckemann

An Investigation into Game Based Learning Using High Level Programming Languages 99
Ragab Ihnissi and Joan Lu

Combining Artificial Bee Colony and Genetic Algorithms to Enhance the GPGPU-based ANN Classifier for Identifying Students with Learning Disabilities 105
Tung-Kuang Wu, Shian-Chang Huang, Ying-Ru Meng, Chin-Yu Hsu, and Chih-Han Tai

Analysis of GLCM Parameters for Textures Classification on UMD Database Images 111
Alsadegh Mohamed and Joan Lu

Copyright in Multiscale Cancer Modeling 117
Iryna Lishchuk and Marc Stauch

Assessing Creativity: A Test for Drawing Production using Digital Art Tools 123
Aber Aboalgasm and Rupert Ward

Nonparametric Estimation of Demand Structures in Airline Revenue Management 130
Johannes Ferdinand Jorg and Catherine Cleophas

Environmental Performance Evaluation of Different Glazing-Sunshade Systems Using Simulation Tools 134
Adeeba A. Raheem, Raja R. A. Issa, Svetlana Olbina, and Ian Flood

Comparison of Artificial Neural Networks and Support Vector Machines for Weigh-In-Motion Based Truck Type Classification 140
Yueren Wang, Ian Flood, and Raja R A Issa

Modeling Interactive Digital TV Users Behavior 146
Samuel Basilio

Advancements in Tape outpace Disk Technology

Isabel Schwerdtfeger

System Services, Global Technology Services

IBM Germany

Hamburg, Germany

email: schwerdtfeger@de.ibm.com

Abstract—Tape technologies are able to effectively address many new data intensive market opportunities. Tape has traditionally been a primary backup device, but it no longer does that alone. Recent tape advances have made greater progress than disk technologies over the past 10 years. Today’s modern tape technology is able to efficiently store huge amounts of data in a cost and efficient way. This paper examines the hypothesis that tape shall be viewed as an active or “near-line” storage component, essential in any petabyte (PB) or exabyte (EB) storage environment.

Keywords—Archive; Capacity; Costs; Disk; High Performance Computing (HPC); High Performance Storage System (HPSS); Linear Tape Open (LTO); Performance; Tape; Total Cost of Ownership (TCO).

I. INTRODUCTION

By 2019, 75 percent of organizations will treat archived data as an active and "near line" data source, and not simply as a separate repository to be viewed or searched periodically, up from less than 10 percent today [1]. Therefore, storage systems continue to play a central role in strategic planning in organizations despite effective optimization tools for storage software, tiered storage architectures, and effective data management policies in place. Back in 2006, the worldwide storage market including the total disk and tape storage hardware segment increased to \$28.2 billion [2]. At that time, IBM was the #1 seller in storage technology, exactly 50 years after the first IBM System RAMAC 350 storage product was launched in 1956 with a maximal capacity of 5 megabytes (MB) (for 5 million characters) [3].

Despite the enormous data growth in the last ten years, the storage market has come down to a fragmented market view, where today the measurements are based solely on Big Data, Virtualization, Enterprise Information Archives, or just by Tape Drive level. The Big Data market is expected to have a compound annual growth rate of 28.5 percent that will reach \$5.89 billion in 2018 in disk storage systems [4]. That is a pure fraction of what has been shipped and installed into disk systems since 2006 and it shows how sharp the costs for disk storage systems have come down, despite the huge data growth needs spurred on by digital-natives, smartphones, and supercomputers up to this day.

Already in 2006, and long before it was widely known, tape technology was already "green", to conserve energy, power, and cooling costs. Tape is the cheapest storage for storing large amount of data. Tape cartridges only consume power when read or written [3]. In a recent memo by the Tape Storage Council (TSC) the current trends, usages, and technology innovations occurring within the tape storage industry were described as cartridge capacity increase, much longer media life, and improved bit error rates [5]. The publication of this memo by leading vendors suggests that today’s modern tape technology is nothing like the tape of the past.

For tape technologies two main types exists: Linear Tape Open (LTO) and Enterprise. LTO is a joint specification for ½ inch magnetic tape and its corresponding tape drives. IBM, HP and Seagate jointly worked on this development since the late 1990s. Due to the sale of Seagate to Quantum, Quantum is now part of the LTO consortium. The LTO consortium issues in regular intervals the LTO Ultrium Roadmap. On September 10th, 2014, the newest “Generation 10” Roadmap was released. The newest LTO Generation 10 will be able to store up to 120 Terabyte (TB) on a single tape cartridge and a read/write performance up to 2750 megabyte per second (MB/s) [6]. As for Enterprise tape, that is IBM 3592, a series of tape drives and corresponding magnetic tape data storage media formats developed by IBM.

Significant technology innovations in tape technologies address the constant demand for improved reliability, higher capacity, power efficiency, ease of use and the lowest cost per gigabytes (GB) of any storage solution [5]. A management of PB or multi-PB disk infrastructures is very difficult to build and maintain as things break much more frequently [7]. Furthermore, high demands for upcoming EB storage systems, and the need for efficient data management solutions were already addressed by the U. S. Department of Energies (DOEs) in 2009 and summarized in a report to the U.S. Government. Alone one PB of storage is roughly equivalent of 210,000 DVDs [8]. Fig. 1 shows the exponential data growth by the DOEs of archived data projected until 2022. It shows that the DOEs will have single archive systems that retain between 2 and 50 EB of data.

Numerous sources, including analysts, consultants, IBM and Oracle, predict in a 10-year projection for technology costs per terabytes a steady decline in price for flash, disks and tape from 2014 until 2023. The compound growth rate

(CGR) for flash is predicted to reduce by -30%, for disk by -15% and for tape -23% [9].




Figure 1. Archived Data stored by the DOE Lab through 2018-2022 [13]

While daily data grows exponentially, there will be an increasing gap, if this data should be managed on disks systems alone. In three years two and a half more disks are necessary to cope with storage needs. Today a disk backup with 600TB with about 300 disks at 7,5 kilowatt (kw) will be in three years 1,8 PB with ca. 750 disks at 18,7kw [10]. Together including the recent tape technology advancements, this gap could be better addressed with tape instead.

The introduction in Section I illustrates the dynamics of huge data growth and describes market and trends of storage systems. In addition, basic tape technologies are being described. Section II lists the recent developments of tape advancements, while Section III compares tape benefits to disk technologies including the aspect of Total Cost of Ownership (TCO). Section IV describes an example of an active archive by its benefit of using tape technologies with use of optimized software. Big Data software, such as the High Performance Storage System (HPSS) used by scientists and researchers today, supports the growing data needs as illustrated in Section I. This paper closes in Section V with a short summary of the examination and conclusion.

II. LIST OF RECENT DELVEOPMENTS OF TAPE TECHNOLOGIES

In 2013 and 2014, at least eight new developments and products for the tape business were announced and available for the market. The memo of the TSC lists the recent developments in chronological order [5]:

- Sept. 16, 2013 Oracle announced a new enterprise tape drive StorageTek T10000D with 8,5TB native capacity and a data rate of 252 MB/s native.
- Jan. 16, 2014 Fujifilm Recording Media U.S.A. reported over 100 million manufactured LTO Ultrium data cartridges, equivalent to over 53 EB in data capacity.
- Apr. 30, 2014 Sony Corporation independently developed a soft magnetic under layer to successfully demonstrate the world’s highest areal recording density for tape storage media of 148 GB/inch².

- May 19, 2014 Fujifilm and IBM successfully demonstrated a record areal data density of 85.9 GB/inch² on linear magnetic particulate tape using the NANOCUBIC and Barium Ferrite (BaFe) particle technologies. This breakthrough equates to a standard LTO cartridge capable of storing up to 154 TB of uncompressed data, making it 62 times greater than today’s available LTO-6 cartridge capacity.
- Sept. 9, 2014, IBM announced Linear Tape File System (LTFS) Enterprise Edition (EE) Version 2.1.4.4 extending LTFS to the tape library support.
- Sept. 10, 2014, the LTO Consortium announced the extended roadmap with LTO-9 and LTO-10.
- Oct. 6, 2014, IBM announced the TS1150 enterprise drive with a data rate up to 360 MB/sec and a native cartridge capacity of 10 TB.
- Nov. 6, 2014 HP announced its new release of StoreOpen Automation that delivers a solution for LTFS in automation environments with Windows OS, available as free download.

This list includes an impressive number of vendors and suppliers that are committed to investing time and resources to tape technologies. The diversity of vendors and suppliers also suggests that the industry is not solely for a niche group of players but is developing into a storage market as a whole.

It should be widely known, that the following aspects of tape vs. disk apply today:

- Tape is cheaper to acquire,
- Tape is less costly to own and operate,
- Tape is more reliable,
- Tape now has media partitions for faster “disk-like” access,
- The capacity of a tape cartridge is higher than a disk drive’s capacity, and the media life for tape is 30 years or more for all new data [11].

Fig. 2 shows the areal density developments of disks and tape since the 1990s. While for disks the improvement was a 35% improvement from 2003 to 2009, the future outlook of that development is now slowing down, while the tape roadmap stays stable growing with a 33,15% per year until 2022.




Figure 2. Areal Density of Hard Disk and Tape –Laboratory Demonstrations and Products [10] [11]

III. ANALYSIS OF TAPE VS. DISK

As section II listed an overview of recent tape technology advancements, mainly covering the terms of capacity increase and of cost decline, this section brings the aspects of costs per terabyte, capacity and performance together. As already demonstrated, costs per terabyte in disk and tape will decrease over the next 10 years. While data volumes continue to grow, at least at the same or even higher speed, a combination of these two developments show, that in four years the break even for tape in terms of cost effectiveness per terabyte will be reached.

Fig. 3 provides the clear picture of this development, reached by 2019: the red trend line is the decrease in cost per terabyte for disks, the blue trend line is the decrease in cost per terabyte for tape, while the yellow trend line shows, as an example, the archived data growth from Oak Ridge National Laboratory.




Figure 3. Development of Technology Cost per Terabyte vs. Terabyte of data growth [9] [13]

Not just for high performance research organizations, such as Jülich Supercomputing Centre (JSC), the third aspect of performance is vital for IT-managers to define their storage strategy and layout. The Fig. 3 combines only cost per terabyte and data growth rate. However, the criteria of performance must be viewed as another third aspect of a tape vs. disk comparison.

A tape data rate using a 10 TB native capacity on the IBM TS1150 is up to 360 MB/s [5]. A standard commercially available hard disk drive of 4 TB SAS has a performance up to 6 gigabytes per second. Comparing the aspect of performance alone for tape vs. disk, the disk will always win this criteria. As a result, tape will need a “little helper”, realized by tape storage software, to perform much better than disk. As the access and interchange capabilities of tape have to be improved, Linear Tape File System (LTFS), a long awaited file system specification for LTO, is available since mid of 2010 [11].

LTFS provides a dual partitioning functionality, allowing the tape to be self-describing. Metadata operations, such as browse directory tree structures and file-name search, are

performed more quickly and do not require physical tape movement [11].

Using the Linear Tape File System, files can be created on tape and accessed similar to the process of creating and accessing files on an external hard drive or a USB flash drive. Applications, such as file browsers, image viewers and media players can directly browse and access files on tape. LTFS enables easy and simple use of tapes in desktop computers and embedded systems. One drive can read and write data at an impressive 1.2 Gigabytes per second, (approximately two times the speed of a hard drive), when using LTFS [12]. This is still not quite that fast than a single HDD, but tape comes close. The following section IV describes, how performance can be highly achieved by using tape and respective tape software technologies.

IV. BIG DATA STORAGE SOFTWARE

In Spring 2010, a couple of leading storage vendors and suppliers formed a collaborative industry alliance called the Active Archive Alliance in order to educate end user organizations on the evolving new technologies that enable reliable, online and efficient access to their archived data [14]. “Active Archive is a combined solution of open systems applications, disk, and tape hardware that allows users to access all of their data, and gives you an effortless solution that stores and manages all of your data.” [15] The National Center for Supercomputing Applications (NCSA) successfully deployed an active archive for the world’s largest active file repository in 2014. This active archive has a 99,99% availability, consists of 224 TS1140 tape drives at 240 MB/s each with a 52,7GB/s aggregate I/O throughput and an ability to grow up to 380 PB of data storage over the course of five years [16]. Data-intensive needs of scientists and engineers are driving this massive scale of a data archive in conjunction with the use of its Blue Waters high performance supercomputer at NCSA. The solution consists of a File-based active archive to enterprise tape, including the following hardware and software components:

- IBM HPSS software solution with DataDirect Networks disk cache
- Spectra Logic T-Finity enterprise tape library
- IBM TS1140 3592 JC enterprise data tape with 4.0 TB native capacity

“As the size of the archives increases, the tiering in these systems could be a big driver for a variety of digital storage technologies, from flash memory to hard disk drive, to magnetic tape or optical discs and even cloud storage (which can be any combination of these technologies)”. [18] A leading concept has been long-time established in the storage industry for this: it is called hierarchical storage management (HSM).

Due to lack of performing software components (or very custom designed software in combination with an proprietary file system at a high price or at high installation complexity levels), difficulties in ensuring data end-to-end integrity, high implementation efforts needed for a high availability deployment, and access to data with an adequate speed in

Gigabytes per second hindered IT Departments to introduce in their storage concepts also tape technologies as active part in their storage infrastructure.

Simpler HSM deployments in using a combination of higher valued fibre channel disks with lower commodity type of SATA disks and only using tape as the backup of this storage infrastructure is a very common deployment in most of organizations today. “Easy-Tiering” software technologies are in place to move within the disk pools older data to lesser expensive disk hardware or can be found in file systems supporting policies for hierarchical storage management (i.e., in IBM General Parallel File System (GPFS)). In periodically terms that data gets then stored onto tapes in order to be archived. When those data arrived at the tape pool, that data was very sporadic or only on special events accessed (i.e., huge data loss due to catastrophic damages, or for special audit purposes, etc.). A high velocity of that data access was not needed frequently or the data was not a very business critical item. This will change.

Even with the development of Flash technology arrived, this type of technology has not solved all performance problems for small block files [15]. For increased metadata performance in i.e., for database storage systems, Flash and solid-state-disks (SSDs) will play in future a vital role and will be commercially priced in an attractive way. But for massive data scale up to multiple PB and EB with the need to access that data in a faster way, the current standard HSM concepts will fail to provide a professional solution and/or surge higher costs due to higher software licenses fees while storage capacity increase. It is time that IT Department managers look at tape in a different way. For archival storage systems, Flash could improve metadata performance. It also has potential use as a low latency cache for user data to be used in the hierarchy of storage that most archival storage systems offer today with disk and tape alone [13].

IBM is the leader for deploying a proven highly scalable archive and Hierarchical Storage Management system with its HPSS offering. HPSS has been used successfully for very large digital image libraries, scientific data repositories, university mass storage systems, and weather forecasting systems, as well as defense and national security applications. Major compute-intensive and data-intensive sites such as Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), Brookhaven National Laboratory (BNL), Oakridge National Laboratory (ORNL), Argonne National Laboratory (ANL), the European Center for Medium-range Weather Forecasts (ECMWF), the German Climatic Research Center (DKRZ), the German Meteorological Service (DWD), and the National Oceanic and Atmospheric Administration (NOAA) have all chosen HPSS for their mission-critical data assets.

HPSS is a network-centered, cluster-based software offering that provides for stewardship and access of many petabytes of data. HPSS is capable of concurrently accessing hundreds of disk arrays and tape drives for extremely high aggregate data transfer rates, thus enabling HPSS to easily meet otherwise unachievable demands of total storage capacity, file sizes, data rates, and number of objects stored. In recent installations HPSS demonstrates an aggregated

throughput for read/write from clients from disk to the tape archive with up to 18Gigabytes per second [8]. HPSS is designed to scale horizontally and consists of a HPSS Core Server, that manages the metadata, and HPSS Data Movers that move the data between the disk cache and the tape pool. Fig. 4 shows the architecture of HPSS.

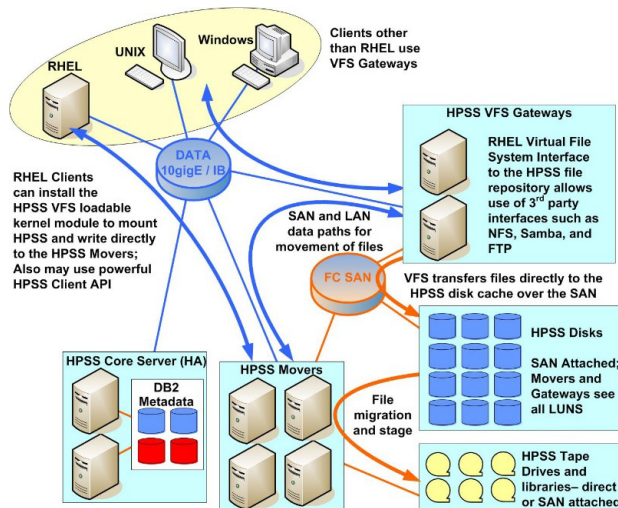


Figure 4. HPSS Architecture, IBM Corp.

In order to come to a start configuration for an HPSS installation, it is necessary to provide a “rule-of-thumb” on how to measure the demands for a HSM concept for an individual organization. The DOE Lab report states that there is a correlation of capacity with the amount of main memory of the server systems [13].

However, today business organizations that have increasing data growth needs, with smaller but high capable server environments, need to consider on what journey they want to continue with their storage environment if they reach certain storage capacity levels. Starting with an active data storage capacity of 3 PB and an expected exponentially data growth rate for the next 3 to 5 years, a professional HSM concept should be evaluated.

Fast and secure access to data is a key requirement in business environments today, as more data is being created, reused or reactivated for a different purpose. Those storage environments should be enabled with a flexible HSM concept in order to drive efficiencies. HPSS is being installed on commodity x86_64 node servers and with a standard storage disk system i.e., an IBM DS3860 or with a NetApp’s E-Series 5500 disk storage system or similar, with the operating system Red Hat Enterprise Linux.

Table I shows for HPSS performance values for a very small designed HSM configuration. It consists of one core server with optional manual cold-standby functionality, and two standard x86 Intel server as Data Movers. Three PB of data would be efficiently managed by this HSM configuration. The HPSS configuration would achieve already with this rather simple server hardware configuration the “standard” disk performance of today’s 4 TB HDDs.

TABLE I. AGGREGATE THROUGHPUT RATES OF HPSS, P. SCHAEFER, HPSS ENGINEER, IBM CORP.

Configuration	Disk*	Fibre*	10GigE network*	Tape system*
Small LTO/6	2458	1638	3414	896
Small TS1150	2458	1638	3414	1152
5Gbps LTO/6	6145	4095	9104	3072
5Gbps TS1150	6145	4095	9104	3168
* all values in Megabytes/second (Mb/s)				

As a start for a small HPSS, i.e., as a start basic configuration having one core server, one data mover and a disk cache with dual controllers of 844 TB of usable space (8+2p LUNs and 10% sparing), and for a HPSS system with a moderate 5Gigabyte per second, using one core server, two data movers and a disk cache with 985 TB of usable space (8+2p LUNs and 10% sparing), both disk cache each including the performance turbo option, achieve an reasonable aggregated throughput performance at lower costs.

V. CONCLUSION

There is no doubt, that data, especially unstructured data, will continuously grow, as individuals, organizations and researchers want more capacity for storage, access information almost everywhere, and conduct more valuable data experiments at any time with any interface. Tape is still being seen as a laggard in its development in the past five years. But in the past 18 months, tape technology vendors and developers have reached new levels of improvement, so that tape can be viewed for future storage system requirements differently.

From 2019 tape will play a greater role, as leading analysts, consultants and vendor predict. Fig. 3 clearly indicates the breakeven point, when to consider tape as an active storage component and not just as a backup device. Section III described the development in regard of cost per terabyte and terabyte of data growth.

For tape and performance, intelligent Big Data storage software is being used to improve the performance and speed of such tape systems. In addition to that, new data center service models are evolving, i.e., Cloud-based data centers. For that, the next stage is to enhance those cloud data centers operated by service providers with a “Big Data Cloud Storage” offering. This would be consisting of large tape environments and the use open source software such as OpenStack, to provide an easy to use enterprise-class cloud based storage service. This would be very suitable for very large archives (petabytes to exabytes) with occasional or significant retrieval needs, and it would offer very low cost storage together with affordable and predictable retrieval costs.

REFERENCES

[1] A. Dayley, G. Landers, A. Kros, and J. Zhang, “Magic Quadrant for Enterprise Information Archiving 2014”, Gartner, Inc., Nov. 2014.

[2] IBM Press Release, 2007, Report: IBM Passes All Major Vendors to Take #1 Position in Disk and Tape Storage Hardware Market Share Worldwide for 2006, URL: <http://www-03.ibm.com/press/us/en/pressrelease/21643.wss>, [accessed 2015-02-12].

[3] T. Pearson, 2007, “IBM passes all major vendors in storage”, URL: https://www.ibm.com/developerworks/community/blogs/InsideSystemStorage/entry/ibm_passes_all_major_vendors?lang=en, [accessed 2015-02-12].

[4] L. Dubois and A. Nadkarni, “Worldwide Storage in Big Data 2014 – 2018 Forecast”, IDC Opinion, Aug. 2014.

[5] Tape Storage Council, 2014, “2014 TSC Memo: Data Growth and Technology Innovation Fuels Tape’s Future”, URL: <http://tapestorage.org/wp-content/uploads/2014/12/2014-Tape-Storage-Council-Memo.pdf>, [accessed 2015-02-11].

[6] LTO@ Ultrium, (2014), LTO@ Program Further Extends Product Roadmap Through Generation 10, URL: <http://www.lto.org/2014/09/lto-program-further-extends-product-roadmap-through-generation-10/>, [accessed 2015-02-13].

[7] T. Olavsrud, 2012, “How to Implement Next-Generation Storage Infrastructure for Big Data”, URL: <http://www.computerworld.com/article/2503174/enterprise-applications/how-to-implement-next-generation-storage-infrastructure-for-big-data.html>, [accessed 2015-05-10].

[8] M. Boettinger, J. Cadmus, J. Meyer, and J. Nachtsheim, 2015, “German Researchers Use IBM Big Data Solution To Manage World’s Largest Trove Of Climate Data”, URL: <http://www.prnewswire.com/news-releases/german-researchers-use-ibm-big-data-solution-to-manage-worlds-largest-trove-of-climate-data-300080089.html>, [accessed 2015-05-08].

[9] D. David Floyer, 2014, “The Emergence of a New Architecture for Long-term Data Retention”, URL: http://wikibon.org/wiki/v/The_Emergence_of_a_New_Architecture_for_Long-term_Data_Retention, [accessed 2015-05-14].

[10] S. Weingand, 2014, “Tape ist lebendiger denn je”, URL: <http://de.slideshare.net/JosefWeingand/ibm-dpr-tape-ist-lebendiger-denn-je-0814>, [accessed 2015-05-13].

[11] F. Moore, 2014, “Tape. New Game. New Rules”, Horison Information Strategies, URL: <http://horison.com/wp-content/uploads/2014/10/Fujifilm-Tape-New-Game-New-Rules-21.pdf>, [accessed 2015-02-13].

[12] StorageDNA, 2012, Facts about LTO and LTFS, URL: http://www.storagedna.com/media/documents/StorageDNA_whitepaper_Facts-about-LTO-LTFS_web.pdf, [accessed 2015-05-14].

[13] National Energy Research Scientific Computing (NERSC) Facility, Extreme Scale Workshop, 2009, HPSS in the Extreme Scale Era, URL: <http://www.nersc.gov/assets/HPC-Requirements-for-science/HPSSExtremeScaleFINALpublic.pdf>, [accessed 2015-02-13].

[14] Active Archive Alliance, 2015, Introduction, URL: <http://www.activearchive.com/>, [accessed 2015-02-13].

[15] Spectra Logic, 2015, What is Active Archive?, URL: <http://www.spectralogic.com/index.cfm?fuseaction=solutions.showContentAndChildren&CatID=1690>, [accessed 2015-02-13].

[16] Active Archive Alliance Case Study, 2014, (NCSA) National Centers for Supercomputing Applications, NCSA Employs Active Archive for World’s Largest Active File Repository, URL: <http://www.activearchive.com/common/pdf/AA-Case-Study-NCSA.pdf>, [accessed 2015-02-13].

[17] T. Coughlin, 2014, Archiving Storage Tiers Part 2, URL: <http://www.forbes.com/sites/tomcoughlin/2014/12/30/archiving-storage-tiers-part-2/>, [accessed 2015-02-13].

[18] T. Coughlin, 2015, Flash Memory in Archives, URL: <http://www.forbes.com/sites/tomcoughlin/2015/01/31/flash-memory-in-archives/>, [accessed 2015-02-13].

An Adaptive Load-balancer for Task-scheduling in FastFlow

Md Moniruzzaman*, Kamran Idrees*, Michael Rossbory†, José Gracia*

*High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany

†Software Competence Center Hagenberg GmbH, Austria

moniruzzaman@hlrs.de, idrees@hlrs.de, Michael.Rossbory@scch.at, gracia@hlrs.de

Abstract—Balancing the computational load of multiple concurrent tasks on heterogeneous architectures is one of the critical requirements for efficient usage of such systems. Load-imbalance is inherently present if the computation load is distributed non-uniformly across various tasks or if execution time for the same kind of tasks varies from one class of processing element to the other. Load-imbalance may however also arise from causes that are beyond the control of the user, as for instance operating system jitter, over-subscription of the available workers, interference and resource contention by concurrent tasks, etc. Writing a balanced parallel application requires careful analysis of the problem and good understating of various hardware architectures of the computing nodes. FastFlow is a C++ library that offers high-level parallel pattern abstractions on the user side, and lowers those onto efficiently implemented architecture specific skeletons. The default FastFlow scheduler, however, assigns tasks to workers in a round-robin fashion and is thus not well suited to handle load-imbalance. In this paper, we present an adaptive load-balancing task scheduler for FastFlow, a model for the expected relative performance of our adaptive scheduler over the default round-robin scheduler, and finally evaluate the quality of the implementation with low-level as well as two specific application benchmarks. We find that the adaptive load-balancer does not introduce additional overheads if load-imbalance is not present, and that our scheme is particularly efficient in mitigating the effect of thread over-subscription. Finally, we show that the proposed scheduler can lead to substantial performance gain for real industrial applications.

Keywords—task scheduling; load balancing; heterogenous architecture; NUMA; FastFlow.

I. INTRODUCTION

The efficient utilisation of heterogeneous computing systems, in particular the balancing of application load on the available processing elements or workers, is a non-trivial tasks. One of the key factors for load-imbalance is the heterogeneity of processing elements architecture – a GPGPU will be faster than a CPU for certain workloads (and slower for others) [7]. Therefore, one has to consider the computing power at the time of assigning work to a particular worker. Ideally, the faster worker should not wait for the slower one in a heterogeneous hardware environment.

A further, often underestimated source for load-imbalance is variation of the execution time due to the very presence of other concurrent tasks as for instance operating system jitter, over-subscription of the available workers, interference and resource contention by concurrent tasks, etc. A particular example is inefficient access to memory and caches due to non-uniform memory architectures (NUMA), which are common

already on multi-core systems.

The most obvious source for load-imbalance, however, is that the task itself exhibits non-uniform task execution time. The execution time of the task is not always the same, and the exact time requirement by the task, even in the case of homogeneous hardware environment, is not easy to determine. For instance, root-finding algorithms to minimize complex mathematical functions will converge much faster if the initial guess is close to the final minimum, but might take much longer if the initial guess is in an off region of the parameter space. In this sense, the execution time of the tasks is non-uniformly distributed.

If we consider the block of code in Figure 1, one easily notices that the execution time of the task `aTask` varies depending upon which section of if-then-else clause is being executed. Now, the challenge is that we cannot determine the execution time of a task prior to its execution. However, load-imbalance would be obvious, if we blindly distribute a number of tasks equally to different workers without having prior knowledge of the required time to finish the task. Clearly, equal task distribution to all workers is not an optimal solution in this scenario because the non-uniform task execution time is caused by the task itself and the underlying hardware architectures as discussed earlier. This non-uniform execution behaviour is one of the primary reasons of load-imbalance and non-optimal solution [8]. In order to mitigate this situation, we need an adaptive load-balancing task scheduler.

In this paper, we propose a novel idea for an adaptive, load-balancing scheduler and realize it for FastFlow programming model [1]. The details of the approach are described in the Section II along with FastFlow’s existing task scheduler. Our approach is evaluated by a series of benchmarks in Section III. The benchmarks include not only synthetic ones, but also two real world applications, namely a Molecular Dynamic code [5] and a Material Optimization Process [10]. Finally, we draw our conclusion in Section IV.

```
def aTask(condition) {
  if (condition)
    longCalculation()
  else
    shortCalculation()
}
```

Figure 1. An example of non-uniform task




Figure 2. a FastFlow Farm pattern.

II. RELATED WORK AND OUR APPROACH

FastFlow is a C++ framework to write parallel programs [1]. The framework consists of parallel programming patterns such as farm, pipeline, map, etc, which are composed of various so-called nodes. The farm pattern for instance, is composed by an emitter node, several worker nodes, and a collector node as shown in the Figure 2. Note that in FastFlow, all nodes are implemented as individual OS threads. For instance in the farm pattern, each of the various worker nodes as well as the emitter and collector nodes, respectively, are separate threads. The nodes have input and output queues to connect to each other thus forming basic patterns. An emitter node, with help of an embedded load balancer, distributes tasks to the workers through their input queues. These workers, in principle, do the job concurrently, and the collector collects the results.

Task scheduling for parallel resources is an active area of research [13],[14]. However, the default scheduling policy of the emitter’s embedded load-balancer is round-robin. Therefore, we refer to this default scheduler as round-robin scheduler (RRS). As argued in the introduction, we believe that the round-robin scheduling policy is not suitable for non-uniform tasks executing environment. Therefore, we propose an new load-balancing scheme, namely adaptive round-robin scheduler (ARRS). The pseudo-code of the ARRS is shown in Figure 5; the basic working is explained in the following.

The basic idea of the ARRS is to respect the computing capability of workers, i.e., faster workers will receive more tasks than slower workers. We do so by keeping the height of the input queues roughly at the same level. Whenever, a task is ready for being emitted onto the workers, it is issued to the worker with the smallest input queue length, i.e., with the lowest queue level.

Determining the length of input queues in FastFlow, however, is a comparatively expensive operation. In fact, FastFlow aims to take scheduling decisions within a few clock cycles. Recording the queue level of all workers, would thus incur a prohibitively expensive overhead. Instead, we do not record queue levels at each task scheduling event, but only at one in `CHUNK_SIZE` events, thus reducing the overhead by the same factor.

In many applications, tasks are created in bursts, with relatively few tasks emitted between bursts. Scheduling all tasks, which are of unknown duration, could lead to a situation, where the queue levels are balanced in terms of number of tasks, but very unbalanced in terms of total duration




Figure 3. Queue length difference at a certain point.




Figure 4. ARRS tries to balance the queues of the workers.

of the tasks in the respective queues. In order to mitigate this situation, we use a throttling technique. Tasks are not necessarily scheduled as soon as they become available, instead the scheduler defers scheduling decisions until at least one of the input queues of the workers is nearly empty. Only then we proceed with scheduling a block of tasks before throttling again. For simplicity, the size of scheduling block is set to `CHUNK_SIZE`, i.e., the frequency of queue level recording.

As said earlier, continuously monitoring the length of the queues is prohibitively expensive. Therefore, we inject a so called *fake-task* into the queues of each worker towards the end of the block-wise scheduling episode. As soon as the first worker executes this fake-task, the throttling mechanism will start the next block-wise scheduling episode. Our scheduler is thus asynchronously driven by the length of the queues and does not require further explicit synchronisation. Note, that the queue of the fastest worker is not allowed to run fully empty. The fake-tasks are placed in the worker queues as soon as only a small number, `NEAR_END`, of tasks remain to be scheduled in the current scheduling episode. This parameter is set to a small multiple of the number of workers.

Within a scheduling episode, incoming tasks are scheduled onto the shorter queues in such a way that differences in queue lengths, which have been precomputed at the time of queue level recording, are leveled out, as illustrated in Figures 3 and 4. If the differences in queue length are smaller the `CHUNK_SIZE` to begin with, the queue lengths will balance well within the scheduling episode. The scheduler will then revert to a round-robin policy for the rest of the scheduling episode. Notably, ARRS will also revert to a round-robin policy at the very start of the application.

The proposed adaptive round-robin scheduler is a drop-in replacement for the default FastFlow round-robin scheduler and fully transparent to the user. ARRS is required to perform as well as RRS when there is no load-imbalance of tasks, i.e., the overhead on ARRS over RRS needs to be small. The efficiency of ARRS in the case of load-imbalance is determined mostly by the value of the parameter `CHUNK_SIZE`, i.e., the frequency of queue level recording and the duration over which it is assumed that relative queue levels do not change significantly. The other parameter `NEAR_END`, which gives the lead time of starting a scheduling episode over the emptying of the queues has been found to have a negligible effect as long as it is in the order of a few times the number of workers.

III. EVALUATION

The proposed load-balancer ARRS has been integrated into FastFlow version 2.0.0 [6]. The same version of FastFlow with RRS has been compared with the performance of ARRS.

```

#define CHUNK_SIZE 300
#define NEAR_END 4*NUMBER_OF_WORKERS
int n_scheduled = 0

def arrsScheduler(the_task) {
    if (n_scheduled < CHUNK_SIZE) {
        // hand-out tasks

        if (queuesUnbalanced()) {
            scheduleLowestQueue(the_task)
        } else {
            scheduleRoundRobin(the_task)
        }

        if (CHUNK_SIZE - n_scheduled == NEAR_END) {
            sendFakeTaskToAllWorkers()
        }

        n_scheduled++
        return true
    } else {
        // wait for queues to run empty
        waitForFakeTask()

        recordQueueLevels()
        n_scheduled = 0
        return false
    }
}
    
```

Figure 5. Pseudo code of the scheduling algorithm

During the experiments no other user process was running on the machine.

The evaluation steps are divided in to four categories. Firstly, we measured the overhead of the scheduler by creating a synthetic benchmark. Here, we want to see the possibility of using ARRS regardless of task type (uniform or non-uniform) as replacement for RRS. Secondly, we have defined a performance model for heterogenous hardware which has explicit non-uniform task execution behavior, and we compared the performance of the scheduler with the model by using a hardware simulation technique. Thirdly, we assessed the scheduler with implicit producer of load-imbalance such as NUMA and thread over-subscription. Finally, two real world applications have been exercised with the scheduler. During the experiment, the default `CHUNK_SIZE` was 300, if it is not mentioned explicitly.

Throughout the paper, we will discuss the overhead and benefit of ARRS over RRS in terms of the relative performance (improvement), S , for a given benchmark. The experimental relative performance is calculated as the ratio of average execution times using RRS, t_R , and ARRS, t_A , respectively, i.e.

$$S = \frac{\langle t_R \rangle}{\langle t_A \rangle} \quad (1)$$

In all benchmarks, the averages have been calculated over at least 10 runs. Error are estimated from the standard error of mean over the samples. Error bars are shown in plots only if they are relatively large.




Figure 6. Relative performance of ARRS over RRS with uniform tasks.

A. Evaluation environment

The experiments have been conducted on an Intel Nehalem machine. This machine has two NUMA domain each with 4 cores. However, since hyper-threading was enabled, operating systems see (8+8) 16 logical cores. Each core is equipped with 32KB L1 data cache and 256KB L2 cache. A 8MB L3 cache is shared by all cores on a socket. Each socket has 6GB memory. The two sockets are connected by a Quick Path Interconnect (QPI) bus controller.

B. Overhead

The purpose of this section is to measure the scheduling overhead by ARRS. To measure the overhead, a FastFlow application with uniform task is devised. The uniform task can be described as a task that requires, in principle, the same amount of time every time it is executed by a worker.

The FastFlow application consist of a farm, i.e., emitter, collector and workers. The experiment was exercised with 2, 4, 8 and 16 workers separately. The emitter of the farm distributes tasks to the workers through a scheduler (RRS or ARRS). The uniform task consists in calculating all prime numbers up to 75,00, which takes ~ 0.28 ms time to execute on this machine.

Figure 6 shows the relative performance of ARRS over RRS as a function of number of scheduler tasks for this FastFlow application running with 2, 4, 8 and 16 workers, respectively. If the value of the relative performance, S , is 1, both schedulers perform equally; if the value is higher than 1, ARRS performs better than RRS. The error estimate for the relative performance is shown at each point of the experiment with error bars. The relative performance is on average always larger than 1 for this benchmark. Therefore, we can conclude there is no significant overhead of ARRS compared to RRS for uniform task. In fact, the trend seems to indicate that ARRS performs slightly better than RRS. This is most likely due to the presence of OS jitter which causes slight imbalance of tasks [3], which do affect the execution time of with RRS, but are leveled out with ARRS.

C. Explicit load-imbalance of task execution

This section presents an idealized model for the relative performance of ARRS over RRS in case of heterogeneous architectures. The model is used to evaluate the efficiency of the implementation. Let us consider that number of tasks to schedule is n . All tasks are uniform, i.e., they take the same




Figure 7. Performance comparison in hardware simulation with 1 faster worker.

time to execute on given architecture. However, the execution time on different architectures may differ. Let p be the number of workers. All workers participate in the execution of tasks and do nothing else. We assume there is no overhead by the scheduler. Each worker i , with $1 \leq i \leq p$, shows relative speed s_i for the task to execute. The higher the value of s_i , the faster a given worker executes. Without loss of generality, assume worker $i = 1$ is the slowest worker.

The time to execute all n tasks by an ideal round-robin scheduler is $t_R \propto \max(\frac{n}{p s_i}) = \frac{n}{p s_1}$, where the maximum is taken over the group of workers, and the last equality is due to the assumption that worker $i = 1$ is the slowest one.

In the case of an ideal adaptive round-robin scheduler, each worker executes $n_i = n s_i / \sum_p s_i$ out of n total tasks. Here, $\sum_p s_i$ can be considered as aggregated computing power of the workers. If the work load is perfectly balanced, then the time to execute all n tasks is given by $t_A \propto n_i / s_i = n / \sum_p s_i$.

Finally, the expected potential relative performance, S , of the ideal adaptive round-robin scheduler over the non-loadbalancing round-robin scheduler is

$$S = t_R / t_A = \frac{\sum_p s_i}{p s_1} \quad (2)$$

The main difference between the assumed ideal adaptive round-robin scheduler and our ARRS is that the former takes scheduling decisions for each incoming tasks, while ARRS does so block-wise for `CHUNK_SIZE` tasks.

In the following, we simulate the effect of a heterogenous system consisting of two types of workers and contrast the experimental relative performance with the ideal model in (2). We simulate heterogeneity by consistently decreasing the execution time of tasks on exactly one specific worker. Specifically, the tasks on slower workers calculate all prime numbers up to 7500, the fast worker only up to 2000. In our system, the execution times are ~ 0.28 ms and ~ 0.048 ms, i.e., $s_1 = 1, s_f \sim 5.8$.

The result of the simulation in terms of measured relative performance (see (1)) as a function of total number of tasks is shown in Figure 7. The experiment was conducted with 4, 8 and 16 workers, respectively, out of which one is considered to be faster than the others. The plot also indicates the expected relative performance calculated from the model (2). The data demonstrates, that in general ARRS always schedules tasks




Figure 8. Performance comparison in hardware simulation with 1 faster worker.




Figure 9. Relative performance improvement by ARRS over RRS by mitigating NUMA effect.

more efficiently than RRS, as the relative performance is always larger than unity.

Notably, at low number of tasks the relative performance is rather low, but increases with increasing number of tasks as it approaches the expected theoretical value (2) at larger number of tasks. As anticipated the parameter `CHUNK_SIZE` has an impact of the efficiency of our (non-ideal) ARRS. Figure 8 shows the same benchmark as above with varying `CHUNK_SIZE`. The plot shows clearly that the theoretical curve is approached faster for smaller values of the parameter. Our implementation of ARRS, unlike the ideal theoretical one, schedules tasks in blocks. The first block, however, is scheduled in a round-robin fashion resulting in load-imbalanced queues on heterogenous systems. The imbalance needs to be corrected by successive scheduling event; however, the magnitude of the initial load-imbalance increases with on the block size and thus takes longer to correct for.

D. Implicit load-imbalance of task execution

There are many possible effects giving rise to an implicit imbalanced situational to a parallel application. In this section, we discuss 1) non-uniform memory architectures, 2) thread or worker over-subscription, as implicit cause of load-imbalance.

1) *NUMA effect*: In a NUMA architecture, threads running on cores which belong to one NUMA domain, suffer performance penalties if they access data which resides in another NUMA domain. This phenomenon may result in non-uniform task execution times [11], particularly if data access patterns do not allow to cache data efficiently.

We have implemented a benchmark consisting of only 2




Figure 10. Relative performance of ARRS over RRS with uniform tasks & thread over-subscription.

workers, for simplicity. Both workers are pinned to different NUMA domains using a tool called *likwid* [4]. Each time the task is invoked, it accesses and updates a different small region of memory in a cache in-efficient manner. All the data regions used in the benchmark have been initialized by only one of the workers and thus reside in its entirety in this worker's NUMA domain.

Figure 9 shows relative performance of ARRS over RRS for this NUMA scenario. The figure shows very clearly that NUMA effects have a significant impact on the performance, which is mitigated by the ARRS scheduler. In fact, one can estimate the ratio, $s_2/s_1 = 2S - 1 \approx 3.4$, of the execution time of a task on both workers, respectively, from the maximum of the relative performance, $S \approx 2.2$, taken from the figure. Again we observe a slow rise-up of the efficiency. Notably, for very large number of tasks, the relative performance goes down again. Here, the memory used by the benchmark is so large, that it no longer fits into a single NUMA domain and is thus distributed among both NUMA domains. This results in overall less load-imbalance as both workers – not only one as before – have to access data from remote domains.

2) *Thread over-subscription*: Typical FastFlow applications consist of dozens of FastFlow nodes organized in (possibly nested) patterns as farm, pipeline, map, etc. Each of these nodes is mapped to its own separate thread. Thus in general, the number of threads of a FastFlow application is larger than the number of processing elements, say cores, on the system it executes on. Moreover, all of these threads compete for their share of on-core-time at the same time. We refer to this scenario as thread or worker over-subscription. The situation is worsened by the fact that FastFlow workers are usually pinned to specific processing elements.

We have exercised the same uniform-tasks benchmark code as in Section III-B for 20 workers. In our system, 8 of these workers share a core, while the remaining 12 are assigned to a core exclusively. Figure 10 shows the experimental result of this benchmark. Again, ARRS mitigates the effect of worker over-subscription although the variation of the benchmark's execution time is relatively large (as evident from the error bars), possibly due to the larger impact of OS jitter and other factors when the system is under such increased stress.

E. Real world applications

1) *Molecular Dynamics*: We have tested ARRS with a Molecular Dynamics simulation code called *CMD* [5]. Molec-




Figure 11. CMD Speedup (with ARRS over RRS) versus number of task/molecules.




Figure 12. CMD Speedup (with ARRS over RRS) versus number of workers.

ular Dynamics (MD) is a simulation methodology used for modeling of molecules in a large range of scientific fields as material science, chemistry, theoretical physics, and biology. The code spends most of its time in the calculation of the force acting between any pairs of molecules or atoms [5]. This part was parallelised with a FastFlow farm pattern [12] in such a way, that a task consists of the force calculation for single pair of molecules. The tasks takes the same (very short, ≈ 8 ms for 68000 pair of molecules using 2 workers) amount of time for each pair [9] and is thus considered uniform in the sense used in this paper. We therefore do not expect any performance improvement, but would rather aim to assert that the overhead of ARRS is negligible also for this real application exhibiting very short task duration.

We have run CMD on 8 workers with realistic parameters similar to those that would be used in smallish MD production runs. Figure 11 indeed shows that ARRS does not introduce any measurable overhead over RRS. In fact, it performs consistently better than RRS by a few percent.

While CMD, due to its simple structure, would in general not be used under conditions of thread over-subscription, we have nonetheless benchmarked it. Figure 12 illustrates the relative performance for number of workers from 2 to 20. The relative performance of ARRS over RRS increases slightly with the number of workers up to 16. This increase likely stems from load-imbalance originating in the higher likelihood with increasing thread count of adverse thread-synchronization. At 20 workers, the application is load-imbalanced as some of the CPU cores need to serve two workers, while other only one. This scenario is much more efficiently handled by ARRS than by RRS.




Figure 13. Material Optimization Process speed-up by ARRS over RRS.

2) *Material Optimization Process*: The purpose of this application is to minimize the consumption of raw material in an industrial production process, and thus, reduce the production cost. The application is a complex mathematical optimization process subject to multiple constraints. In essence, it maintains a pool of given size of acceptable mathematical solutions. Solutions are taken from the pool and passed on to one of various different pipelines for further optimization. Each of these optimization pipelines follows different strategies, all of which take different time to finish. Solutions which potentially improve on the optimization goal are returned to the pool for further processing.

This application has been parallelized using a FastFlow farm. Each of the farm's workers follows a specific optimization strategy resulting in varying execution time across workers. The farm's emitter issues solutions from the task pool to the workers, while the collector replenishes the pool with promising solutions. The application continues until a desired optimal solution is found, the pool is empty, or another termination criterion is triggered [10]. In this case the application, running with 2 workers to compute 1000 task, needs ≈ 70917 ms with RRS whereas ≈ 39865 ms with ARRS. Figure 13 shows the relative performance of ARRS over RRS as a function of the size of the solution pool, which is a parameter that needs to be adapted by the user of the application to find the best trade-off between long application runtime and quality of the material optimization. While the pool size is not directly proportional to the total number of tasks executed by the application, in general the latter tends to increase with the pool size. Clearly, ARRS outperforms RRS for all values of the pool size. For realistic pool sizes ARRS can be more than three times faster than RRS if two workers are used.

IV. CONCLUSION

In this paper, we have presented an adaptive load-balancing task-scheduler, ARRS, and applied it to the FastFlow pattern-based parallel programming model. We have also presented an idealized theoretical model for its performance. We have benchmarked ARRS with the aim to assert 1) the overheads under conditions where it will not shown performance improvements by design, and 2) the efficiency of our implementation in terms of the idealized model.

Our benchmarks support the conclusion that ARRS overheads are negligibly small and that our implementation can

reach the theoretically expected efficiency. However, the efficiency of our implementation depends on the number of tasks that are scheduled; it approaches the ideal value asymptotically as the number of tasks increases. Careful analysis of our implementation reveals that the cause for this behaviour is an initial phase – whose length is given by the block size parameter – of round-robin scheduling of tasks. This phase leads to initial load-imbalances that need to be balanced out over successive scheduling periods.

Nonetheless, our real world application benchmarks show that the load-balancing scheduler leads to substantial performance increases in all cases considered. For an industrial relevant Material Optimization Problem, ARRS is up to three times faster than the default scheduler. The proposed load-balancer is well capable of dealing with non-uniformity of tasks and heterogeneity of hardware. We also show that ARRS may mitigate the performance impact of NUMA effects, load-imbalances due to OS jitter, etc., and thread over-subscription. We claim that ARRS is drop-in replacement for FastFlow's default scheduler that increases the speed of most applications, particularly on heterogenous systems, at no additional development cost for the FastFlow user.

Future work will address the short-comings of our approach – which are relevant only at (possibly intermittent) low number of tasks – by shortening the initial round-robin phase and by adapting the scheduler's block size to the number of task that are ready for scheduling.

ACKNOWLEDGEMENT

This work has been supported by the European Commission Seventh Framework Programme through the projects *ParaPhrase: Parallel Patterns for Adaptive Heterogeneous Multicore Systems* (contract no.: 288570) and *Polca: Programming Large Scale Heterogeneous Infrastructures* (contract no.: 619686), and by the German Research Foundation (DFG) through Priority Programme 1648 *Software for Exascale Computing (SPPEXA)*.

REFERENCES

- [1] Aldinucci, Marco, Massimo Torquati, and Massimiliano Meneghin. "FastFlow: Efficient parallel streaming applications on multi-core." arXiv preprint arXiv:0909.1187 (2009).
- [2] Torquati, Massimo. "Single-producer/single-consumer queues on shared cache multi-core systems." arXiv preprint arXiv:1012.1824 (2010).
- [3] Vicente, Elder, and R. Matias. "Exploratory study on the linux os jitter." In Computing System Engineering (SBESC), 2012 Brazilian Symposium on, pp. 19-24. IEEE, 2012.
- [4] Treibig, Jan, Georg Hager, and Gerhard Wellein. "LIKWID: Lightweight Performance Tools." In Competence in High Performance Computing 2010, pp. 165-175. Springer Berlin Heidelberg, 2012.
- [5] Idrees, Kamran, Mathias Nachtmann, and Colin W. Glass. "Evaluation of FastFlow Technology for Real-World Application." In Sustained Simulation Performance 2013, pp. 77-88. Springer International Publishing, 2013.
- [6] FastFlow Source: <http://sourceforge.net/projects/mc-fastflow/files/fastflow-2.0.0.tar.bz2/download> 2015.04.20
- [7] Lee, Victor W., Changkyu Kim, Jatin Chhugani, Michael Deisher, Daehyun Kim, Anthony D. Nguyen, Nadathur Satish et al. "Debunking the 100X GPU vs. CPU myth: an evaluation of throughput computing on CPU and GPU." In ACM SIGARCH Computer Architecture News, vol. 38, no. 3, pp. 451-460. ACM, 2010.

- [8] Goli, Mehdi, John McCall, Christopher Brown, Vladimir Janjic, and Kevin Hammond. "Mapping parallel programs to heterogeneous CPU/GPU architectures using a monte carlo tree search." In Evolutionary Computation (CEC), 2013 IEEE Congress on, pp. 2932-2939. IEEE, 2013.
- [9] Niethammer, Christoph, Colin W. Glass, and José Gracia. "Avoiding serialization effects in data/dependency aware task parallel algorithms for spatial decomposition." In Parallel and Distributed Processing with Applications (ISPA), 2012 IEEE 10th International Symposium on, pp. 743-748. IEEE, 2012.
- [10] Rossbory, Michael, and Werner Reisner. "Parallelization of Algorithms for Linear Discrete Optimization Using ParaPhrase." In DEXA Workshops, pp. 241-245. 2013.
- [11] Thomadakis, Michael E. "The architecture of the Nehalem processor and Nehalem-EP SMP platforms." Resource 3 (2011): 2.
- [12] Brown, Christopher, Vladimir Janjic, Kevin Hammond, Holger Schoner, Kamran Idrees, and Colin Glass. "Agricultural reform: more efficient farming using advanced parallel refactoring tools." In Parallel, Distributed and Network-Based Processing (PDP), 2014 22nd Euromicro International Conference on, pp. 36-43. IEEE, 2014.
- [13] Thoman, Peter, Herbert Jordan, Simone Pellegrini, and Thomas Fahringer. "Automatic OpenMP loop scheduling: a combined compiler and runtime approach." In OpenMP in a Heterogeneous World, pp. 88-101. Springer Berlin Heidelberg, 2012.
- [14] Wang, Zheng, and Michael FP O'Boyle. "Mapping parallelism to multi-cores: a machine learning based approach." In ACM Sigplan Notices, vol. 44, no. 4, pp. 75-84. ACM, 2009.

Adaption of the n-way Dissemination Algorithm for GASPI Split-Phase Allreduce

Vanessa End
and Ramin Yahyapour

Gesellschaft für wissenschaftliche
Datenverarbeitung mbH Göttingen
Göttingen, Germany
Email: <vanessa.end>,
<ramin.yahyapour>@gwdg.de

Christian Simmendinger
and Thomas Alrutz

T-Systems Solutions for Research GmbH
Stuttgart/Göttingen, Germany
Email: <christian.simmendinger>,
<thomas.alrutz>@t-systems-sfr.com

Abstract—This paper presents an adaption of the n-way dissemination algorithm, such that it can be used for an allreduce operation, which is - together with the barrier operation - one of the most time consuming collective communication routines available in most parallel communication interfaces and libraries. Thus, a fast underlying algorithm with few communication rounds is needed. The dissemination algorithm is such an algorithm and already used for a variety of barrier implementations due to its speed. Yet, this algorithm is also interesting for the split-phase allreduce operations, as defined in the Global Address Space Programming Interface (GASPI) specification, due to its small number of communication rounds. Even though it is a butterfly-like algorithm, significant improvements in runtime are seen when comparing this implementation on top of *ibverbs* to different message-passing interface (MPI) implementations, which are the de facto standard for distributed memory computing.

Keywords—GASPI; Allreduce; Partitioned Global Address Space (PGAS); Collective Communication.

I. INTRODUCTION

A main aspect in distributed application programming is the communication of data. An emerging communication interface designed for high performance computing (HPC) applications is the Global Address Space Programming Interface (GASPI) specification [1]. It is based on one-sided communication semantics, distinguishing it from message-passing paradigms, libraries and application programming interfaces (API) like the Message-Passing Interface (MPI) standard [2]. In the spirit of hybrid programming (e.g., combined MPI and OpenMP communication) for improved performance, GASPI's communication routines are designed for inter-node communication and leaves it to the programmer to include another communication interface for intra-node, i.e., shared-memory communication. Thus, one GASPI process is started per node or cache coherent non-uniform memory access (ccNUMA) socket.

To enable the programmer to design a fault-tolerant application and to achieve perfect overlap of communication and computation, GASPI's non-local operations are equipped with a timeout mechanism. By either using one of the predefined constants `GASPI_BLOCK` or `GASPI_TEST` or by giving a user-defined timeout value, non-local routines can either be called in a blocking or a non-blocking manner. In the same way, GASPI also defines split-phase collective communication routines, namely `gaspi_barrier`, `gaspi_allreduce`

and `gaspi_allreduce_user`, for which the user can define a personal reduce routine. The goal of our research is to find a fast algorithm for the allreduce operation which has a small number of communication rounds and whenever possible uses all available resources for the computation of the partial results computed in each communication round.

Collective communication is an important issue in high performance computing and thus research on algorithms for the different collective communication routines has been pursued in the last decades. In the area of the allreduce operation, influences from all other communication algorithms can be used, e.g., tree algorithms like the binomial spanning tree [3] or the tree algorithm of Mellor-Crummey and Scott [4]. These are then used to first reduce and then broadcast the data. Also, more barrier related algorithms like the butterfly barrier of Brooks [5] or the tournament algorithm described by Debra Hensgen et al. in the same paper as the dissemination algorithm [6] influence allreduce algorithms.

Yet, none of these algorithms seems fit for the challenges of split-phase remote direct memory access (RDMA) allreduce, with potentially computation-intense user-defined reduce operations over an InfiniBand network. The tree algorithms have a tree depth of $\lceil \log_2(P) \rceil$ and have to be run through twice, leading to a total of $2\lceil \log_2(P) \rceil$ communication rounds. In each of these rounds, a large part of the participating ranks remain idling, while the *n*-way dissemination algorithm only needs $\lceil \log_{n+1}(P) \rceil$ communication rounds and involves all ranks in every round. Also, the butterfly barrier has $k = \lceil \log_2(P) \rceil$ communication rounds to do besides only being fit for 2^k participants.

One of the most heavily used algorithms for barrier operations is the dissemination algorithm presented by Hensgen et al. in 1988 [6]. It is used in different programming APIs and libraries like the MPICH implementation [7] of MPI, the global address space networking API (GASNet) [8] and the second generation of the Global address Programming Interface (GPI-2) v1.1.0 [9] for barrier implementations due to its speed. In 2006 Torsten Hoefler et al. [10] have based the *n*-way dissemination algorithm on this work to exploit implicit parallelism of the InfiniBand network.

However, until today the dissemination algorithm and the developments have not been used for allreduce operations, partly due to the problems arising when using the dissemination algorithm with a number of nodes $P \neq 2^k$. The same

problem arises for the n -way dissemination algorithm if the number of involved nodes does not equal $(n + 1)^k$. This will be elaborated in more detail below.

There are two key features which make (n -way) dissemination based allreduce operations very interesting for both split-phase implementations as well as user-defined reductions, like they are both defined in the GASPI specification [1].

- 1) Split-phase collectives either require an external active progress component or alternatively progress has to be achieved through suitable calls from the calling processes. Since the underlying algorithm for the split-phase collectives is unknown to the enduser, all participating processes have to repeatedly call the collective several times. Algorithms for split-phase collectives hence ideally both involve all processes in every communication step and moreover ideally require a minimum number of steps (and thus a minimum number of calls). The n -way dissemination algorithm exactly matches these requirements. It requires a very small number of communication rounds of order $\lceil \log_{n+1}(P) \rceil$ and additionally involves every process in all communication rounds.
- 2) User-defined collectives share some of the above requirements in the sense that CPU-expensive local reductions ideally should leverage every calling CPU in each round and ideally would require a minimum number of communication rounds (and hence a minimum number of expensive local reductions).

In the following, the reasons why the n -way dissemination algorithm is not suitable for the allreduce collective operation are described. The cases in which the algorithm may be used as is are identified and also those cases, in which the algorithm must be adapted to receive correct results are described. Based on the data movement within the algorithm *data boundaries* can be identified with which the algorithm can then be adapted. Through this adaption, the amount of data to be transferred in the last round is reduced, in some cases it is even possible to omit the last round. Even though the shown adaption of the algorithm was done for allreduce operations performed with non-idempotent functions, the benefits regarding runtime are also measurable when used for, e.g., a barrier.

Some more related work will be presented in the next section. The n -way dissemination algorithm is introduced in Section III and then, in Section III-A, the reasons necessitating the adaption are described. An adaption to the n -way dissemination algorithm is proposed in Section IV to resolve these problems. First results, achieved with an implementation of the adapted algorithm on top of *ibverbs*, are presented in Section V. The conclusion and an outlook of future work are given in Section VI.

II. RELATED WORK

Some related work, especially in terms of developed algorithms, has already been presented in the introduction. Still to mention is the group around Jehoshua Bruck, which has done much research on multi-port algorithms, hereby developing a k -port algorithm with almost the same communication scheme as the n -way dissemination algorithm has [11][12]. These works were found relatively late in the implementation phase, such that an extensive comparison to this work has not been

made yet. But considering the results achieved with this work, also Bruck's algorithm is an interesting candidate for the GASPI allreduce.

In the past years, more and more emphasis has been put on RDMA techniques and algorithms [13][14] due to hardware development, e.g., InfiniBandTM[15] or RDMA over Converged Ethernet (RoCE). While Panda et al. [13] exploit the multicast feature of InfiniBandTM, this is not an option for us because the multicast is a so called unreliable operation and in addition an optional feature of the InfiniBandTM architecture [15]. Congestion in fat tree configured networks is still a topic in research, where for example Zahavi is an active researcher [16]. While a change of the routing tables or routing algorithm is often not an option for application programmers, the adaption of node orders within the API is a possible option.

III. THE n -WAY DISSEMINATION ALGORITHM

The n -way dissemination algorithm, as presented in [10] has been developed for spreading data among the participants, where n is the number of messages transferred in each communication round. As the algorithm is not exclusive to nodes, cores, processes or threads, the term *ranks* will be used in the following. The P participants in the collective operation are numbered consecutively from $0, \dots, P-1$ and this number is their rank. With respect to rank p , the ranks $p+1$ and $p-1$ are called p 's *neighbors*, where $p-1$ will be the *left-hand neighbor*.

Let P be the number of ranks involved in the collective communication. Then $k = \lceil \log_{n+1}(P) \rceil$ is the number of communication rounds the n -way dissemination algorithm needs to traverse, before all ranks have all information. In every communication round $l \in \{1, \dots, k\}$, every process p has n peers $s_{l,i}$, to which it transfers data and also n peers $r_{l,j}$, from which it receives data:

$$\begin{aligned} s_{l,i} &= p + i \cdot (n + 1)^{l-1} \pmod{P} \\ r_{l,j} &= p - j \cdot (n + 1)^{l-1} \pmod{P}, \end{aligned} \quad (1)$$

with $i, j \in \{1, \dots, n\}$. Thus, in every round p gets (additional) information from $n(n + 1)^{l-1}$ participating ranks - either directly or through the information obtained by the sending ranks in the preceding rounds.

A. Using the n -way Dissemination Algorithm for Allreduce

When using the dissemination algorithm for an allreduce, the information received in every round is the partial result the sending rank has computed in the round before. The receiving rank then computes a new local partial result from the received data and the local partial result already at hand.

Let S_l^p be the partial result of rank p in round l , \circ be the reduction operation used and x_p be the rank's initial data. Then rank p receives n partial results $S_{l-1}^{r_{l,i}}$ in round l and computes

$$S_l^p = S_{l-1}^p \circ S_{l-1}^{r_{l,1}} \circ S_{l-1}^{r_{l,2}} \circ \dots \circ S_{l-1}^{r_{l,n}}, \quad (2)$$

which it transfers to its peers $s_{l+1,i}$ in the next round. This data movement is shown in Table I for an allreduce based on a 2-way dissemination algorithm. First for 9 ranks, then for 8 participating ranks. By expanding the result of rank 0 in round 2 from the second table, it becomes visible, that the reduction operation has been applied twice to x_0 :

$$S_2^0 = (x_0 \circ x_7 \circ x_6) \circ (x_5 \circ x_4 \circ x_3) \circ (x_2 \circ x_1 \circ x_0). \quad (3)$$

TABLE I. ROUND-WISE COMPUTATION OF PARTIAL RESULTS IN A 2-WAY DISSEMINATION ALGORITHM

rank	round 0	round 1	round 2
0	x_0	$S_1^0 = x_0 \circ x_8 \circ x_7$	$S_2^0 = S_1^0 \circ S_6^1 \circ S_1^3$
1	x_1	$S_1^1 = x_1 \circ x_0 \circ x_8$	$S_2^1 = S_1^1 \circ S_7^1 \circ S_1^4$
2	x_2	$S_1^2 = x_2 \circ x_1 \circ x_0$	$S_2^2 = S_1^2 \circ S_1^8 \circ S_1^5$
3	x_3	$S_1^3 = x_3 \circ x_2 \circ x_1$	$S_2^3 = S_1^3 \circ S_1^0 \circ S_1^6$
4	x_4	$S_1^4 = x_4 \circ x_3 \circ x_2$	$S_2^4 = S_1^4 \circ S_1^1 \circ S_1^7$
5	x_5	$S_1^5 = x_5 \circ x_4 \circ x_3$	$S_2^5 = S_1^5 \circ S_1^2 \circ S_1^8$
6	x_6	$S_1^6 = x_6 \circ x_5 \circ x_4$	$S_2^6 = S_1^6 \circ S_1^3 \circ S_1^0$
7	x_7	$S_1^7 = x_7 \circ x_6 \circ x_5$	$S_2^7 = S_1^7 \circ S_1^4 \circ S_1^1$
8	x_8	$S_1^8 = x_8 \circ x_7 \circ x_6$	$S_2^8 = S_1^8 \circ S_1^5 \circ S_1^2$

rank	round 0	round 1	round 2
0	x_0	$S_1^0 = x_0 \circ x_7 \circ x_6$	$S_2^0 = S_1^0 \circ S_5^1 \circ S_1^2$
1	x_1	$S_1^1 = x_1 \circ x_0 \circ x_7$	$S_2^1 = S_1^1 \circ S_1^6 \circ S_1^3$
2	x_2	$S_1^2 = x_2 \circ x_1 \circ x_0$	$S_2^2 = S_1^2 \circ S_1^7 \circ S_1^4$
3	x_3	$S_1^3 = x_3 \circ x_2 \circ x_1$	$S_2^3 = S_1^3 \circ S_1^0 \circ S_1^5$
4	x_4	$S_1^4 = x_4 \circ x_3 \circ x_2$	$S_2^4 = S_1^4 \circ S_1^1 \circ S_1^6$
5	x_5	$S_1^5 = x_5 \circ x_4 \circ x_3$	$S_2^5 = S_1^5 \circ S_1^2 \circ S_1^7$
6	x_6	$S_1^6 = x_6 \circ x_5 \circ x_4$	$S_2^6 = S_1^6 \circ S_1^3 \circ S_1^0$
7	x_7	$S_1^7 = x_7 \circ x_6 \circ x_5$	$S_2^7 = S_1^7 \circ S_1^4 \circ S_1^1$

In general, if $P \neq (n+1)^k$, the final result will include data of at least one rank twice: In every communication round l , each rank receives n partial results each of which is the reduction of the initial data of its $(n+1)^{l-1}$ left-hand neighbors. Thus, the number of included initial data elements is described through

$$\sum_{i=1}^l n(n+1)^{i-1} + 1 = (n+1)^l \quad (4)$$


for every round l .

In the cases of the maximum or minimum operation to be performed in the allreduce, this does not matter. In the case of a summation though, this dilemma will result into different final sums on the participating ranks. In general, the adaption is needed for all operations, where the repeated application of the function to the same element changes the final result, so called *non-idempotent* functions.

IV. ADAPTING THE n -WAY DISSEMINATION ALGORITHM

The adaption of the n -way dissemination algorithm is mainly based on these two properties: (1) in every round l , p receives n new partial results. (2) These partial results are the result of the combination of the data of the next $\sum_{i=0}^{l-1} n(n+1)^{i-1} + 1$ left-hand neighbors of the sender. This is depicted in Figure 1 through boxes. Highlighted in green are those ranks, whose data view is represented, that is rank 0's in the first row and rank 2's in the second row. Each box encloses those ranks, whose initial data is included in the partial result the right most rank in the box has transferred in a given round. This means for rank 0, it has its own data, received S_0^6 and S_0^7 in the first round (gray boxes) and will receive S_1^5 and S_1^1 from ranks 2 and 5 in round 2 (white boxes).

As each of the boxes describes one of the partial results received, the included initial data items can not be retrieved by the destination rank. The change from one box to the next is


 Figure 1. The data boundaries g and received partial results $S_i^{r,l,j}$ of ranks 0 and 2.

thus defined as a *data boundary*. The main idea of the adaption is to find data boundaries in the data of the source ranks in the last round, which coincide with data boundaries in the destination rank's data. When such a correspondence is found, the data sent in the last round is reduced accordingly. To be able to do so, it is necessary to describe these boundaries in a mathematical manner. Considering the data elements included in each partial result received, the data boundaries of the receiver p can be described as:

$$g_{l_{rcv}}[j_{rcv}] = p - n \sum_{i=0}^{l_{rcv}-2} (n+1)^i - j_{rcv}(n+1)^{l_{rcv}-1} \bmod P, \quad (5)$$

where $j_{rcv}(n+1)^{l_{rcv}-1}$ describes the boundary created through the data transferred by rank $r_{l_{rcv},j_{rcv}}$ in round l_{rcv} .

Also, the sending ranks have received partial results in the preceding rounds, which are marked through corresponding boundaries. From the view of rank p in the last round k , these boundaries are then described through

$$g_{l_{snd}}^s[j_{snd}] = p - s(n+1)^{k-1} - n \sum_{i=0}^{l_{snd}-2} (n+1)^i - j_{snd}(n+1)^{l_{snd}-1} \bmod P, \quad (6)$$

with $s \in \{1, \dots, n\}$ distinguishing the n senders and j_{snd}, l_{snd} corresponding to the above j_{rcv}, l_{rcv} for the sending rank. To also consider those cases, where only the initial data of the sending or the receiving rank is included more than once in the final result, we let $l_{snd}, l_{rcv} \in \{0, \dots, k-1\}$ and introduce an additional *base border* g_B in the destination rank's data.

These boundaries are also depicted in Figure 1 for the previously given example of a 2-way dissemination algorithm with 8 ranks. The figure depicts the data present on ranks 0 and 2 after the first communication round in the gray boxes with according boundaries $g_B, g_0, g_1[1]$ and $g_1[2]$ on rank 0 and $g_0^2, g_1^2[1]$ and $g_2^2[1]$ on rank 2. Since the boundaries g_B and $g_1^2[1]$ coincide, the first sender in the last round, that is rank 5, transfers its partial result but rank 2 only transfers a reduction $S' = x_2 \circ x_1$ instead of $x_2 \circ x_1 \circ x_0$.

More generally speaking, the algorithm is adaptable, if there are boundaries on the source rank that coincide with

boundaries on the destination rank, i.e.,

$$g_{l_{\text{snd}}}^s [j_{\text{snd}}] = g_{l_{\text{rcv}}} [j_{\text{rcv}}] \quad (7)$$

or $g_{l_{\text{snd}}}^s [j_{\text{snd}}] = g_B$. Then the last source rank, defined through s , transfers only the data up to the given boundary and the receiving rank takes the partial result up to its given boundary out of the final result. Taking out the partial result in this context means: if the given operation has an inverse \circ^{-1} , apply this to the final result and the partial result defined through $g_{l_{\text{rcv}}} [j_{\text{rcv}}]$. If the operation does not have an inverse, recalculate the final result, hereby omitting the partial result defined through $g_{l_{\text{rcv}}} [j_{\text{rcv}}]$. Since this boundary is known from the very beginning, it is possible to store this partial result in the round it is created, thus saving additional computation time at the end.

From this, one can directly deduce the number of participating ranks P , for which the n -way dissemination algorithm is adaptable in this manner:

$$\begin{aligned} P &= g_{l_{\text{snd}}}^s [j_{\text{snd}}] - g_{l_{\text{rcv}}} [j_{\text{rcv}}] \\ &= s(n+1)^{k-1} + n \sum_{i=0}^{l_{\text{snd}}-2} (n+1)^i + j_{\text{snd}}(n+1)^{l_{\text{snd}}-1} \\ &\quad - n \sum_{i=0}^{l_{\text{rcv}}-2} (n+1)^i - j_{\text{rcv}}(n+1)^{l_{\text{rcv}}-1}. \end{aligned} \quad (8)$$

For given P , a 5-tuple $(s, l_{\text{snd}}, l_{\text{rcv}}, j_{\text{snd}}, j_{\text{rcv}})$ can be precalculated for different n . Then this 5-tuple also describes the adaption of the algorithm:

Theorem 1: Given the 5-tuple $(s, l_{\text{snd}}, l_{\text{rcv}}, j_{\text{snd}}, j_{\text{rcv}})$, the last round of the n -way dissemination algorithm is adapted through one of the following cases:

- 1) $l_{\text{rcv}}, l_{\text{snd}} > 0$
The sender $p - s(n+1)^{k-1}$ sends its partial result up to $g_{l_{\text{snd}}}^s [j_{\text{snd}}]$ and the receiver takes out its partial result up to the boundary $g_{l_{\text{rcv}}} [j_{\text{rcv}}]$.
- 2) $l_{\text{rcv}} > 0, l_{\text{snd}} = 0$
The sender $p - s(n+1)^{k-1}$ sends its own data and the receiver takes out its partial result up to the boundary $g_{l_{\text{rcv}}} [j_{\text{rcv}}]$.
- 3) $l_{\text{rcv}} = 0, l_{\text{snd}} = 0$
The sender $p - (s-1)(n+1)^{k-1}$ sends its last calculated partial result. If $s = 1$ the algorithm ends after $k-1$ rounds.
- 4) $l_{\text{rcv}} = 0, l_{\text{snd}} = 1$
The sender $p - s(n+1)^{k-1}$ sends its partial result up to $g_{l_{\text{snd}}}^s [j_{\text{snd}} - 1]$. If $j_{\text{snd}} = 1$, the sender only sends its initial data.
- 5) $l_{\text{rcv}} = 0, l_{\text{snd}} > 1$
The sender $p - s(n+1)^{k-1}$ sends its partial result up to $g_{l_{\text{snd}}}^s [j_{\text{snd}}]$ and the receiver takes out its initial data from the final result.

Proof: We show the correctness of the above theorem by using that at the end each process will have to calculate the final result from P different data elements. We therefore look at (8) and how the given 5-tuple changes the terms of relevance. We will again need the fact, that the received partial results are always a composition of the initial data of neighboring elements.

1) $l_{\text{rcv}}, l_{\text{snd}} > 0$:

$$\begin{aligned} P &= s(n+1)^{k-1} + n \sum_{i=0}^{l_{\text{snd}}-2} (n+1)^i + j_{\text{snd}}(n+1)^{l_{\text{snd}}-1} \\ &\quad - n \sum_{i=0}^{l_{\text{rcv}}-2} (n+1)^i - j_{\text{rcv}}(n+1)^{l_{\text{rcv}}-1} \\ &= g_{l_{\text{snd}}}^s [j_{\text{snd}}] - g_{l_{\text{rcv}}} [j_{\text{rcv}}]. \end{aligned} \quad (9)$$

In order to have the result of P elements, the sender must thus transfer the partial result including the data up to $g_{l_{\text{snd}}}^s [j_{\text{snd}}]$ and the receiver takes out the elements up to $g_{l_{\text{rcv}}} [j_{\text{rcv}}]$.

2) $l_{\text{rcv}} > 0, l_{\text{snd}} = 0$:

$$\begin{aligned} P &= s(n+1)^{k-1} - n \sum_{i=0}^{l_{\text{rcv}}-2} (n+1)^i - j_{\text{rcv}}(n+1)^{l_{\text{rcv}}-1} \\ &= s(n+1)^{k-1} - g_{l_{\text{rcv}}} [j_{\text{rcv}}] \end{aligned} \quad (10)$$

and thus we see that the sender must send only its own data, while the receiver takes out data up to $g_{l_{\text{rcv}}} [j_{\text{rcv}}]$.

3) $l_{\text{rcv}} = 0, l_{\text{snd}} = 0$:

$$P = s(n+1)^{k-1}. \quad (11)$$

In the first $k-1$ rounds, the receiving rank will already have the partial result of $n \sum_{i=1}^{k-1} (n+1)^i = (n+1)^{k-1} - 1$ elements. In the last round it then receives the partial sums of $(s-1)(n+1)^{k-1}$ further elements by the first $s-1$ senders and can thus compute the partial result from a total of $(s-1)(n+1)^{k-1} + (n+1)^{k-1} = s(n+1)^{k-1} - 1$ elements. Including its own data makes the final result of $s(n+1)^{k-1} = P$ elements. If $s = 1$ the algorithm is done after $k-1$ rounds.

4) $l_{\text{rcv}} = 0, l_{\text{snd}} = 1$:

$$P = s(n+1)^{k-1} + j_{\text{snd}} \quad (12)$$

Following the same argumentation as above, the receiving rank will have the partial result of $s(n+1)^{k-1} - 1$ elements. It thus still needs

$$\begin{aligned} &P - (s(n+1)^{k-1} - 1) \\ &= s(n+1)^{k-1} + j_{\text{snd}} - s(n+1)^{k-1} + 1 \\ &= j_{\text{snd}} + 1 \end{aligned} \quad (13)$$

elements. Now, taking into account its own data it still needs j_{snd} data elements. The data boundary $g_1 [j_{\text{snd}}]$ of the sender includes j_{snd} elements plus its own data, i.e., $j_{\text{snd}} + 1$ elements. The $j_{\text{snd}}^{\text{th}}$ element will then be the receiving rank's data, thus it suffices to send up to $g_1 [j_{\text{snd}} - 1]$.

5) $l_{\text{rcv}} = 0, l_{\text{snd}} > 1$:

$$\begin{aligned} P &= s(n+1)^{k-1} \\ &\quad + n \sum_{i=0}^{l_{\text{snd}}-2} (n+1)^i + j_{\text{snd}}(n+1)^{l_{\text{snd}}-1} \end{aligned} \quad (14)$$

In this case, the sender sends a partial result which necessarily includes the initial data of the receiving rank. This means that the receiving rank has to take out its own initial data from the final result. Due to $l_{snd} > 1$ the sender will not be able to take a single initial data element out of the partial result to be transferred. ■

Note that the case where a data boundary on the sending side corresponds to the base border on the receiving side, i.e., $g_{l_{snd}}^s[j_{snd}] = g_B$, has not been covered above. In this case, there is no 5-tuple like above, but rather $P - 1 = g_{l_{snd}}^s[j_{snd}]$ and the adaption and reasoning complies to case 4 in the above theorem.

V. EXPERIMENTS AND EVALUATION

In the following, it is assumed that an optimal implementation for an allreduce function on modern ccNUMA architectures will always take the form of a hybrid implementation where communication within the node or socket is using shared socket memory and communication across nodes is performed via distributed memory. This not only complies to the GASPI specification but also follows the approach of Panda et al. [13]. The n -way dissemination algorithm was implemented on top of `ibverbs` and then compared to different MPI implementations, as this is the de facto standard of distributed memory communication. The implementation of the adapted n -way dissemination algorithm was tested on two different systems:

Cluster 1: A system with 96 nodes, each having two sockets with 12 core Ivy Bridge E5-2695 v2 @2.40GHz processors and an InfiniBand ConnectX FDR interconnect in fat tree configuration. On this system, the algorithm was compared to the `allreduce` and `barrier` of the Intel MPI 4.1.3 implementation, as this was the fastest MPI implementation available on the system.

Cluster 2: A system with 36 nodes, each having two sockets with 6 core Westmere X5670 @2.93GHz processors and an InfiniBand ConnectX QDR network in fat tree configuration. On this system, the algorithm was compared to the `allreduce` and `barrier` routines of MVAPICH 2.2.0 and OpenMPI 1.6.5, because no Intel MPI implementation is available on this system.

In order to obtain the following average runtimes, the algorithms were run 10^3 times on the larger system and 10^5 times on the smaller system to balance single higher runtimes which may be caused through different deterministically irreproducible aspects like jitter, contention in the network or similar. Timings were taken right before the call and again right after the call returned to obtain the average runtimes.

Similar to [10] it was observed, that the choice of n influences the runtime of the routine. The choice of n for the n -way dissemination algorithm will primarily depend on message rate, latency and bandwidth of the underlying network. Ultimately n will have to be determined as a function of these parameters or - much simpler, but less elegant - in the form of static (but network dependent) lookup tables. For the experiments in this section, different $n \in \{1, \dots, 7\}$ were tested in the first call of the routine. The routine was internally run 15 times for each n and timed. The n with the lowest average runtime was chosen for the following runs. This overhead is included in the following runtime plots.

Comparison of Averaged Runtimes of the n -way Barrier and two MPI Versions
2 x 6 core Westmere X5670, 2.93GHz, InfiniBand ConnectX QDR

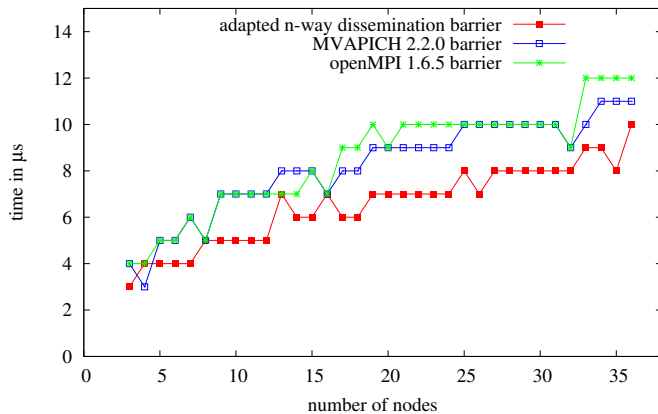


Figure 2. Comparison of the adapted n -way dissemination barrier to available MPI barrier implementations: OpenMPI 1.6.5 and MVAPICH 2.2.0 on Cluster 2.

Comparison of Averaged Runtimes of the Barrier
2 x 12 core Ivy Bridge E5-2695 v2 2.40GHz, InfiniBand ConnectX FDR




Figure 3. Comparison of the adapted n -way dissemination barrier to the Intel MPI 4.1.3 barrier implementation on Cluster 1.

Even though from the theoretical side it is not necessary to adapt the n -way dissemination algorithm for a correct implementation of the barrier operation or idempotent `allreduce` functions, the adaption of the last round brings benefits in the runtime for these routines also. In Figure 2, two MPI barrier implementations (MVAPICH 2.2.0 and OpenMPI 1.6.5) are compared to the adapted n -way dissemination barrier on Cluster 2. The same was done on Cluster 1, the results of the averaged runtimes are shown in Figure 3. The runtime of the n -way barrier implementation is not much lower than that of the MPI implementations on the QDR system but on the larger FDR system a great improvement to the Intel MPI implementation can be seen. On the QDR system, the barrier reaches an average runtime of approximately $8 \mu s$ when running on 32 nodes. On the larger system this runtime decreases to approximately $5 \mu s$. This significant decrease has to be related to the underlying network and the much higher message rate available on FDR. Here a higher n can be chosen, which correspondingly increases the parallelism in

communication and decreases the number of communication rounds needed.

In Figure 4, the first comparisons between MPI allreduces and the adapted n -way dissemination allreduce on the QDR connected system are plotted. There are two different setups: In the first case the allreduce routine was called with one integer as the initial data element to be reduced per process. (upper plot in Figure 4). While there was some improvement compared to the MPI implementations when testing the barrier, in this case no improvement can be seen. This can be explained through the additional time needed for computation in the allreduce.

In the second case the reduce operation works element-wise on an array of 255 elements (lower plot in Figure 4). Here the difference in the runtimes of the OpenMPI allreduce and the MVAPICH allreduce are almost neglectable, while the n -way dissemination allreduce has a much faster averaged runtime. This seems somewhat surprising, as butterfly-like algorithms are known for congesting the network. While it is to be expected that congestion becomes more probable, when increasing the message size, the opposite seems to be the case.

In both test settings the averaged times of the MPI implementations seem more uniform than those of the n -way dissemination algorithm. The main reason for this is that the n -way algorithm provides a higher parallelism per communication step and hence is less robust against jitter.

In Figure 5, a similar picture for the allreduce comparisons on the larger system as for the barrier is seen, i.e., on node level substantially lower runtimes are achieved by the n -way dissemination algorithm than those of Intel MPI 4.1.3. Similar to the tests on the first cluster, the improvements in the barrier runtimes are much higher than those in the allreduce runtimes. Potentially even better results can be achieved for the allreduce operation by internally overlapping the communication and a computation of partial results. We will investigate this as part of future work.

While butterfly-like algorithms often suffer from congestion in the network, this problem is mitigated in the adaption of the n -way dissemination algorithm. The high potential of congestion especially arises through the normally symmetric communication scheme of algorithms like the butterfly itself, the pairwise exchange or the n -way dissemination algorithm with $P = (n + 1)^k$. By using the algorithm for $P \neq (n + 1)^k$, this symmetry is dissolved. Two further arguments for the use of this algorithm comes straight from the GASPI specification. First of all the message sizes are very limited. According to the GASPI specification the largest messages in an allreduce may be 255 doubles. This amounts to only 2040B, which is less than the limit of 4KB of an InfiniBand packet. In addition, only one allreduce operation may be active per GASPI group at a time. Assuming that applications will not create multiple identical groups, the paths for the packets to be transmitted will not be identical either in multiple allreduces.

VI. CONCLUSION AND FUTURE WORK

In this paper we have presented an adaption to the n -way dissemination algorithm, such that it is usable for (user-defined) allreduce operations within GASPI. The main advantage of this modified algorithm is not just an excellent performance, but also the smaller depth of the communication

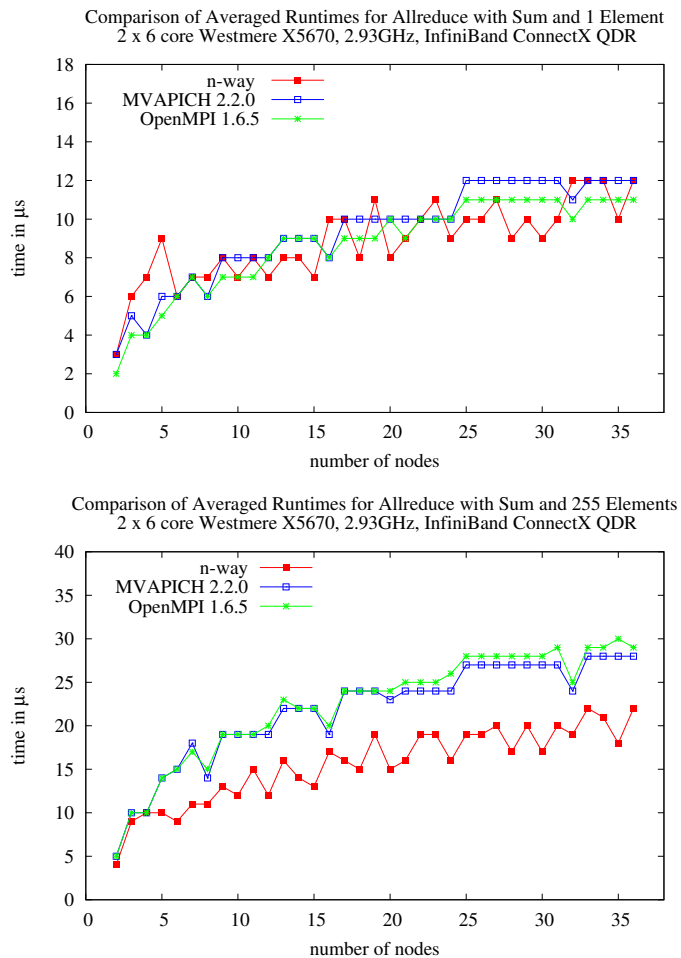


Figure 4. Comparison of the adapted n -way dissemination allreduce to two MPI allreduce implementations with one element per rank to be reduced in the upper plot and 255 elements to be reduce in the lower plot.

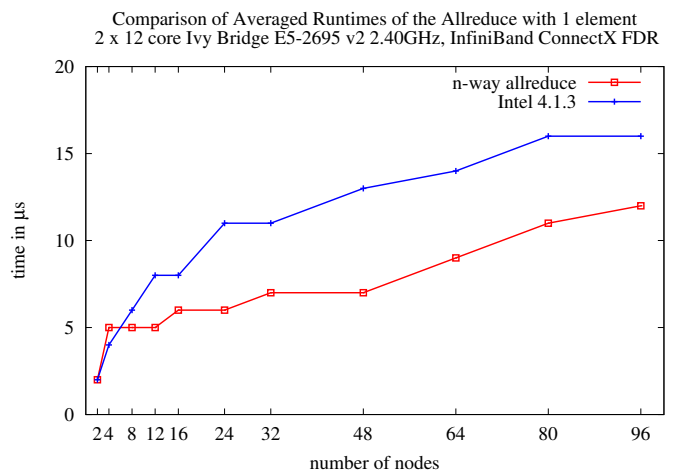


Figure 5. Comparison of the averaged runtimes of the Intel MPI 4.1.3 MPI_Allreduce and the n -way dissemination algorithm based on ibverbs.

tree and the fact that this algorithm allows for a participation of all ranks in every step of the communication tree. Both

features are key requirements for split-phase collectives and user-defined allreduces of the GASPI PGAS API. Since the n -way dissemination algorithm only requires $\lceil \log_{n+1}(P) \rceil$ communication rounds, a split-phase collective (with progress only within the call) will achieve faster progress than algorithms with deep communication trees mentioned in Section I.

While butterfly-like communication schemes have been abandoned in the past, new hardware with much higher throughput might again make these algorithms with small numbers of communication rounds more attractive in the future. Also, topology developments in high performance interconnects like the Cray Aries interconnect make these communication schemes interesting again [17]. Future work will correspondingly investigate the influence of these new developments on butterfly-like algorithms.

As already mentioned in Section II, future work will also deal with an in-depth analysis and comparison of the adapted n -way dissemination algorithm to other algorithms with similar communication schemes.

REFERENCES

- [1] GASPI Consortium, "GASPI: Global Address Space Programming Interface, Specification of a PGAS API for communication, Version 1.0.1," November 2013. [Online]. Available: <http://www.gaspi.de/fileadmin/GASPI/pdf/GASPI-1.0.1.pdf>, 2015.05.07
- [2] Message-Passing Interface Forum, MPI: A Message Passing Interface Standard, Version 3.0. High-Performance Computing Center Stuttgart, 09 2012. [Online]. Available: <http://www.mpi-forum.org/docs/mpi-3.0/mpi30-report.pdf>, 2015.05.07
- [3] N.-F. Tzeng and H. lung Chen, "Fast compaction in hypercubes," IEEE Transactions on Parallel and Distributed Systems, vol. 9, 1998, pp. 50–55.
- [4] J. M. Mellor-Crummey and M. L. Scott, "Algorithms for scalable synchronization on shared-memory multiprocessors," ACM Transactions on Computer Systems, vol. 9, no. 1, Feb. 1991, pp. 21–65.
- [5] E. D. Brooks, "The butterfly barrier," International Journal of Parallel Programming, vol. 15, no. 4, 1986, pp. 295–307.
- [6] D. Hensgen, R. Finkel, and U. Manber, "Two algorithms for barrier synchronization," International Journal of Parallel Programming, vol. 17, no. 1, Feb. 1988, pp. 1–17.
- [7] Argonne National Laboratory, "Barrier implementation of MPICH 3.1.2." [Online]. Available: <http://www.mpich.org/static/downloads/3.1.2/>, 2015.05.07
- [8] "README File of GASNet." [Online]. Available: <http://gasnet.lbl.gov/dist/README,2015.05.07>
- [9] "Barrier implementation of GPI2-1.1.0." [Online]. Available: <https://github.com/cc-hpc-itwm/GPI-2/releases/tag/v1.1.0>, 2015.05.07
- [10] T. Hoefler, T. Mehlan, F. Mietke, and W. Rehm, "Fast barrier synchronization for infiniband," in Proceedings of the 20th International Conference on Parallel and Distributed Processing, ser. IPDPS'06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 272–272.
- [11] J. Bruck and C.-T. Ho, "Efficient global combine operations in multi-port message-passing systems," Parallel Processing Letters, vol. 3, no. 04, 1993, pp. 335–346.
- [12] J. Bruck, C. tien Ho, S. Kipnis, E. Upfal, and D. Weathersby, "Efficient algorithms for all-to-all communications in multi-port message-passing systems," in IEEE Transactions on Parallel and Distributed Systems, 1997, pp. 298–309.
- [13] S. P. Kini, J. Liu, J. Wu, P. Wyckoff, and D. K. Panda, "Fast and scalable barrier using rdma and multicast mechanisms for infiniband-based clusters," in Recent Advances in Parallel Virtual Machine and Message Passing Interface. Springer, 2003, pp. 369–378.
- [14] V. Tipparaju, J. Nieplocha, and D. Panda, "Fast collective operations using shared and remote memory access protocols on clusters," in Proceedings of the 17th International Symposium on Parallel and Distributed Processing, ser. IPDPS '03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 84.1–.
- [15] InfiniBand Trade Association, InfiniBand architecture specification: release 1.3, R. Dance, Ed. InfiniBand Trade Association, 03 2015, vol. 1. [Online]. Available: http://www.infinibandta.org/content/pages.php?pg=technology_public_specification,2015.05.07
- [16] E. Zahavi, "Fat-tree routing and node ordering providing contention free traffic for mpi global collectives," J. Parallel Distrib. Comput., vol. 72, no. 11, Nov. 2012, pp. 1423–1432.
- [17] B. Alverson, E. Froese, L. Kaplan, and D. Roweth, "Cray XC series network," [Online]. Available from: <http://www.cray.com/Assets/PDF/products/xc/CrayXC30Networking.pdf>, 2015.05.07.

Gauss-Jordan Matrix Inversion Speed-Up using GPUs with the Consideration of Power Consumption

Mahmoud Shirazi

Mehdi Kargahi

Farshad Khunjush

School of Computer Science
Institute for Research in
Fundamental Sciences (IPM)
Tehran, Iran

School of Computer Science
Institute for Research in
Fundamental Sciences (IPM)
Tehran, Iran

School of Computer Science
Institute for Research in
Fundamental Sciences (IPM)
Tehran, Iran

Email: m.shirazi@ipm.ir

School of ECE, Faculty of Engineering
University of Tehran, Tehran, Iran
Email: kargahi@ut.ac.ir

School of Electrical and Computer Engineering
Shiraz University, Shiraz, Iran
Email: khunjush@shirazu.ac.ir

Abstract—Matrix inversion is an important requirement for many scientific and practical applications. Parallel architectures like Graphics Processing Units (GPUs) can have noticeable impacts on accelerating such computations, however, at the expense of more power consumption. Both computation speed and power become more challenging when the size of matrices gets larger. This paper proposes an accelerated version of the Gauss-Jordan matrix inversion method using GPUs. The experimental results show that the proposed method is faster than a recent baseline method. As an additional property, the configurations have been made in a manner to consume less power.

Keywords—Matrix Inversion; GPU; Power Consumption.

I. INTRODUCTION

Matrix inversion is too important in a diverse spectrum of applications including graph theory, computer graphics, cryptography, etc. There are several known methods like Gauss-Jordan [1], Strassen [2], Strassen-Newton [3], Cholesky decomposition [4], and Lower Upper Decomposition [5] for matrix inversion.

As calculating the inverse of a matrix takes $O(n^3)$ (where n is the dimension of the matrix), the usage of many core architectures, like GPU is too effective to reduce the computation time of inverting large matrices. In fact, the highly parallel architecture of GPUs makes them more suitable than general-purpose Central Processing Units (CPUs) for algorithms with large parallelizable blocks of data.

The recursive nature of some matrix inversion algorithms such as Strassen-Newton reduces the effectiveness of parallelization [6]. In recursive algorithms, each iteration needs to wait for the result of previous iteration. In Gauss-Jordan method, each iteration calculates the new values of all elements in the matrix that cause this method to be more appropriate for parallelization. GPUs has demonstrated high computing power in several application fields and we can use GPUs to speedup the Gauss-Jordan matrix inversion method. On the other hand, GPU also produces high power consumption and has been one of the largest power consumers in desktop and supercomputer systems [7].

Girish Sharma et al. [6] proposed a fast GPU-based Parallel Gauss-Jordan method (PGJM). This algorithm calculates the

inverse of matrix in two steps using Gauss-Jordan method. Their method uses n and $n \times (n - 1)$ threads for steps 1 and 2, respectively. As the matrix size increases, the parallel Gauss-Jordan method will need more threads, blocks, streaming multiprocessors (SMs) and thus more processing cores. In addition, using more processing cores results more power consumption.

Comparing to the method that proposed in [6], here, we improve this method using one step for calculating the inverse of matrix. The improved parallel Gauss-Jordan method (I-PGJM) is faster than PGJM. In fact, using different grid and block dimensions it can affect the power consumption. Thus, We can make some trade-off between the speed-up and power consumption of I-PGJM.

The remainder of this paper is organized as follows. Section II presents a background of Gauss-Jordan and parallel Gauss-Jordan methods, and addresses the highlights to achieve a program that consumes less power. Section III proposes a fast parallel Gauss-Jordan method that has lower power consumption. Experimental results are presented in Section IV. Section V discusses the related work and Section VI concludes the paper.

II. BACKGROUND

This section discusses the Gauss-Jordan method [1], the parallel Gauss-Jordan method [6] and the programming-level techniques for reducing the power consumption.

A. The Gauss-Jordan Method

Gauss-Jordan is an accurate method for matrix inversion that is based on Gauss-Jordan elimination. This method for calculating the inverse of matrix A of size $n \times n$ augment the matrix with the identity matrix (I) of the same size. Thus we have a matrix $C = \{c_{ij}\}$ of size $n \times 2n$ that its left half is A and its right half is I . This method converts the left half of C to identity matrix after some row operations. At the end, the right half of C is equal to A^{-1} .

Converting the left half of C can be broken down into two steps per each column of the left half matrix. The first step converts c_{ii} to 1 by dividing all elements in i th row by c_{ii} . The second step convets all elements in i th column to 0 except

for the i th one. Following these two steps for each column sequentially, the left half of C becomes the unit matrix while the right half becomes the desired inverse of A .

The first step needs to $O(n)$ divisions and the second step needs to $O(n^2)$ computations. Thus, this method takes $O(n^3)$ for calculating the inverse of matrix.

B. The Parallel Gauss-Jordan Method

Girish Sharma et al. [6] proposed the parallel version of Gauss-Jordan method. They implement their method with CUDA on Nvidia GTX 260. Their method used two steps. The first one uses n threads (equal to number of rows or columns in matrix A) that each divides one element of current row ($c_{rr}, \dots, c_{r(n+r)}$, where r is the current row index) by c_{ii} . Preventing division by zero exception can be done by checking the value of c_{ii} . If its value was zero then the elements of current row added to elements of k th row such that $c_{kk} \neq 0$. To avoid the extra GPU cycles, resulting from if...else condition, this summation is done prematurely for all elements. The second step uses $n \times (n - 1)$ threads to convert the elements of r th columns to zero except for c_{rr} . In this step, they minimize the computations by processing only the first n columns (starting from the r th columns).

This algorithm uses n threads in one block to avoiding division by zero, n threads per one block for step 1 and n threads in n blocks for step 2 that each block handles the computations of one column at a time. Since the maximum number of threads per block in modern GPU are 1024, if the number of columns is greater than 1024, then the matrix splits vertically.

This algorithm was programmed using CUDA C and was tested against various sizes of the several types of matrices, e.g., identity, sparse, band, random and hollow matrix. The results demonstrate that if the size of the matrix (n) is less than 100, the computation time is linear to the size of the matrix because the maximum number of active threads that can be scheduled at a time is 9216 and for $n = 100$ the number of threads that need for computations at a time is 10100. Thus values around $n = 100$ displaying the quadric nature of the algorithm due to hardware limitation.

In here, we propose a method that uses one step per each column to set c_{rr} to 1 and the remain elements in the r th column to 0. This method that was described in Section III, uses a total number of $n \times (n - 1)$ threads and is faster than PGJM.

C. GPU Power-aware Programming

The power consumption of GPUs can be divided into two parts, namely, leakage power and dynamic power [8]. Dynamic power is the power that is consumed by a device when it is actively switching from one state to another. The main concern with leakage power is when the device is in its inactive state. Different components, such as SMs and memories, e.g., local, global, shared, etc. contribute to dynamic power consumption.

Techniques for reducing the GPU power consumption are classified into five categories [8]: dynamic voltage/frequency scaling (DVFS), CPU-GPU workload division-based techniques, architectural techniques, techniques that exploit workload variation to dynamically allocate resources and programming-level techniques.

For reducing the power consumption of Gauss-Jordan method for large matrices, we use programming-level techniques. Yi Yang et al. [9] identify the common code patterns that lead to inefficient use of GPU hardware and increase the power consumption. These code segments that they call it 'bug' are grouped into following categories:

- Global memory data types and access patterns
- The thread block dimensions
- Portability across different GPUs
- The use of constant and texture memory
- Floating-point number computations

Using data types either smaller than 'float' or larger than 'float4' (the data with a size of 4 floats) violates the memory coalescing requirements. The results show that GTX480 delivers much lower bandwidth when the data type is not float, float2, or float4.

In modern GPUs, threads are organized in a thread hierarchy, multiple threads forming a 1D/2D/3D thread block, and thread blocks forming a thread grid. Different size in the thread hierarchy results different performance and power consumption in different GPUs. The experimental results in [9] demonstrate that if the thread ids are used as memory access addresses, the size of the x-dimension in a thread block needs to be at least 32 for GTX480 to maximize the global memory access bandwidth and the optimal configuration to improve data reuse using shared memory is application and data size dependent.

Since different GPUs have different hardware features, the configuration of the program must be altered with different GPUs. However, different hardware features in different GPUs are not necessitating significantly different performance considerations.

Proper use of constant and/or texture memory can achieve high performance and low power consumption due to the on-chip constant and texture caches. Thus, if one method uses constant values, it is better to use constant or texture memory.

If we use a constant float, such as 4.0, the CUDA compiler will treat the constant as a double-precision number that results in less performance and higher power consumption. For fixing this bug if single-precision provides sufficient accuracy, we can use the explicit single-precision floating-point number (4.0f).

We used these hints for reducing the power consumption of I-PGJM. The details of I-PGJM power consumption consideration are shows in Subsection III-B.

III. IMPROVED PARALLEL GAUSS-JORDAN METHOD

In this section, we first propose the improved method for matrix inversion and then consider the power consumption of this method.

A. I-PGJM

As mentioned in Section II, Girish Sharma et al. [6] proposed a two-step method for calculating the inverse of a matrix using Gauss-Jordan method. In this section, we present an improved method that combines these two steps into one and calculates the inverse of a matrix faster than this method.

Suppose we want to calculate the inverse of matrix A of size $n \times n$. The Gauss-Jordan method as mentioned in Section

Function FixAll
Input:

Matrix: Input matrix
colld: Current column of input matrix

Begin

Set T to thread id in x dimension
Set B to block id in x dimension
Set D to Matrix [colld, colld]
Set SharedRow[T] to Matrix[colld, T+colld]
 $C \leftarrow \text{Matrix}[B, \text{colld}] / D$
Sync all threads
IF B is equal to colld THEN
 $\text{Matrix}[B, T+\text{colld}] \leftarrow \text{Matrix}[B, T+\text{colld}] / D$
ELSE
 $\text{Matrix}[B, T+\text{colld}] \leftarrow \text{Matrix}[B, T+\text{colld}] - C * \text{SharedRow}[T]$

End

 Figure 1. The I-PGJM method for matrix inversion in GPU ($n \leq 1022$).

It creates a matrix $C = \{c_{ij}\}$ of size $n \times 2n$ so that its left half is A and its right half is the identity matrix. The right half of matrix C After n iterations that each consists of two steps is the inverse of A .

The first step of iteration r , sets $c_{rj} = \frac{c_{rj}}{c_{rr}}$ where $j = r, \dots, n+r$ (Note that there is no need to update the values for $j < r$ and $j > n+r$ [6]) and the second step sets

$$c_{ij} = c_{ij} - c_{ir} \times c_{rj} \quad (1)$$

where $i = 1, \dots, n$, $i \neq r$ and $j = r, \dots, n+r$. The step 1 and 2 must run sequentially as the step 2 needs the new value of c_{rj} calculated after step 1. To run these two steps in a time, we can define a coefficient $co_{ir} = \frac{c_{ir}}{c_{rr}}$ and rewrite equation 1 as:

$$c_{ij} = c_{ij} - co_{ir} \times c_{rj} \quad (2)$$

where c_{rj} is the original value of r th element in C . Then, in iteration r we can share the elements of r th row among threads. Thus the elements of matrix C in iteration r are calculated as follows:

$$c_{ij} = \begin{cases} c_{ij} - co_{ir} \times \text{sharedRow}_j & i \neq r \\ \frac{c_{ij}}{c_{rr}} & i = r \end{cases} \quad (3)$$

where $j = r, \dots, n+r$ and the values of c_{rr} and sharedRow are calculated before updating the values of C using *syncthreads* in CUDA. Using n threads each in n blocks can calculate the inverse of matrix A . Figure 1 displays the I-PGJM method implemented using CUDA C.

To avoid the division by zero and extra GPU cycles resulting from if...else condition, before the above computations the summation of each row to next row is calculated. Thus, we need $n+2$ threads in each iteration. As each block calculates the values of one row and the maximum number of threads in each block are 1024, for $n+2 > 1024$ we need to break down the calculation of each row into several blocks. Thus for $n > 1022$ the grid dimension is $N \times n$ where $N = \lfloor \frac{n}{1024} \rfloor + 1$ and each block has $\lceil \frac{n+2}{N} \rceil$ threads. Figure 2 shows the method for calculating the inverse of large matrices ($n > 1022$).

The total number of threads can be larger than $n+2$. Then, the 'if' condition at the beginning of the method in Figure 2 is necessary to avoid accessing the out of range matrix

Function FixAll
Input:

Matrix: Input matrix
colld: Current column of input matrix

Begin

$T \leftarrow \text{blockId}.x * \text{blockDim}.x + \text{threadId}.x$
IF T+colld is greater than Matrix dimension THEN
 Do nothing and return
Set TIndex to thread id in x dimension
Set B to block id in y dimension
Set D to Matrix [colld, colld]
Set SharedRow[TIndex] to Matrix[colld, T+colld]
 $C \leftarrow \text{Matrix}[B, \text{colld}] / D$
Sync all threads
IF B is equal to colld THEN
 $\text{Matrix}[B, T+\text{colld}] \leftarrow \text{Matrix}[B, T+\text{colld}] / D$
ELSE
 $\text{Matrix}[B, T+\text{colld}] \leftarrow \text{Matrix}[B, T+\text{colld}] - C * \text{SharedRow}[T\text{Index}]$

End

 Figure 2. The I-PGJM method for calculating the inverse of large matrices ($n > 1022$) in GPU.

dimension. As each block calculates at most 1024 elements of the current row in the matrix, the size of shared variables can be set to 1024.

Comparing the I-PGJM and PGJM [6], I-PGJM used $(n+2) \times \text{sizeof}(\text{float})$ shared data rather than $(3n+2) \times ((n+1) \times \text{sizeof}(\text{float}))$ in step 1 and $(2n+1) \times \text{sizeof}(\text{float})$ in step 2). As mentioned in [6] if the matrix size is greater than 1024 the proposed algorithm will use more blocks for computing the inverse of large matrices. Thus, for large matrices these two methods need to have more blocks.

The experimental results in Section IV show that I-PGJM is faster than the method that proposed by Sharma et al. [6].

B. Power consumption considerations

In the I-PGJM, we try to reduce the use of registers, shared memory and global memory. But increasing the matrix dimension lead to use more blocks for calculating the matrix inversion and hence more power consumption. As mentioned in Subsection II-C, using different data types and access patterns can affect the power consumption. We used different data types for matrix inversion in K20Xm and find that it has similar behaviour as GTX480. Also, a few common data access patterns exist that may cause performance degradation. This performance degradation in K20Xm (that we used in here) is negligible.

Using constant or texture memory can reduce the power consumption. In Gauss-Jordan method, the matrix elements should be updated iteratively. Then, we can not use constant or texture memory.

The grid and block size are important for performance and power consumption. We experiment with different grid and block sizes to achieve less power consumption. Achieving the suitable grid and block dimensions lead to have the fast method that can calculate the inverse of large matrices and has less power consumption.

IV. EXPERIMENTAL RESULTS

The PGJM [6] and I-PGJM were programmed using CUDA C and were executed by Tesla K20Xm GPU. This

GPU has a shared memory size of 48KB and the maximum threads per block is 1024. The block and grid dimensions (x, y, z) are respectively restricted to $(1024, 1024, 64)$ and $(65535, 65535, 65535)$. For comparing the results to the sequential program, we used Intel Core i7 CPU (2.67GHz) with 8192 KB and 6GB of cache and memory size, respectively.

We used random matrices that have different dimensions from 2^1 to 2^{13} . the experimental results are shown in two subsections. The first one demonstrates the results of comparing these methods in term of execution time, and the second one shows the power consumption of proposed method for different grid and block dimensions to get the optimal dimensions for reducing the power consumption.

A. Matrix Inversion Execution Time

In this subsection, we compare the execution time of PGJM, I-PGJM and the sequential Gauss-Jordan method with different matrix dimensions (n) . The execution time is calculated using `cudaEventRecord` function and is the difference in time before passing the data from host to device and after getting the results from device. For $n > 1022$, the grid and block dimensions of proposed method are (N, n) and $(\lceil \frac{n+2}{N} \rceil, 1)$, respectively and for PGJM (for step 2) are $(\frac{n}{1024}, n+2)$ and $(1024, 1)$, respectively. For sequential algorithm, the execution time for $n < 2^5$ is less than one millisecond (see Table I). Figure 3 shows the speedup of I-PGJM with respect to PGJM. The values around 100 for matrix dimension display the quadric nature due to hardware limitation [6]. This behaviour results more speedup for proposed method as matrix dimension increased ($speedup = 9.02$ for $n = 2^{13}$). Thus, I-PGJM is suitable for large matrices. For $n > 1024$ these two methods need to use more blocks, then the speedup slightly reduced for $n = 2048$.

TABLE I. THE EXECUTION TIME (MS) OF I-PGJM, PGJM AND THE SEQUENTIAL ALGORITHM FOR DIFFERENT MATRIX DIMENSIONS.

n	I-PGJM	PGJM	Sequential Algorithm
2	0.245	0.255	0
2^2	0.273	0.284	0
2^3	0.310	0.353	0
2^4	0.409	0.488	0
2^5	0.575	0.784	0
2^6	0.955	1.412	3
2^7	1.844	3.536	10
2^8	4.761	9.964	120
2^9	21.570	62.314	950
2^{10}	155.691	497.304	7550
2^{11}	1410.690	4209.350	60500
2^{12}	9750.890	44424.500	482619
2^{13}	72648.400	655863.000	3860219

Figure 4 shows the speedup of I-PGJM with respect to the sequential Gauss-Jordan method. The speedup of I-PGJM with respect to the sequential program is around 53 for $n = 2^{13}$. Since the maximum number of threads per block are 1024, for $n > 1024$ the algorithm need to use more blocks and hence the performance drop occurred for $n = 2048$.

B. Power Consumption

The execution time results demonstrate that I-PGJM is more suitable for calculating the inverse of large matrices. As increasing the computation results in more power consumption, the proposed method needs to use a configuration that leads




Figure 3. The speedup of the proposed method related to method in [6].




Figure 4. The speedup of I-PGJM with respect to the sequential algorithm.

to less power consumption. As mentioned in Subsection II-C and Section III, we used 'float' data type, fewer threads and shared memory to reducing the power consumption. Also, we experiment with different grid and block dimensions to reducing the power consumption.

We measure the run-time power of the proposed method with the NVIDIA Management Library (NVML) [10] by running the proposed method on a thread and NVML on another thread using Pthreads. NVML is a high level utility that called `nvidia-smi`. NVML can be used to measure power when running the kernel but since `nvidia-smi` is a high level utility the rate of sampling power usage is very low and unless the kernel is running for a very long time we would not notice the change in power [11]. The `nvmlDeviceGetPowerUsage` function in the NVML library retrieves the power usage reading for the device, in milliwatts. This is the power draw for the entire board, including GPU, memory, etc. The reading is accurate to within a range of +/- 5 watts error.

In the proposed method, all blocks in one row of a grid




Figure 5. The execution time and power consumption of I-PGJM and PGJM for different grid dimensions. The matrix size is $2^{12} \times 2^{12}$.

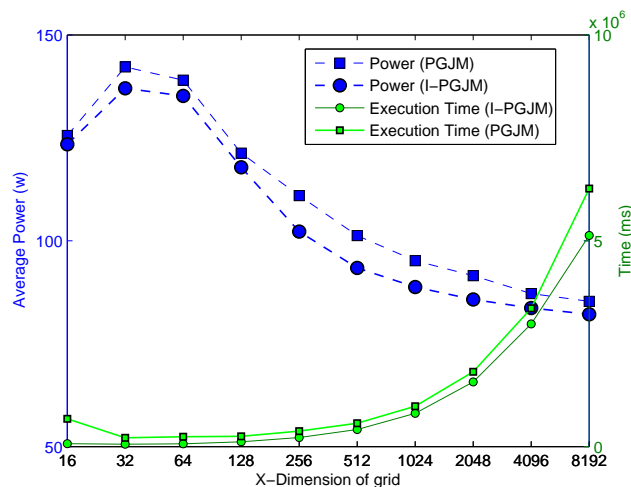


Figure 6. The execution time and power consumption of I-PGJM and PGJM for different grid dimensions. The matrix size is $2^{13} \times 2^{13}$.

compute the elements of one row in matrix and each block has $\lceil \frac{n+2}{N} \rceil$ threads related to a subset of elements in the corresponding row. For constant n , we calculate the execution time and power consumption of the proposed method with different grid and block dimensions. The x-dimension of grid can be changed from N to n . The grid sizes that used for experiments are from 2^i to 2^j where $2^i > N$ and $2^j < n + 2$ for $n > 1022$. The block dimension can get different values such that the total number of threads per block remain constant, i.e., $\lceil \frac{n+2}{N} \rceil$.

Figures 5 and 6 show the execution time and power consumption of I-PGJM and PGJM for calculating the inverse of matrix A of size $n \times n$ for different values of grid dimensions for $n = 2^{12}$ and $n = 2^{13}$, respectively. Increasing the number of blocks in x-dimension of grid results in lower threads in each block (Matrix dimension is constant). Thus, each block used less resources that lead to lower power consumption. On the other hand, the maximum number of blocks that can be scheduled in each SM is limited. Thus, increasing the number of blocks results more execution time.

Figures 7 and 8 show the experimental results for different block dimensions of I-PGJM for calculating the inverse of matrix A of size $n \times n$ for $n = 2^{12}$ and $n = 2^{13}$, respectively. As can be seen in these figures, the difference between maximum and minimum power consumption is negligible. Note that, the power usage reading of NVML is accurate within a range of ± 5 wats error. But, when the x-dimension of blocks is larger than 32, the power consumption increases. Thus, using less than 32 threads in x-dimension of blocks leads to less power consumption. Also, by simply changing the x-dimension of blocks to 32, we have better performance.

V. RELATED WORK

There has been much recent work decreasing the execution time of Gauss-Jordan method for matrix inversion. Some researchers focus on using multiple GPUs [12]. Although splitting the matrix between GPUs and ability to use them is important, in this study, we focus on one GPU and propose a




Figure 7. The execution time and power consumption of I-PGJM for different block dimensions. The matrix size is $2^{12} \times 2^{12}$.

method that has a considerable reduction of the computational time and power consumption.

Peter Benner et al. [13] have studied the inversion of large and dense matrices, on hybrid CPU-GPU platforms, with application to the solution of matrix equations arising in control theory. They have presented several matrix inversion algorithms, based on common matrix factorizations and the GJE method, similar to [12], but for solving equations in control theory.

Kaiqi Yang et al. [14] proposed a parallel matrix inversion algorithm based on Gauss-Jordan elimination with pivoting. This method divides the matrix into some sub-matrices and assign each sub-matrices to one core. Each core updates the sub-matrix using Gauss-Jordan elimination. This method experimented in at most four cores and has communication overhead. Then, this method can not used in many core platforms such as GPU.

The most related work to this paper is the method that is




Figure 8. The execution time and power consumption of I-PGJM for different block dimensions The matrix size is $2^{13} \times 2^{13}$.

proposed in [6]. In here, we proposed a method that reduces the computation times of this method. Also with the power consumption hints that described in [8] and [9], we used the configuration that reduces the power consumption of the proposed method.

VI. CONCLUSION

In this paper, we improved the execution time of Gauss-Jordan matrix inversion method on GPU. In the case of large matrices, as increasing the computations, the power consumption requires more attention. The proposed method used fewer threads and can be configured with different grid and block dimensions. Thus, we can find the configuration that has lower power consumption. The results show that the block dimensions has negligible effect on the power consumption of proposed method and trade-off between performance and power can improve the power consumption of I-PGJM on GPU.

REFERENCES

- [1] S. C. Althoen and R. McLaughlin, "Gauss-Jordan Reduction: A Brief History," *Mathematical Association of America*, vol. 94, 1987, pp. 130–142.
- [2] V. Strassen, "Gaussian elimination is not optimal," *Numer Math*, vol. 13, 1969, pp. 354 – 356.
- [3] D. H. Bailey and H. R. P. Gerguson, "A Strassen-Newton algorithm for high-speed parallelizable matrix inversion," in *In Supercomputing '88 Proceedings of the 1988 ACM/IEEE conference on Supercomputing*, 1988, pp. 419–424.
- [4] A. Burian, J. Takala, and M. Ylinen, "A fixed-point implementation of matrix inversion using Cholesky decomposition." in *Proceedings of the 46th international Midwest symposium on circuits and systems*. Cairo, 2013, pp. 1431–1433.
- [5] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling, Eds., Section 2.3: LU decomposition and its applications, numerical recipes in FORTRAN: the art of scientific computing. New York: Cambridge University Press, 2007.
- [6] G. Sharma, A. Agarwala, and B. Bhattacharya, "A fast parallel Gauss Jordan algorithm for matrix inversion using CUDA," *Computers and Structures*, vol. 128, 2013, pp. 31–37.

- [7] G. Wang, Y. Lin, and W. Yi, "Kernel Fusion : an Effective Method for Better Power Efficiency on Multithreaded GPU," in *IEEE/ACM International Conference on Green Computing and Communications IEEE/ACM International Conference on Cyber, Physical and Social Computing*, 2010, pp. 344 – 350.
- [8] S. Mittal and J. S. Vetter, "A Survey of Methods for Analyzing and Improving GPU Energy Efficiency," *ACM Computing Surveys*, vol. 47, 2014, article No. 19.
- [9] Y. Yang, P. Xiang, M. Mantor, and H. Zhou, "Fixing Performance Bugs: An Empirical Study of Open-Source GPGPU Programs," in *Proceedings of the 41th International Conference on Parallel Processing (ICPP)*, 2012, pp. 329 – 339.
- [10] "NVIDIA Management Library (NVML)," 2014, URL: <https://developer.nvidia.com/nvidia-management-library-nvml> [accessed: 2015-05-10].
- [11] K. Kasichayanula, D. Terpstra, P. Luszczek, and S. Tomov, "Power Aware Computing on GPUs." in *Symposium on Application Accelerators in High Performance Computing (SAAHPC)*, 2012, pp. 64 – 73.
- [12] P. Ezzatti, E. S. Quintana-Ort, and A. Remon, "High Performance Matrix Inversion on a Multi-core Platform with Several GPUs," in *Proceedings of the 19th International Euromicro Conference on Parallel, Distributed and Network-Based Processing*, 2011, pp. 87 – 93.
- [13] P. Benner, P. Ezzatti, E. S. Quintana-Ort, and A. Remon, "Matrix inversion on CPUGPU platforms with applications in control theory," in *Concurrency and Computation: Practice and Experience*, 2013, pp. 1170 – 1182.
- [14] K. Yang, Y. Li, and Y. Xia, "A Parallel Method for Matrix Inversion Based on Gauss-jordan Algorithm," *Journal of Computational Information Systems*, vol. 14, 2013, pp. 5561 – 5567.

Towards Data Persistency for Fault-tolerance Using MPI Semantics

José Gracia,* M. Wahaj Sethi,[†]* Nico Struckmann,* Rainer Keller[‡]

*High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany

[†]e.solutions GmbH, Ingolstadt, Germany

[‡]University of Applied Sciences, HfT Stuttgart, Germany

Abstract—As the size and complexity of high-performance computing hardware, as well as applications increase, the likelihood of a hardware failure during the execution time of large distributed applications is no longer negligible. On the other hand, frequent checkpointing of full application state or even full compute node memory is prohibitively expensive. Thus, application-level checkpointing of only indispensable data and application state is the only viable option to increase an application’s resiliency against faults. Existing application-level checkpointing approaches, however, require the user to learn new programming interfaces, etc. In this paper we present an approach to persist data and application state, as for instance messages transferred between compute nodes, which is seamlessly integrated into Message Passing Interface, i.e., the de-facto standard for distributed parallel computing in high-performance computing. The basic idea consists in allowing the user to mark a given communicator as having special, i.e., persistent, meaning. All communication through this persistent communicator is stored transparently by the system and available for application restart even after a failure.

Keywords—*Message Passing Interface; MPI, fault-tolerance; application-level checkpointing; data persistency*

I. INTRODUCTION

Numerical simulation on high-performance computing (HPC) systems is an established methodology in a wide range of fields not only in traditional computational sciences as physics, chemistry, astrophysics, but also becoming more and more important in biology, economic sciences, and even humanities. The total execution time of an application is rapidly approaching the mean time between failures of large HPC systems. Commonly, only a small part of the system will be affected by the hardware fault, but usually all of the application will crash. Application developers can therefore no longer ignore system faults and need to take fault-tolerance and application resiliency into account as part of the application logic. A necessary step is to store intermediate result as well as the current internal state of the application to allow restarting the application at a later time, which is commonly referred to as checkpointing.

In practice, however, the sheer size of simulation data and the limited I/O bandwidth prohibit dumping all intermediate results at high frequency [1]. Checkpoints are therefore chosen to satisfy requirements of the scientific analysis of the simulation data. However, most computational experiments, i.e., simulations, are by definition sufficiently robust to allow drawing similar or equal scientific conclusions if initial or boundary conditions – and by extension intermediate

results – are changed slightly within use-case specific limits. Application-level checkpointing of suitably aggregated intermediate results is therefore being considered as a promising technique to improve the resiliency of scientific applications at relatively low cost of resources. The application developer or end user, purposefully discards most of the intermediate data and checkpoints only those data which are absolutely essential for later reconstruction of a sane state. An example would be to store mean values of given quantities, other suitable higher-order moments of the distribution of the quantities, or leading terms of a suitable expansions. Note however, that the nature of the reconstruction data is fully application and even use-case specific. The application at its restart will use this data to reconstruct the state in the part of the application that was lost to the failure, while keeping the full, precise data in the reset of system which was not affected by the fault.

In this paper we present a method for persisting intermediate results and internal application state. Our proposed method uses idioms and an interface borrowed from the Message Passing Interface (MPI) [2], which is the most widely used programming model for distributed parallel computing in HPC. This allows users of MPI to integrate our method seamlessly into existing applications at minimal development cost.

This paper is organized into a brief overview of related work in Section II, followed by a our approach to data persistency through MPI semantics in Section III, and finally a short summary of this work in Section IV

II. RELATED WORK

SafetyNet [3] is an example of checkpointing at the hardware level. It keeps multiple, globally consistent checkpoints of the state of a shared memory multiprocessor. This approach has the benefit of lower overhead of runtime but it as additional power and monetary cost. Right now, this approach provides checkpointing solution for a single node only.

In the kernel-level approach, the operating system is responsible for checkpointing, which is done in the kernel space context. It uses internal kernel information to capture the process state and further important information required for a process restart. Berkeley Lab Checkpoint/Restart (BLCR) [4][5] and Checkpoint/Restore In Userspace (CRIU) [6] are two examples of this class. This approach provides a transparent solution for checkpointing but files generated by this approach are large and moving checkpoint files to stable storage takes more time. Another problem associated with this approach

is that it requires considerable maintenance and development effort as internals of process state, etc. vary greatly from one OS to another and are prone to change over time.

Checkpointing at the user-level solves the problem of high maintenance effort due to kernel diversity. In this case, checkpointing is done in user-space. All relevant system calls are trapped to track the state of a given process. However, due to the overhead of intercepting system calls it takes more time to complete. Similar to kernel-level, user-level checkpointing needs to save complete process state. So, this approach also suffers from the problem of large file size.

In contrast to the schemes mentioned above, which are transparent to user, application-level checkpointing requires explicit user action. The application developer provides hints to the checkpointing framework. By means of these hints, additional checkpoint code is added to the application. This additional code saves required information and restarts the application in case of failure. Application level checkpointing normally creates smaller size checkpoints as they have knowledge about program state.

One such application-level checkpointing scheme is the library Scalable Checkpoint/Restart (SCR) [7][8]. SCR stores checkpoints temporarily in the memory of neighboring compute nodes before writing them to stable storage. It also includes a kind of scheduler which determines the exact checkpointing time according to system health, resource utilization and contention, and external triggers. SCR is designed to interoperate with MPI. The application developer uses SCR functions to mark important data which is then checkpointed transparently in the background at a suitable point in time. The drawback is that application developers have to learn yet another programming interface and add additional, possibly complex code, which is not related to their numerical algorithm.

Previous extensions to MPI, such as FT-MPI [9] offered the application programmer several possibilities to survive, e.g., leave a hole in the communicator in case of process failure. This particular MPI implementation has been adopted in Open MPI [10]. The Message Passing Interface standard in its current form, i.e., MPI-3 [2], does not provide fault-tolerance. Typically, if a single process of a distributed application fails due to, for instance, catastrophic failure of the given compute node, all other processes involved will eventually fail as well in an unrecoverable manner. Recently, several proposals [11][12][13] have been put forward to mitigate the issue by allowing an application to request notification about process failures and by providing interfaces to repair vital MPI communicators. The application, in principle, can use this interface to return the MPI stack to a sane state and continue operation. However, any data held by the failed process is lost. Notably, this includes any messages that have been in flight at the time of the failure.

III. PERSISTENT MPI COMMUNICATION

In this paper we present an approach that allows application developers to persist, both, essential locally held data and the content of essential messages between processes. Unlike other models, we use idioms that are familiar to any MPI developer. In fact, we add a single function which returns

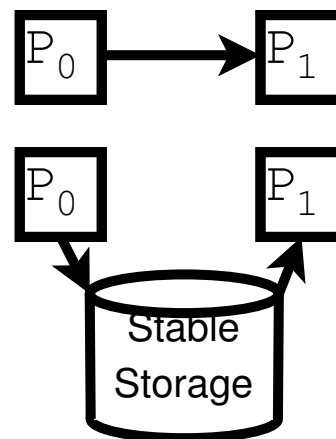


Figure 1. Illustration of communication between two processes, P0 and P1, through regular communicators (top) versus through persistent communicators (bottom).

a MPI communicator with special semantic meaning. Then, the programmer continues to use familiar send and receive MPI calls or collective operations to store data and messages persistently or to retrieve them during failure recovery.

A. Background

In MPI, any process is uniquely identified by its *rank* in a given *communicator*. A communicator can be thought of as an ordered set of processes. At initialization time, MPI creates the default communicator, `MPI_COMM_WORLD`, which includes all processes of the application. New communicators can be created as a subset of existing ones to allow logically grouping processes as required by the application. Collective MPI operations, as for instance a broadcast or scatter, take a communicator as argument and necessarily require the participation of all the processes of the given communicator. In addition, some collective operations single out one process which is identified by its rank in the respective communicator. Also point-to-point communication routines take a communicator as argument. In send operations, the target of a message is passed as rank relative to the given communicator argument. The destination of receive operations is given analogously.

In addition, most MPI communication routines require the specification of the so-called *tag* which allows the programmer to classify different message contents. A tag may be thought of as a P.O. Box or similar. Finally, MPI messages are delivered in the same order they have been issued by the sender. Any MPI message can thus be uniquely identified by the signature tuple $(comm, src, dst, tag)$ and a sequence number that orders messages with the same signature. The signature is composed of a communicator, `comm`, the rank of the message source, `src`, the rank of the destination, `dst`, and the message tag `tag`.

B. Persistent communicators and proposed idioms

The basic idea of our approach is very simple. The user marks a communicator as having a special, i.e., persistent, semantics. Any communication issued through a persistent communicator is stored transparently by the MPI library and is available for application restart even after failure (see


```

1 #define VALUE_TAG
2 const char mykey[] = "Run_A,_June_6_2015";
3 MPI_comm persistent;
4 int value = 3;
5
6 MPI_Comm_persist(MPI_COMM_SELF, info,
7     mykey, persistent);
8
9 if (failure())
10     MPI_Recv(&value, 1, MPI_INT, rank,
11         VALUE_TAG, persistent, &status);
12 else
13     MPI_Send(&value, 1, MPI_INT, rank,
14         VALUE_TAG, persistent);
    
```

Figure 2. Simple example of data persistency

Figure 1). In contrast to no-persistent communicators the message is not immediately delivered. An MPI process may thus persist any data and application state by sending it to itself through a persistent communicator. In case of failure the data is simply restored by posting a receive operation on the persistent communicator. Moreover, a process may persist data for any other process by sending a message targeted to the other process through a persistent communicator.

A communicator is marked as persistent by calling the routine `MPI_Comm_persist` which has the following signature

```

int MPI_Comm_persist(MPI_Comm comm, char *key,
    MPI_Info info, MPI_Comm *persistentcomm)
    
```

Here, `persistentcomm` is a pointer to memory which will hold the newly created persistent communicator. It is derived from the existing communicator `comm` and will consist of the same processes, etc. A user-provided string `key` shall uniquely identify this particular application run or instance in case part of it needs to be restarted after a fault occurred. Essentially this serves as a kind of session key. Finally, the info object `info` may hold additional information for the MPI library, as for instance hints where to store data temporarily, or the size of the expected data volume.

A simple example how to persist data is shown in Figure 2. On line 6, the persistent communicator `persistent` is derived from `MPI_COMM_SELF` which is a pre-defined communicator consisting of just the given process. The routine `failure()` shall return `TRUE` if this process is being restarted after a fault. If this is not the case, the application will persistently store the content of the variable `value` by sending a message to itself on line 13. If a fault occurred the application instead will restore the content of the variable `value` by receiving it from itself on line 10.

Our approach also allows to replay or log communication between processes in the case of faults. The programmer simply derives persistent communicators from all relevant communicators and then mirrors every send operation done on a non-persistent communicators with the persistent one. Receive operations are posted on the persistent communicator as necessary by the failed process only. The reduce the amount of additional code one could also allow transparent persistency. In this case a persistent communicator would persist data and also actually deliver data as expected from a non-transparent

```

1 #define SIZE VERY_LARGE
2 #define SESSION 1001
3
4 int rank, other;
5 float data[SIZE], boundary, seed=321;
6 int iter = 0;
7 MPI_comm persistent, world;
8
9 MPI_Init();
10 world = MPI_COMM_WORLD;
11 MPI_Comm_rank(world, &rank);
12 if (rank==0)
13     other = 1;
14 else
15     other = 0;
16 MPI_Comm_persist(world, &info, SESSION,
17     &persistent)
18
19 if (failure()) {
20     // retrieve seed and iter
21     MPI_Recv(&seed, rank, SEEDTAG, persistent);
22     MPI_Recv(&iter, rank, ITERTAG, persistent);
23 }
24
25 init_data(data, seed);
26
27 if (failure()) {
28     // retrieve boundary conditions
29     MPI_Recv(data[SIZE-1], other, BNDDTAG,
30         persistent);
31 }
32
33
34 for (int i = iter; i<10, i++) {
35     compute(data);
36
37     boundary = data[0];
38     MPI_Sendrecv(&boundary, other, BNDDTAG,
39         data[SIZE-1], other, BNDDTAG,
40         MPI_COMM_WORLD);
41     // store boundary for recovery
42     MPI_Send(&boundary, other, BNDDTAG,
43         persistent);
44
45     seed = aggregate(data);
46     printf("%i_%i_%f\n", rank, i, seed);
47
48     // store state and aggregate for recovery
49     MPI_Send(&seed, rank, SEEDTAG, persistent);
50     MPI_Send(&i, rank, ITERTAG, persistent);
51 }
52
53 printf("Final:_%i_%f", rank, seed);
54 MPI_Finalize();
    
```

Figure 3. A simple MPI program with persistency for 2 processes

one. For simplicity, we will not use this facility for the rest of the paper and use persistent communication explicitly.

The core of a somewhat more elaborated example is shown in Figure 3. For the sake of simplicity, we assume that the application is executed with only two processes. This fictitious algorithm evolves for several iterations a very large array of data through a complex calculation `compute` (line 35). At any given point in time, one can however aggregate the

data into a single value `seed` (line 45). In turn, `seed` can be used to reconstruct the data array with sufficient accuracy by calling `init_data(seed)` (line 25). The algorithm requires to exchange boundary conditions between processes. The first element of the local array is sent to the other process, where it replaces the last element (line 38). The system shall provide a function `failure()` which notifies fault conditions.

The state of the application is given by the iteration counter `i` of the fore loop on line 34. This value is persisted by sending a message to oneself (line 50) at the end of each iteration. The algorithm requires the persistence of the `seed`, again by a message to oneself on line 49. Finally, the exchange of boundary conditions is logged on line 42.

In case of failure, the failed process is restarted and restores its internal state (line 22), the aggregate (line 21) which is used to reconstruct the data array (line 25). The same initialization operation had been executed by the surviving process with initial values at the original start of the application. The failed process also retrieves the last boundary value received from the other process (line 29). Then it enters the main loop with the correctly restored iteration counter and resumes normal operation in parallel to the surviving process.

C. Implementation concerns

Our proposed persistent communicator semantics is relatively easy to implement. As explained in III-A, any given MPI message is uniquely identified by its signature and the sequential ordering. In addition, the user has specified a unique session key at the time of creation of the persistent communicator. Together these are used to store any persistent message content in a suitable stable storage. This could be for instance the memory of one (or several for redundancy) neighbor MPI processes, remote network filesystems, or any data base. In fact, one could persist the data using the SCR library and leave the details to its automatics. After the failure, the application is restarted with the same session key and thus allow to map messages to the state before the fault.

Incoming persistent messages with the same signature, and thus different sequence number, shall overwrite the previously stored one. However, one could also implement a stack of user-defined depth and store a history of messages which are retrieved in order of storage or in reverse. Such schemes could be facilitated by additional parameters provided in the info object at the time of creation of the persistent communicator.

IV. CONCLUSIONS

In this paper we have presented work in progress on the a method to allow persisting of application data and internal state for fault recovery. Unlike other methods, our approach uses well known MPI semantics. The only addition to MPI is a routine that allows to mark a communicator as persistent. All messages to such a communicator are stored on a stable storage for later usage during failure recovery. We have shown basic idioms of storing and retrieving not only application data, but also internal state of the application and to use message logging to recover messages that have been exchanged with other MPI processes just prior to the fault.

ACKNOWLEDGMENT

This work was supported by the German Federal Ministry of Education and Research (BMBF) through the project FEToL (Grant Number 01IH11011F).

REFERENCES

- [1] Y. Ling, J. Mi, and X. Lin, "A variational calculus approach to optimal checkpoint placement," *IEEE Trans. Computers*, vol. 50, no. 7, 2001, pp. 699–708.
- [2] MPI Forum, "MPI: A Message-Passing Interface Standard. Version 3.0," September 21st 2012, available at: <http://www.mpi-forum.org> [retrieved: May, 2015].
- [3] D. J. Sorin, M. M. K. Martin, M. D. Hill, and D. A. Wood, "Safetynet: Improving the availability of shared memory multiprocessors with global checkpoint/recovery," *SIGARCH Comput. Archit. News*, vol. 30, no. 2, May 2002, pp. 123–134.
- [4] J. Duell, P. Hargrove, and E. Roman, "The Design and Implementation of Berkeley Lab's Linux Checkpoint/Restart," *Future Technologies Group*, white paper, 2003.
- [5] J. Cornwell and A. Kongmunvattana, "Efficient system-level remote checkpointing technique for bldr," in *Proceedings of the 2011 Eighth International Conference on Information Technology: New Generations*, ser. ITNG '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 1002–1007.
- [6] CRIU project, "Checkpoint/Restore In Userspace – CRIU," 2015, available at: http://http://www.criu.org/Main_Page/ [retrieved: May, 2015].
- [7] A. Moody, G. Bronevetsky, K. Mohror, and B. R. d. Supinski, "Design, modeling, and evaluation of a scalable multi-level checkpointing system," in *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 1–11.
- [8] K. Mohror, A. Moody, and B. R. de Supinski, "Asynchronous checkpoint migration with mnet in the scalable checkpoint / restart library," in *IEEE/IFIP International Conference on Dependable Systems and Networks Workshops, DSN 2012, Boston, MA, USA, June 25-28, 2012*. IEEE, 2012, pp. 1–6.
- [9] G. E. Fagg et al., "Fault tolerant communication library and applications for high performance," in *Los Alamos Computer Science Institute Symposium, Santa Fe, NM, Oct. 2003*, pp. 27–29.
- [10] E. Gabriel et al., "Open MPI: Goals, concept, and design of a next generation MPI implementation," in *Proceedings of the 11th European PVM/MPI Users' Group Meeting*, ser. LNCS, D. Kranzlmüller, P. Kacsuk, and J. Dongarra, Eds., vol. 3241. Budapest, Hungary: Springer, Sep. 2004, pp. 97–104.
- [11] W. Bland, A. Bouteiller, T. Héroult, J. Hursey, G. Bosilca, and J. J. Dongarra, "An evaluation of user-level failure mitigation support in MPI," in *Recent Advances in the Message Passing Interface - 19th European MPI Users' Group Meeting, EuroMPI 2012, Vienna, Austria, September 23-26, 2012*. *Proceedings*, 2012, pp. 193–203.
- [12] W. Bland, A. Bouteiller, T. Héroult, G. Bosilca, and J. Dongarra, "Post-failure recovery of MPI communication capability: Design and rationale," *IJHPCA*, vol. 27, no. 3, 2013, pp. 244–254.
- [13] J. Hursey, R. Graham, G. Bronevetsky, D. Buntinas, H. Pritchard, and D. Solt, "Run-through stabilization: An mpi proposal for process fault tolerance," in *Recent Advances in the Message Passing Interface*, ser. *Lecture Notes in Computer Science*, Y. Cotronis, A. Danalis, D. Nikolopoulos, and J. Dongarra, Eds. Springer Berlin Heidelberg, 2011, vol. 6960, pp. 329–332.

QoS versus QoE for Database Tuning

Jaroslaw Kurpanik, Malgorzata Pankowska

Department of Informatics

University of Economics in Katowice

Email: jaroslaw.kurpanik@ue.katowice.pl, pank@ue.katowice.pl

Abstract—Quality in Information Communication Technology (ICT) domain is understood as the totality of characteristics of an entity that are important for its ability to satisfy users. In the paper, authors present a consistent set of concepts concerning Quality of Experience (QoE) and Quality of Service (QoS) modelling for information system development. In the last part, they focus on database application tuning and give an example of how the selected QoS measures (i.e., Cost, %CPU and physical read) can be used for quality improvement.

Keywords—QoE, QoS, quality measures, quality management, database tuning.

I. INTRODUCTION

Last years developed servitization is considered as a business strategy to get competitive advantage in IT sector. Servitization as a process is developed within Service Science, which is an interdisciplinary approach to study, create, improve a service [1]. The generic model of services process lifecycle is argued to be derived from the basic phases of a generic business process lifecycle and encompasses part of the possible service outcomes, captured in the Interact-Service-Propose-Agree-Realize (ISPAR) model [2]. In services, values are co-created by user and service provider. Value is determined by results, quality of the customer experience, service price and access costs. Service has a different meanings in different disciplines. In information systems engineering, the service concept is used to separate the external and internal behaviour of a system. In enterprise architecture approach, service is perceived a bidge that different layers of an enterprise [3].

In the paper, servitization is assumed to include all the activities for Information Communication Technology (ICT) services and application quality management. The paper aims to emphasize the relations, differences and applicability of Quality of Experience (QoE) and Quality of Service (QoS) for business application development. The first part of the paper covers explanation of QoE and QoS in the ICT sector. The second part includes discussion on usage of QoS approach for tuning the databases for business.

II. QOS VERSUS QOE

The quality management is to maintain the internal quality of an entity. In order to achieve this, it aims at four subjects, i.e., the product quality, the process quality, the ICT facilities quality and the business organization quality,

mostly aiming at matters like quality of people and their expertise and competencies.

According to Deora et al. [4], quality can be perceived in different aspects, i.e., quality as functionality, as reputation, and as conformance, which is synonymous with meeting specifications. These different views require quality to be monitored and measured differently. Quality as a functionality characterizes, for example, the design of a software and can be measured by comparing the piece of software against others offering similar functionalities. Quality as conformance is monitored for each software individually, and usually requires the user's experience in order to measure the promise and expectations against the delivery. Quality as reputation means that evaluated software is regarded as a reference over time because of its functionalities and conformance qualities. Beyond that quality can be considered as a continuous improvement of non-functional features and can be measures by comparing the modified software against the same piece of software before modification, assuming it keeps the same functionalities. Considering service quality means to focus on the gaps that influence the service quality. At first, the expected service as perceived by the service provider is different from what is expected by the customer. Second, the service specification as used by the service provider differs from the expected service as perceived by the service provider. Third, the actual service delivery differs from the specified services, because of the deficiencies in human resource policies, and failures to match demand and supply. Fourth, communication about the service does not match the actual service delivery because of ineffective management of customer expectations, overpromising, and inadequate horizontal communications. Fifth, the actual service performance differs from the customers' expectations [15].

Service quality as an intangible characteristics is expressed by other qualitative characteristics:

- reliability, understood as the ability of a service operation to be executed within the maximum expected time frame,
- response time to handle the user requests,
- availability, understood as the probability for a service being accessed properly in an observation period,
- integrity, understood as the quality aspect of how the service maintains the correctness of the interaction in respect to the source.

According to ISO 9000:2000, quality is the ability of a set of inherent characteristics of a product, system or process to fulfill the requirements of customers and other interested parties. Nowadays, researchers and practitioners differentiate QoE from QoS. The relationship between these two concepts are formulated in different ways:

- QoE as a counterbalance form for QoS.
- QoE as a contradiction of QoS.
- QoE as a linear combination of QoS and user experience.
- QoE as an interchangeable concept of QoS.

The first approach emphasizes the user's perceptions and experiences. Therefore, the service providers are requested to deliver the software and services not with a high QoS, but with a high QoE to their customers [6].

The QoE concept is particularly important for human-machine interface design, where there is a focus shift from classical usability, understood in terms of effectiveness and efficiency, towards the design of experiences that people have through the use of the interfaces. QoE is defined as the degree of delight or annoyance of a person whose experiencing involves an application, service, or system. It results from the person's evaluation of the fulfilment of their expectations and needs with respect to the utility and/or enjoyment in the light of the person's context, personality and current state [7]. QoE means the overall acceptability of an application or service, as perceived subjectively by the end user and QoE is supplemented by QoS, which means totality of characteristics of a service that bear on its ability to satisfy stated and implied needs of the user of service.

In the second approach, Varela et al. [8] contrast QoE with QoS. According to them, QoS means the totality of characteristics of a telecommunication services, but QoE is not limited to only such services. QoE is defined entirely from the user's perspective. In QoS evaluation process, the focus is mostly on performance of telecommunication systems, while QoE focuses on ICT services and applications. For QoS evaluation, methods are technology oriented, empirical and include simulation results' measurements. QoE evaluation is a multi-disciplinary and multi-methodological approach. In this contradictive process, Weiss et al. [9] add that QoE itself is an important measure for addressing the instantaneous user reaction to changes in the QoS (e.g., types of transmission degradations). Assessment methods of momentary QoE include the measures for truly instantaneous assessment, but also the assessment for short-term samples which exhibit microscopic variation [9].

For Garcia et al. [10], QoE is influenced by user expectations, context and personal preferences. The perception and understanding of the users' QoS are typical for video quality evaluation, but they are going beyond purely measuring video quality. They cover new subjective studies and methodologies for episodic subjective quality evaluation in realistic environments. The quality of the multimedia content as perceived by the user is referred by Timmerer et al. [11] as an example of QoE application. As such, the QoE is defined for a certain modality (e.g., image) or simply for a combination of audio and visual elements and

QoE is targeting two human senses, namely hearing and seeing. They argue that browsing multimedia content may also stimulate other senses such as olfaction, mechanoreception, equilibrioception, or thermoreception. Users are located in a certain environment, where there are other sensory information, e.g., light, wind, vibration, and scent, so the users perceive also these additional sensory effects and even consider them as a part of the particular multimedia content and enrich their personal experience. Therefore, although there is a wide range of factors influencing user experience and QoE, the most important factors can be classified into three main categories; the context around the user, the user's state, and the system properties [7].

Reiter et al., specify ten context categories as follows: personal context, social context concerning interpersonal action and relations, economic context including cost and brand of software applications, event-based context, application based context, historical context, physical context describing the user locations and movements, temporal context, task context, and technical context. The task context concerns the user activities realized in parallel or as a sequence, and interruptions. The technical context describes the relationships among the system of interest and other relevant devices, applications, networks and other informational artefacts [12].

In the third approach, QoS deals with the performance aspects of physical systems. QoE deals with the users' assessment of system performance, as colored by context, culture, the users' expectations with respect to the system or service and their fulfilment, socio-economic issues, and psychological profiles, among other factors.

According to Beyer and Moller [13], QoE is defined as weighted linear combination of QoS and user experience. For Wennerheimer and Robinson, there is no simple correlation of QoE and QoS [14]. However, a generic correlation model between QoS and QoE does not lead to any improvements. QoE is the overall acceptability of an application or service and as such QoE represents user satisfaction. From the software developer point of view, QoS is the collective effect of software usability and service performance. Therefore, QoE covers nontechnical and human aspects of system performance, i.e., smoothness and interaction maintenance effort, response ability, pace, naturalness, comprehensibility effort, accessibility, integrity, retainability, ergonomics, aesthetics, learnability, intuitivity, and intrinsic motivation of user activity.

The aesthetics is connected with ICT devices parameters concerning resolution, colour brightness, noisiness, loudness, and vibration. QoS concerns the usability of functionalities of software, services and devices. Apart from that, QoS properties, such as performance and reliability are increasingly important to ensure user acceptance.

Generally, QoS parameters can be classified as stimulants and destimulants [15]. Stimulants are variables, which values should be increased for the quality improvement, as it is in the case of reliability, compatibility, interoperability, scalability, effectiveness, efficiency, security, accuracy, coherence, and network throughput.

Destimulants are QoS parameters, which values ought to be decreased in the quality improvement process, as it is in the case of metrics including time delay, jitter, interruptions, errors, discontinuity, response time, packet loss rate, latency and time to repair.

QoS is determined in relation to the fulfilment of the service agreement. The measurement of the QoS needs to take into consideration the input from both the customer as well as the service provider, as the service and software development agreement binds both parties. Measuring the quality of software and services is important in modern service-oriented environments. It is reshaping the world of business by providing trusted business processes and the reputation of service providers and software developers. The QoS characteristics are a quantifiable aspect of QoS and they are developed independently of the means by which they are represented and controlled. The values of QoS characteristics include not only numbers (e.g., Booleans, integers, reals, complex numbers), but also vectors, matrices, ranks and names of states.

Particularly in the ICT sector, the subclasses describing QoS characteristics of general importance to communication and processing are presented by the following characteristics groups: time-related, coherence, capacity, integrity, safety, security. QoS characteristics cover not only metrics, but also policies, which make a concrete software useful.

The policies include scheduling and dispatching policies, admission control of requests, levels of services supported in the server, and management of service requests.

In the last, fourth approach, QoE and QoS terminology are used interchangeably [16]. QoE is the overall performance of a system from the point of view of the users. QoE is a measure of end-to-end performance at the service level from the user perspective and an indication of how well the system meets the user's needs. QoS is a measure of performance from the network perspective. QoS refers to a set of technologies that enable the network administrator to manage the effects of congestion on application performance as well as providing differentiated services to selected network traffic flows or to selected users. Quality management in ICT sector supported by QoS characteristics requires establishing, monitoring and maintenance of the actual QoS characteristics, controlling of the QoS targets, alerting as a result of some events relating to the QoS management [17].

III. QoE APPLICATION EXAMPLE

Top management and company decision makers have an opportunity to measure the quality of ICT services with Information Technology Service Management (ITSM) system. Therefore, they are able to manage business processes and technologies. ITSM system enables ICT providers to adhere to the conditions included in the Service Level Agreement (SLA), expressed by Key Performance Indicators (KPIs). Usually, the SLA must contain quantitative measures that represent a desired and mutually agreed state of a service, provide additional boundaries of a service scope, and describe guaranteed minimal service performance [18].

Furthermore, KPIs are metrics that are used to measure efficiency and effectiveness of a service, and service operation status. The SLA examples would cover, for a particular customer, the reaction time, resolution time, compliance to agreed deadlines. KPIs examples could include average reaction time by users, incoming ticket volume trend, time of system inaccessibility, number of solved problems in an established time, number of correctly provided services in relation to the number of all services.

However, service catalogue is constructed in the context of particular business goals' specification. ICT service design does not concern the whole system in the same way. The quality parameters may be not defined for database, application servers or networks. The quality parameters are defined for particular levels of the information system. They are analysed and maintained inside the ICT department, which focuses on the new solution development, as well as on the existing functionalities extension. However, this approach is not appropriate for business, which requires the effects important for the customers. Therefore, the accountability of ICT services should be based on SLA parameters.

For example, in an information system for sale and distribution, number of cases of exceeding the accepted time (i.e., 2 seconds) for registration of one item in an order during a day. User receives a phone request and inserts data, which are dictated by their customer. Average number of items dictated through the phone covers about 70 elements per one customer request. Request reception time is critical. After the registration of each item in request, information systems verifies product accessibility, controls the product promotions and calculates the price and margins. Information system is to measure the time between product approval and the moment of the presentation of all data on the terminal screen. In case of exceeding the time, the information is registered in log file. Business user accepts no more than 50 cases, where the time of 2 seconds is exceeded. Logs are constantly monitored and in a situation of increasing the number of such cases, suitable actions are undertaken to return to the acceptable level of service. Similar time and acceptance levels are defined by inventory division users working with mobile devices.

IV. QoS APPLICATION EXAMPLE

The consideration on the quality measures refers to the Oracle database efficiency. Efficiency analysis and database tuning requires multidimensional approach, because of the number of influencing variables. For Oracle applications, the database tuning is considered as:

- Good database design and distribution of the database workload across multiple disks to avoid or reduce disk overloading.
- Disk I/O optimization, which is related directly to throughput and scalability.
- Checkpointing to increase the I/O activity and system resource usage.

- Batching multiple operations together and/or increasing the number of operations that run in parallel (increasing concurrency) [19].

In this paper, authors focus on operations, which are important from the point of view of database administrator and which have an impact on database system efficiency and on the speed of query performance. Particularly, they consider input/output (I/O) operations and the number of physical disc readings realized by a database system to find the answer to the user query. If that number is smaller, the waiting time for query results is shorter. The number can be optimized in different ways, for example by suitable index creation [12]. A query execution plan is an ordered set of steps used to access data in an SQL relational database management system.

Query optimizer attempts to determine the most efficient way to execute a given query by considering the possible query plans [20]. They can be rule-based or cost-based. In the first approach, the rule-based optimizer (RBO) chooses an execution plan based on the rules and access paths available and the ranks of these access paths [21].

Oracle's ranking of the access paths is heuristic. If there is more than one way to execute an SQL statement, then the RBO always uses the operation with the lowest rank. Usually, operations of lower rank execute faster than those associated with constructs of higher rank. Cost-based optimizer (CBO) generates a set of potential execution plans for SQL statements, estimates the cost of each plan, calls the plan generator to generate the plan, compares the costs, and chooses the plan with the lowest cost. In this approach the query performance time is estimated for different options and the query for which the time is the shortest is selected [22].

In Oracle, the cost is defined as units of work or resource used. The query optimizer uses disk I/O, CPU usage, and memory usage as units of work. So, the cost used by the query optimizer represents an estimate of the number of disk I/Os and the amount of CPU and memory used in performing an operation. The operation can be scanning a table, accessing rows from a table by using an index, joining two tables together, or sorting a row set [23]. The cost of a query plan is the number of work units that are expected to be incurred when the query is executed and its result produced. The cost of a table scan or a fast full index scan depends on the number of blocks to be scanned and the multiblock read count value. The cost of an index scan depends on the levels in the B-tree, the number of index leaf blocks to be scanned, and the number of rows to be fetched using the rowid in the index keys [7] [18]. Oracle proposes many different indexes applying for minimizing the query performance cost.

Indexes are optional structures associated with tables and clusters. Indexes can be created on one or more columns of a table to speed up SQL statement execution on the table. Indexes are the primary means of reducing disk I/O when properly used. Oracle provides several indexing schemes, e.g.:

- Normal indexes, by default, Oracle Database creates B-tree indexes.
- Reverse indexes.

- Bitmap indexes.

The purpose of an index is to provide pointers to the rows in a table that contain a given key value. In a regular index, this is achieved by storing a list of rowids for each key corresponding to the rows with that key value.

Oracle Database stores each key value repeatedly with each stored rowid. In a bitmap index, a bitmap for each key value is used instead of a list of rowids [24]. Creating reverse key index, compared to a standard index, reverses the bytes of each column indexed (except the rowid) while keeping the column order [11].

In this example, the sh.customers table was used for the verification of impact of indexes applying on the cost of query performance. Before the decision on index creation, certain statistics for sh.customers table were calculated and presented in Table 1.

TABLE I. SEGMENT OF COLLECTED COLUMN STATISTICS

COLUMN_NAME	DENSITY	NUM_NULLS
CUST_CITY_ID	0,00217939065	0
CUST_TOTAL	1,00000000000	0
CUST_STATE_PROVINCE	0,00689655172	0

Table 1 consists of three columns, where the value in "DENSITY" column provides information about the selectiveness of the column. Lower value indicates that the column is a better candidate to create index. Value included in "NUM_NULLS" column indicates how many "null" values are included in the column [11] [18]. The "DENSITY" value for the "CUST_TOTAL" column is equal to 1, so the column contains only one unique value, therefore creating an index on this column is not necessary, because the full table scan can be efficient. For the two other columns, the selectiveness of data is relatively high. Therefore, the columns are good candidates for indexing.

TABLE II. STATISTICS IN THE "WITHOUT INDEX (FULL TABLE SCAN)" CASE

Without index (Full table scan)						
Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		66	990	405 (1)	00:00:01
1	TABLE ACCESS FULL	CUSTO -MERS	66	990	405(1)	00:00:01
28 physical read IO requests 11919360 physical read bytes 28 physical read total IO requests 11919360 physical read total bytes 11 physical read total multi block requests 1455 physical reads 1455 physical reads cache 1427 physical reads prefetch						

TABLE III. STATISTICS IN THE "NORMAL INDEX" CASE

Normal Index						
<i>Id</i>	<i>Operation</i>	<i>Name</i>	<i>Rows</i>	<i>Bytes</i>	<i>Cost (%CPU)</i>	<i>Time</i>
0	SELECT STATEMENT		66	990	24 (0)	00:00:01
1	TABLE ACCESS FULL	CUSTO -MERS	66	990	24 (0)	00:00:01
2	INDEX RANGE SCAN	TEST_ INDEX	23		2(0)	00:00:01
1 physical read IO requests 8192 physical read bytes 1 physical read total IO requests 8192 physical read total bytes 1 physical reads 1 physical reads cache						

TABLE IV. STATISTICS IN THE "REVERSE INDEX" CASE

Reverse Index						
<i>Id</i>	<i>Operation</i>	<i>Name</i>	<i>Rows</i>	<i>Bytes</i>	<i>Cost (%CPU)</i>	<i>Time</i>
0	SELECT STATEMENT		66	990	4543 (1)	00:00:01
1	TABLE ACCESS BY INDEX ROWID	CUSTO -MERS	66	990	4543 (0)	00:00:01
2	INDEX FULL SCAN	TEST_ INDEX	4751		125 (0)	00:00:01
24 physical read IO requests 101580 physical read bytes 14 physical read total IO requests 1015808 physical read total bytes 124 physical reads 124 physical reads cache						

TABLE V. STATISTICS IN THE "BITMAP INDEX" CASE

Bitmap Index						
<i>Id</i>	<i>Operation</i>	<i>Name</i>	<i>Rows</i>	<i>Bytes</i>	<i>Cost (%CPU)</i>	<i>Time</i>
0	SELECT STATEMENT		66	990	17 (0)	00:00:01
1	TABLE ACCESS BY INDEX ROWID	CUSTO -MERS	66	990	17 (0)	00:00:01
2	BITMAP CONVERSION TO ROWIDS					
3	BITMAP INDEX RANGE SCAN	TEST_ INDEX				
160 physical read IO requests 10485760 physical read bytes 160 physical read total IO requests 10485760 physical read total bytes 1280 physical reads 1280 physical reads cache 1120 physical reads cache prefetch 1120 physical reads prefetch warmup						

Tables 2, 3, 4, and 5 cover the plan of performance of the query: select cust_city from sh.customers where cust_city_id is between 52379 and 52383.

Application of indexes influences an increase of central processor unit performance, as well as decreases physical read, total IO requests and cost of operations.

In the "Without index (Full table scan)" case, the Cost (%CPU) is smaller than in the "Reverse index" case, because in this first option, the multiblock reading

operations are realized (Table 6). These operations are more promptly than reading operations in the other options.

TABLE VI. ANALYSIS OF INDEX APPLICATION FOR EFFICIENCY IMPROVEMENT

Different Index Types					
<i>Parameter</i>	<i>Without Index (Full table scan)</i>	<i>Normal Index</i>	<i>Reverse Index</i>	<i>Bitmap Index</i>	<i>Efficiency increase in the best option selection situation</i>
Cost (%CPU)	405	24	4543	17	16
Physical read	1455	1	125	1280	1455

V. DISCUSSION AND CONCLUSION

Taking into account the contemporary standards and professional studies some conclusions on the relations among QoS and QoE can be formulated. However, firstly, different interpretations of QoE and QoS are developed and widely accepted. The paper aimed to present that categories belong to two different ICT system development environments. QoE focuses on user perception and satisfaction, therefore, QoE metrics should be proposed, specified, evaluated and estimated by end users. Usually, application developers discuss the acceptable metrics and metrics values with end users. The discussions are not a one-time event, but they are repeatable and these metrics ought to be included in the documentation and contract among ICT solution providers and users. On the other hand, the QoS frameworks were originally developed for network management. However, taking into account the ISO/IEC 13236 standard, the term "service" is understood in a very general sense, therefore the wide application of QoS framework is possible [14]. According to that standard, it includes, among others:

- The provision of processing and information repository functions by entities, objects, applications and application processes.
- The interactions between entities, objects, and applications.
- The confidentiality and lifetime characteristics applied to information stored in the information system.
- The communication services.
- The physical equipment.

In the ISO/IEC 13236 standard, a QoS characteristic represents some aspect of a system, service or resource that can be identified and quantified. The QoS parameters are usually, but in unique way, specified in the SLA. These parameters are of different kinds, including:

- A desired level of characteristics, specified by system developers.
- A minimum or maximum level of a characteristic.
- A measured value, used to convey historical information.
- A threshold level.
- A warning or a signal to take corrective actions.

- A request for operations on managed objects relating to QoS.

Taking into account the above mentioned specification of parameter kinds it is difficult to define suitable parameters for database quality improvement. However, the general important QoS characteristics in the ISO/IEC 13236 standard are grouped in the following subclasses:

- Time-related characteristics.
- Coherence characteristics.
- Capacity-related characteristics.
- Integrity-related characteristics.
- Safety-related characteristics.
- Security-related characteristics.
- Reliability-related characteristics.

Within the capacity-related characteristics, capacity, throughput, processing capacity and operation loading are considered. In ISO/IEC 13236, the application information throughput characteristic is defined as an amount of data transferred between applications in a period of time and it is quantified as a rate expressed in bits/second or bytes/second. Beyond that the system throughput characteristic is defined as an amount of processing performed in a period of time. In this paper, authors proposed two metrics, i.e., cost (%CPU) and physical read for database throughput improvement. The measurements are strongly important only for database application developers and not visible to the end users. In the future, authors propose to further develop the system quality measurement from the different ICT system stakeholder point of view.

REFERENCES

- [1] Y. Lin, Y. Shi, and S. Ma, "Servitization Strategy: Priorities, Capabilities, and Organizational Features," in *Service Science, Management and Engineering, Theory and Applications*, G. Xiong, Z. Liu, X. Lu, F. Zhu, and D. Shen, Eds. Amsterdam, Elsevier, pp.11-34, 2012.
- [2] P.P. Maglio, S.L. Vargo, N. Caswell, and J. Spohrer, "The service system is the basic abstraction of service science," *Information Systems and e-business Management* 7, 2009, pp.395-406.
- [3] G. R. Gangadhran and E. Field, "Analyzing Requirements and Approaches for Sourcing Software Based Services," *Service Science and Logistics Informatics*, Luo Zongwei, Ed. Hershey, New York, Information Science Reference, IGI Global, pp. 195-205, 2010.
- [4] V. Deora, J. Shao, W. A. Gray, and N. J. Fiddian, "A Quality of Service Management Framework Based on User Expectations," in *Service Oriented Computing*, M.E. Orłowska, S.Weerawarana, M.P. Papazoglou, and J. Yang, Eds. Berlin, Springer-Verlag, pp. 104-114, 2003.
- [5] C. Mendes and M. Mira da Silva, "DEMO-Based Service Level Agreements," in *Exploring Services Science*. M.Snene, Ed. Belin, Springer, pp.227-242, 2012.
- [6] A. Raake and S. Egger, "Quality and Quality of Experience," In *Quality of Experience*, S.Moller and A.Raake, Eds. Heidelberg, Springer, pp. 11-35, 2014.
- [7] Qualinet White Paper on Definitions of Quality of Experience-Output Version of the Dagstuhl Seminar 12181. [Online] Available from: <https://tel.archives-ouvertes.fr/IRCCYN-IVC/hal-00977812v1> 2015.01.16
- [8] M. Varela, L. Skorin-Kapov, and T. Ebrahimi, "Quality of Service Versus Quality of Experience," in *Quality of Experience*, S.Moller and A.Raake, Eds. Heidelberg, Springer, pp. 85-97, 2014.
- [9] B. Weiss, D. Guse, S. Moller, A. Raake, A. Borowiak and U. Reiter, "Temporal Development of Quality of Experience," in *Quality of Experience*, S. Moller and A.Raake, Eds. Heidelberg, Springer, pp. 133-148, 2014.
- [10] M. Garcia, S. Argyropoulos, N. Staelens, M. Naccari, M. Rios-Quintero, and A. Raake, "Video Streaming," in *Quality of Experience*, S.Moller and A.Raake, Eds. Heidelberg, Springer, pp. 277-298, 2014.
- [11] Ch. Timmerer, M. Walzl, B. Rainer, and N. Murray, "Sensory Experience: Quality of Experience Beyond Audio-Visual," in *Quality of Experience*, S.Moller and A.Raake, Eds. Heidelberg, Springer, pp. 351-366, 2014.
- [12] U. Reiter, K. Brunnstrom, K. De Moor, M. Larabi, M. Pereira, A. Pinheiro, J. You, and A. Zgank, "Factors Influencing Quality of Experience," in *Quality of Experience*, S.Moller and A.Raake, Eds. Heidelberg, Springer, pp. 55-73, 2014.
- [13] J. Beyer and S. Moller, "Gaming," in *Quality of Experience*, S.Moller and A.Raake, Eds. Heidelberg, Springer, pp. 367-382, 2014.
- [14] M. Wennesheimer and D. Robinson, "Service Quality Definition and Measurement, A Technical Report," NGMN Alliance, Next Generation Mobile Networks Ltd. Frankfurt, 2013.
- [15] *Information Technology - Quality of service: Framework*, International Standard, ISO/IEC 13236, ISO, IEC, Geneva, 1998.
- [16] DSL Forum TR-126. Triple-play Services Quality of Experience (QoE). [Online] Available from: <http://www.broadband-forum.org/technical/download/TR-126.pdf> 2015.01.15.
- [17] A. T. Campbell, "Quality of Service Architecture," Computing Department, Lancaster University, Doctor thesis, 1996. [Online]. Available from: <http://www.cs.dartmouth.edu/~campbell/papers/thesis.pdf> 2015. 01. 15.
- [18] User Experience White Paper Bringing clarity to the concept of user experience, Result from Dagstuhl Seminar on Demarcating User Experience, September 15-18, 2010, February 11, 2011. [Online] Available from: <http://www.allaboutux.org/files/UX-WhitePaper.pdf> 2015.01.16.
- [19] Oracle Database Tuning, WebLogic Server Performance and Tuning, 2014. [Online]. Available from: http://docs.oracle.com/cd/E13222_01/wls/docs100/performance/dbtune.html 2014.12.22.
- [20] M. Wojciechowski and M. Zakrzewicz, "Cost Optimizer of Queries" Instytut Informatyki, Politechnika Poznańska, 2015. [Online] Available from: http://www.ploug.org.pl/seminarium/seminarium_III/materialy/kosztowy.pdf, 2015.01.18
- [21] B. Bryla and K. Loney, "Database Administration Guide," Helion, Gliwice, 2010.
- [22] *Database Performance Tuning Guide and Reference*, 2002. [Online] Available from: https://docs.oracle.com/cd/B10501_01/server.920/a96533.pdf , 19.01.2015
- [23] *Database SQL Language Reference*, 2015. [Online] Available from: http://docs.oracle.com/cd/B28359_01/server.111/b28286/statements_5011.htm#SQLRF01209 , 18.01.2015
- [24] Programmer Interview, 2015. [Online] Available from: <http://www.programmerinterview.com/index.php/database-sql/cardinality-versus-selectivity/>, 10.01.2015

Integrated Tool Support in the Context of Security and Network Management Convergence

Felix von Eye, David Schmitz, and Wolfgang Hommel
 Leibniz Supercomputing Centre, Munich Network Management Team
 Garching n. Munich, Germany
 Email: {voneye,schmitz,hommel}@lrz.de

Abstract—The convergence of different types of networks, such as for telecommunication and data transfer, on the hardware layer has an obvious impact on both network management and security management, which especially affects network service providers as well as data centers. This paper argues that also methods, algorithms, and tools from both research domains, network management and security management, should systematically be reviewed for synergies. It first analyzes the current state of the art in both domains and identifies gap areas that require further investigation. Then, the Customer Network Management for the X-WiN (WebCNM) network management tool, which started as a research prototype in the pan-European research and education network, GÉANT, is presented along with selected extensions that were designed and implemented to integrate security management functionality. Several security event visualization options and their use within the European industry-focused Safe And Secure European Routing (SASER) project are discussed.

Keywords—Network management; Security management; Integrated management; Enterprise management; Convergence.

I. INTRODUCTION

In both research and real-world operations, network management and security management are often treated as complementary but still largely independent parts of the overall IT service management performed by network service providers and data centers. However, this view limits the prospects in case of, e.g., handling incidents. In many situations both security events and network events are related and visible within the same time interval, for example when a security event results in an unusual or peak usage of the network. In turn, certain conditions, such as an unreachable service or device, can be caused either by a denial of service attack or a fault in the network.

As security management and network management are typically supported by independent technical management software tools, such contextual relationships are not visible without further ado. This motivates a new approach to combine the information processed by either management tool landscape into one, leading to a convergence of selected network and security management methods and tools. Although it is obvious that it is neither reasonable nor practically possible to fully integrate both disciplines currently, we elaborate on selected event categories that result in operational benefits or have interesting properties for research that motivate further investigation.

The work leading to the presented results have been carried out in the large-scale distributed environment of the SASER-SIEGFRIED project (Safe and Secure European Routing) [1], in which more than 50 European research and industry project

partners design and implement network architectures and technologies for secure future networks. The project's overall goal is to remedy security vulnerabilities of today's IP layer networks and have them ready for deployment in backbone networks by the year 2020. Thereby, security mechanisms are designed based on an analysis of the currently predominant security problems in the IP layer as well as upcoming issues, such as vendor backdoor and traffic anomaly detection. The project focuses on inter-domain network traffic and routing decisions that are based on security metrics, which are derived from aggregating and combining security measurements carried out by multiple involved organizations in a cooperative manner. As this scenario relies on the future internet architecture and software-defined networks (SDNs), we present a prototype of a combined security and network management system, which is independent of any commercial vendor or operator and has already been used in the context of GÉANT, the pan-European research network.

The remainder of this paper is structured as follows: Section II gives an overview of the current state of the art and related work regarding network management, security management, and the convergence of these two disciplines. Section III presents the WebCNM framework, which serves as technical basis for implementing an integrated network and security management platform, and experiences made in the SASER project. Finally, Section IV concludes the paper and gives an outlook to future work.

II. STATE OF THE ART

Both network management and security management are computer science disciplines with a very long, yet only partially overlapping history. In this section, first the focus is clarified by discussing the term *management*. Then, we analyze the status quo of network management in Section II-A and the current state of the art of security management in Section II-B. Finally, Section II-C reviews related work on the convergence of both disciplines and outlines how it has influenced the improved approach presented in Section III.

In general, *management* in the context of IT services refers to any measures and activities that are performed in order to achieve effective and efficient operations of those IT services and the required resources in alignment with an organization's business goals [2]. Management therefore covers the whole life-cycle, including planning, provisioning, setup, configuration, operations and maintenance, and removal; it involves personnel, procedures, processes, technology, and software tools. Management must be performed on any abstraction layer, such as individual physical hardware components or software

properties (e.g., application response time) up to an enterprise-wide and even inter-organizational view. The term *integrated management* refers to approaches that successfully deal with heterogeneity, such as managing hardware by different vendors or across various types of systems, such as network components, servers, and application software. Usually, *management architectures* describe various properties of how management is carried out in an abstract manner. They can be broken down into four models:

- 1) The systems to be managed, referred to as *management objects*, are described by an *information model*.
- 2) The roles of all systems involved in management along with their types of cooperation are described by an *organizational model*.
- 3) The *communication model* describes the exchange of management-related messages between the roles defined by the organizational model.
- 4) The *function model* groups and structures the management-specific functionality.

A software implementation of a *management architecture* is referred to as a *management platform*. Furthermore, a *management system* refers to the sum of organizational structure, policies, planning activities, responsibilities, practices, procedures, processes, and resources [3]; i.e., a management system is not just a piece of software, but management software, such as a *management platform* and various *management tools*, contributes to this overall management system.

A. Network Management Status quo

Network management has evolved over decades with the growth of intra-organizational enterprise networks and the Internet. The most influential foundation for network management is the classical OSI management architecture, which proposed the five functional areas fault, configuration, accounting, performance, and security (FCAPS) [4].

Fault management, which is closely related to the IT service management areas of incident management and problem management as used by the standard ISO/IEC 20000, deals with anticipating, detecting, and reacting to any types of network faults, such as hardware defects, resource exhaustion, and quality of service guarantee violation. To a large degree, fault management is based on monitoring one's network using both active and passive measurements. Inter-domain fault management is still non-trivial due to the heterogeneity of the involved hardware and organizations' restrictive information sharing policies [5].

Configuration management is the function area to actively modify a network component's parameters. It is an area for which integrated management has not been established successfully because of the inherent complexity of the task:

- Various different types of network components, such as routers, switches, and WiFi access points must be supported; they operate on different layers of the ISO/OSI model.
- Hardware and firmware heterogeneity: The same type of network component, such as a router, often has vendor-specific configuration options and still shows a lack of compatibility when components from different vendors or product generations are mixed.

- Scalability: Larger organizations typically must manage several thousands of network components; however, many management tools still lack support for parallelized operations, e.g., to perform even simple tasks like firmware updates.
- Emerging technologies: Given the cost of network equipment, it is not unusual that network components are in use much longer than other types of hardware. For example, while modern network equipment supports new management paradigms, such as SDNs, the majority of components in use today still needs to be managed using Simple Network Management Protocol (SNMP).

As a consequence, network component configuration management is mostly done by using vendor-specific tools in practice, which means that organizations need other means to ensure configuration consistency across network component types and hardware models.

Accounting, like fault management, is based on monitoring network components. Various low-level hardware counters are aggregated and combined to support, e.g., billing processes.

Performance management is closely related to the management of service levels. Quality of service parameters typically include the minimum guaranteed availability of network links along with their bandwidth and upper limits for undesired properties, such as packet loss, delay, and jitter. Similarly to accounting, these parameters may vary with the actual content that is being transported, because, for example, voice-over-IP connections have different requirements than bulk data transfers.

Security management, as far as the traditional network management is concerned, focuses on security properties of individual network components, such as authentication and authorization for management access. While most modern network components support some additional security features, such as port-based access control in switches and access control lists in routers, these features either depend on additional central security management components, such as a RADIUS server for user or device authentication, or are clearly limited by their hardware performance.

B. Security Management Status quo

The classic goal of security management is to ensure the confidentiality, integrity, and availability of sensitive data as well as the systems and services that process it. As in other areas of computer science, there is no single silver bullet security measure to achieve this goal, but numerous modular security controls must be combined in a complementary and partially deliberately overlapping manner, which is often referred to defense-in-depth and graceful-degradation. Security controls can be categorized as *preventing*, *detecting*, or *reacting* in relation to successful attacks.

Derived from the work by Hyland and Sandhu [6], the following areas of security operations must be considered:

- *System security* covers the management of available software updates, hardening each individual system, having malware protection in place, integration with central security services, such as authentication servers, and preventing undesired data leakage.

Network components are also considered to be such systems.

- *Network security* encompasses the definition of security zones based on protection requirements, the application of virtualization and segregation concepts, such as virtual local area networks (VLAN), and dedicated network security components, such as firewalls and network-based intrusion detection systems.
- *Access management* covers the management of users along with their roles and permissions as well as setting up authentication and authorization services. Privileged accounts, i.e., those used by administrators with full control over a system, deserve special attention. Physical access to systems also must be considered.
- *(High) availability* is practically only achieved through some type of redundancy, which includes both hardware and copies of the data.
- *Cryptography* is a complex discipline with endless applications in information security. Flawed implementations and improper application due to lacking know-how are also a large source of major security problems.

Preventive security controls, such as firewalls and access control mechanisms, are intended to enforce policies, i.e., someone must define what is allowed or undesired in an a-priori manner. The proper implementation of these controls can be checked, e.g., by means of penetration tests. The safe general assumption is that at least parts of one's IT infrastructure are insecure and the methods provided by the subdiscipline of vulnerability management assists in identifying those parts. To prioritize options for the improvement of the overall security level, risk management methods need to be applied.

However, security monitoring is quite mature. It mostly relies on aggregating, correlating, and evaluating security events provided by various dedicated sensors as well as systems and applications, e.g., via log files. Intrusion detection systems passively monitor data as it is processed by the involved systems and use signatures of known attacks or outlier detection to register unusual behavior. As most attack attempts are background noise, i.e., they are not specifically targeted against systems known to be vulnerable, correlation with an asset management database and information from one's own vulnerability management greatly help to reduce the number of false positive alerts. Security information and event management (SIEM) systems perform this correlation, prioritize detected security incidents, create reports, and can be considered to be the security management counterpart of management platforms in network management.

C. Related Work on Management Convergence

It is obvious that network management and security management have partially overlapping scopes: Certain network components, such as firewalls, are typically operated by IT security personnel, and security events generated by network components are being processed by security management tools as a matter of course. However, this information flow can traditionally be considered one-way, from network management to security management tools, without feedback loops and with completely separated tool sets used.

Convergence of both network and security management has the meaning that methods, procedures, and tools can jointly be used for both disciplines, bringing them together as one to increase effectiveness and reduce the overhead caused by separate processing of the same data in multiple instances. In this section, related work in this area is analyzed, which has influenced the design of the approach presented in the next section.

In [7], Dawkins et al. presented a novel network security management system for tracking large-scale, multi-step attacks based on data from various specific sensors in different networking domains. Their system provides real-time correlation and analysis of the data; they focus on an interesting novel visualization method that provides a kind of heads-up cockpit display of an entire network that can be used by network as well as security management personnel. While their work is intended for converged networks and sensors specifically designed for such environments, our approach differs in that it makes use of existing data sources.

In [8], Kuklinski and Chemouil discuss management challenges specifically for SDNs. They propose a mapping of the classic FCAPS approach and point out the special role of SDN controllers, for which additional security monitoring mechanisms are proposed. Complementary, Zhu et al. discuss the role of vendor specifications in cross domain communication in [9]. They present key requirements for an inter-domain security infrastructure along with a reference architecture.

Han and Lei compare the policy languages that are used for network and security management in [10]. The identified similarities make it interesting to formulate policies and rule sets that cover both management disciplines. However, despite some widely used policy languages, the use of proprietary policy specification formats in various products still fuels the demand for inter-system conversions. Wang-fei and Qi developed a novel network management system in [11] with an emphasis on security management that covers networking equipment as well as virtual machine servers; it also addresses combined performance management and unifies the access control mechanisms of both network components and IT services.

In sum, despite several approaches towards combining both management disciplines in prior work, there is no thorough analysis of which types of network and security events would be relevant for a unified management approach and no integrated management platform exists yet. This motivates the approach discussed in the next section.

III. EXTENDING THE WEBCNM FRAMEWORK

WebCNM started as a research prototype for vendor-independent, multi-domain, and customer-oriented network monitoring and management visualization. Its core functionality consists of the visualization of network maps that are organized in a tree-based hierarchy. Each network map shows network elements, i.e., nodes and links of different network layers, together with current or historic status and metric information. Detailed historic statistics are provided in a drill-down manner.

WebCNM was originally developed in the pan-European research and education network, GÉANT, as part of the GN3 project (2011-2013). It has been used to visualize

network topologies and respective network metrics of many European and Non-European national research and education networks (NRENs) including DFN (Germany), SWITCH (Switzerland), Uninett (Norway), Pionier (Poland), GARR (Italy), Surfnet (Netherlands), Renater (France), Hungarnet (Hungary), MRen (Montenegro), GRNet (Greece), SEREEN (South-East-Europe), ESNet (Energie Science Network, US), Internet2 (US), RedClara (Pan South/Middle America), RNP (Brazil), and the pan-European core-network of GÉANT as well as the optical private network of CERN's Large Hadron Collider LHC project (LHCOPN). A significantly extended version is in production as of 2015 for all higher education institutions connected to the German NREN. WebCNM's functionalities comprise, among others, customers' network service information, network access information/status, network access accounting, network core status, and network performance management.

WebCNM was specifically designed to be extensible in a flexible and modular manner: It features a JavaScript extension API, which is independent of the implementation technology of both the backend and the GWT-based web client, which allows for a client-side integration with other web pages and web tools. It is therefore a suitable basis for integrating functionality related to security management.

A. Integration of Security Events in WebCNM

Before the integration of security events into any network management tool, it is necessary to analyze which potential security events have an impact on network connections or devices. As a first step, only direct attacks on network devices, attacks on the network protocols or on the availability of systems, services, and the network itself are considered.

As there are a lot of possible attack scenarios, the following implemented examples highlight some core issues:

- *Port scans*: A port scan is often treated as an attack in which a client attempts connections to a range of server ports with the goal of finding active services and preparing the exploitation of their known vulnerabilities.
- *Border Gateway Protocol (BGP) Hijacking*: BGP hijacking is the illegitimate takeover of groups of IP addresses by corrupting inter-domain routing tables. This attack has a very high impact on network security but it is difficult to implement as the attacker has to first compromise central network infrastructure components, such as the backbone routers.
- *Amplification Attack*: In distributed reflective denial-of-service (DRDoS) attacks, adversaries send requests to public servers (e.g., open recursive DNS resolvers) and spoof the IP address of a victim. These servers, in turn, flood the victim with valid responses and – unknowingly – exhaust its bandwidth.
- *Backdoor*: A backdoor in a computer system is a method of bypassing security controls, such as user authentication, in order to enable unsolicited remote access to a system, obtaining access to data, and so on, while attempting to remain undetected. While backdoors on servers and workstations are quite usual and a well researched topic, backdoors on switches and routers are still a huge problem.

Out of these threat scenarios, denial-of-service attacks are the most simple example of security incidents that are also visible in network management because the bandwidth utilization parameter increases. To enable the correlation between network and security events, WebCNM was extended with an interface to use intrusion detection message exchange format (IDMEF) based messages [12], which are designed to exchange security relevant data between systems and domains.

B. Visualization of logins on network devices

In this paper, an example is used to clarify the ideas behind the new approach, in which logins to network devices like backbone routers are analyzed. Regular secure shell (SSH) logs are used as basis, which gives the possibility to detect whether a suspicious login has occurred.

In WebCNM, any network device can be monitored by the network management tool part. There are two ways to determine if there is a security incident. The first one is the alerting by the SIEM system via the IDMEF interface. In this case, the SIEM system sends an IDMEF message as a notification to WebCNM and this event is then displayed inside the network management tool. This method is usually used to inform the network administrators about suspicious events in the network respectively on network devices. In general, in a regular sized network there are too many events, so they can be only analyzed by visualization. In this case, the affected subnetwork or network device is highlighted inside the network management system, which helps to find possible network issues.

Inside the SIEM system, the security administrators usually use a lot of different rulesets, which are able to automatically detect attacks or misuse of components inside the network. These rulesets have to be adjusted to produce only a few false positives and negatives. To detect unknown attacks or to detect events, which are slightly under the radar, it is necessary to also manually analyze the communication flows.

This leads to the other way to display security-relevant events during the direct inspection of the network components. For example, the log messages of a network device can be displayed in the network management tool. As the network management tool is not a SIEM system, it is not able to process the correlation of events in an automated manner, but instead can only visualize them for the administrator. This allows network administrators to see suspicious events, e.g., a login to the network device from outside the management VLAN, which is quite unusual. Figure 1 illustrates an example of a possible visualization of SSH logins that has been developed within the SASER project: It uses GeoIP-based grouping of remote source addresses and makes it easy to visually distinguish between successful and failed logins. Brute force attempts to guess user passwords can be spotted on the right, and clicking on any edge or vertex brings up more details about the selected group of events. When viewing more generic events, advanced color schemes can be used, e.g., to visually identify protocols or VLANs. In this scenario, the server is only accessible via the SSH public key method for the administrators, so it is unusual that there are also connections authenticated with passwords. In other cases this behavior is completely different, as the authentication via password is the only possible way to access a device. This diversity of authentication methods prevents a fully automated generation of SIEM rules, e.g., for anomaly




Figure 1. Visualization of SSH logs in WebCNM.

detection, but instead needs manual classification. This can be done very user-friendly via visualization.

As the information collected in such a way is only some type of metrics from the WebCNM point of view, the used visualization is not limited to the form shown in Figure 1. Figure 2 shows some visualizations that were developed together with the user interface division of FH Potsdam in the SASER project. Especially for data for which no automated processing and alerting rules have been implemented, these visualizations exploit the innate human skill to quickly identify patterns and outliers that do not match one of these patterns; therefore, while many of those developed visualizations do not provide explicitly the necessary information to handle the potential incidents they indicate, they nevertheless assist in quickly picking up those events manually that have to be reviewed further.

These different visualization types were developed, as there are no visualization types generally applicable to all attack scenarios or systems. The benefit of the human analysis of the visualizations is that humans can browse through different visualization views. If there are any outliers visible, they are able to correlate them with other outliers in other views enhanced with their knowledge. Therefore, it is important that there are ways to switch the kind of visualization. The visualizations shown in Figure 2 are specially made for detection of amplification attacks, denial of service detection, portscan detection, the correlation between source and destination addresses in login events, the delays of Transmission Control Protocol (TCP) connections, and the detection of high traffic network nodes. Often, the design goal of visualizations is to highlight differences or similarities of the received datasets. This helps the network administrators to detect security events before the attacker becomes visible to the security administrators and their SIEM system.

C. Evaluation in the SASER Scenario

In the SASER project, the newly designed network management system was introduced to allow the integrated view on network and security. The implementation was tested by

several partners to identify whether the concept also works in real-world scenarios. It turned out in these tests, however, that the focus on port scans is not very useful for huge Internet service providers as there are too many messages to process, which makes the overview very difficult to keep. But they proposed the small change that the messages of lesser important events should only be visible if an administrator is searching for a fault.

On the other side the possible amplification vulnerabilities turned out as very useful. For visualization, there are in general too many events, but it turned out that especially for connections with very strict service level agreements it is useful to know if there is a potential denial of service threat by a vulnerable device inside the local domain, so the routing can be set to minimize the risks.

The function to analyze the log messages of a router or switch with regard to security was determined to be very helpful, as this work is done by security staff, which often has limited knowledge about network-side topology changes, or requires to keep track of those changes manually. Both alternatives are based on the fact that the root causes of faults are hard to determine.

Furthermore, the project’s focus on SDN enables a lot of new possibilities in monitoring and management. As the SDN controller is designed to get extended with applications, it is possible to connect the monitoring functions directly to the controller. This leads to more efficient algorithm implementations because the export and transformation of the information is not needed anymore. There are also no additional delays between the monitoring and the analyzing. As the network management system is directly connected to the SDN controller, it is also possible to get the information directly from the controller.

IV. CONCLUSION AND OUTLOOK

As the integration of security management functionality, including the visualization of security events, as an extension to the WebCNM research prototype has shown, there is an in-

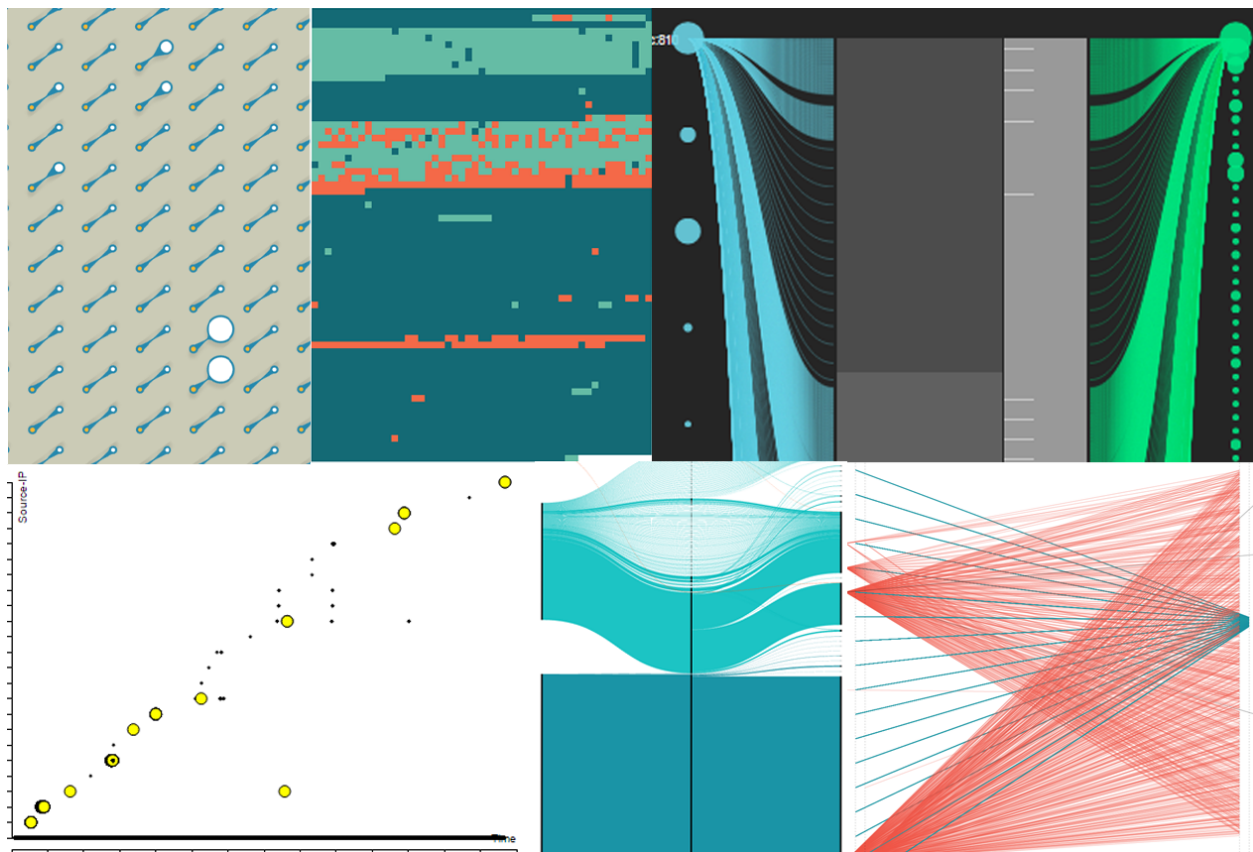


Figure 2. Different visualizations of security events in WebCNM.

interesting potential for converged network and security management tools. Although decisions about which events to include and how to visualize them still require more research, it is obvious that future management software suites can cover both disciplines in integrated manner. While presently only security events from data sources within an organization have been processed, our future work will investigate the inclusion of events and configuration items created by SDN controllers and so-called SDN applications with the goal of enabling WebCNM-based network and security management across organizational borders. The primary security management challenges to this extent are restrictive information sharing policies and technical heterogeneity in the real world, similar to the previous network management challenges that have successfully been overcome. In the long term, self-adapting monitoring data sources that automatically adjust, e.g., their threshold parameters before raising alarms, will also be integrated into WebCNM.

ACKNOWLEDGEMENTS

Parts of this work has been funded by the German Ministry of Education and Research (FKZ: 16BP12309). The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful comments on previous versions of this paper. The MNM-Team, directed by Prof. Dr. Dieter Kranzlmüller and Prof. Dr. Heinz-Gerd Hegering, is a group of researchers at Ludwig Maximilian University of Munich, Technische Universität München, the University of the Federal Armed Forces, and the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities.

REFERENCES

- [1] "The SASER-SIEGFRIED Project Website," <http://www.celtic-initiative.org/Projects/Celtic-Plus-Projects/2011/SASER/SASER-b-Siegfried/saser-b-default.asp>, 2013, [retrieved: 2015-04-10].
- [2] H.-G. Hegering, S. Abeck, and B. Neumair, *Integrated Management of Networked Systems — Concepts, Architectures and their Operational Application*. ISBN 1-55860-571-1, Morgan Kaufmann Pub., 1999.
- [3] ISO/IEC 27001:2005, *Information technology – Security techniques – Information security management systems – Requirements*. International Organization for Standardization and International Electrotechnical Commission, 2005.
- [4] ISO/IEC 7498-4:1989, *Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 4: Management framework*. International Organization for Standardization and International Electrotechnical Commission, 1989.
- [5] P. Marcu, D. Schmitz, W. Fritz, M. Yampolskiy, and W. Hommel, "Integrated monitoring of multi-domain backbone connections," *International journal of Computer Networks & Communications (IJCNC)*, vol. 3, no. 1, 2011, pp. 82–99.
- [6] P. C. Hyland and R. Sandhu, "Management of Network Security Applications," in *Proceedings of the 21st NIST-NCSC National Information Systems Security Conference*. Arlington, Virginia, USA: National Institute of Standards and Technology, Oct. 1998, pp. D-86–D-97.
- [7] J. Dawkins, K. Clark, G. Manes, and M. Papa, "A Framework for Unified Network Security Management: Identifying and Tackling Security Threats on Converged Networks," *Journal of Network and Systems Management*, vol. 13, no. 3, 2005, pp. 253–267.
- [8] S. Kuklinski and P. Chemouil, "Network Management Challenges in Software-Defined Networks," *IEICE Transactions on Communications*, vol. E97-B, no. 1, 2014, pp. 2–9.
- [9] W. Zhu, L. Vizenor, and A. Srinivasan, "Towards a Reference Architecture for Service-Oriented Cross Domain Security Infrastructures,"

- in Internet and Distributed Computing Systems, ser. Lecture Notes in Computer Science, G. Fortino, G. Di Fatta, W. Li, S. Ochoa, A. Cuzzocrea, and M. Pathan, Eds. Springer International Publishing, 2014, vol. 8729, pp. 275–284.
- [10] W. Han and C. Lei, “A survey on policy languages in network and security management,” *Computer Networks*, vol. 56, no. 1, 2012, pp. 477–489.
- [11] W. fei Zhao and Q. Wang, “Study on Network Management and Security Access Control,” *Communications Technology*, vol. 3, 2011, pp. 93–95.
- [12] H. Debar, D. Curry, and B. Feinstein, “The Intrusion Detection Message Exchange Format (IDMEF),” The Internet Engineering Task Force (IETF), Request for Comments 4765, Mar. 2007, [retrieved: 2015-04-10]. [Online]. Available: <http://tools.ietf.org/html/rfc4765>

Design and Implementation of a Scalable SDN-OF Controller Cluster

Min Luo^{*}, Quancai Li⁺, Man Bo⁺, Ke Lin⁺, Xiaorong Wu⁺, Chenji Li⁺, Sheng Lu⁺, Wu Chou^{*}

^{*}Shannon Lab, Huawei Technologies, Inc., Santa Clara, USA

⁺Shannon Lab, Huawei Technologies, Inc., Hangzhou, China

Corresponding Author Email: min.ch.luo@huawei.com

Abstract – Software-defined networking (SDN) is a new paradigm to increase network resource utilization, minimize management complexity, and reduce operational cost. While the logically centralized control in SDN offers many unique advantages that can enable globally optimized routing and resource utilization, the controller in SDN needs to be carefully designed to avoid being a performance bottleneck or a potential single point of failure of the network. This paper presents a scalable clustering approach for the design and implementation of SDN-OF controller. In the proposed approach, controller instances (threads, and processes on a single or multiple physical servers) will mostly run in equal mode for packet-in processing, and only few such instances will serve as the “master” to process special packets that may result in updating network states. A self-learning adaptive mechanism is incorporated to further optimize the number of the controller instances under the varying operating conditions to adapt to changes in network states, controller workloads, and controller data traffic flows dynamically. As a consequence, it automatically enables load balancing and fail-over protection. The proposed approach has been successfully implemented and tested under some stringent conditions, and the performance advantages are observed.

Keywords – Software defined networking; Openflow; Centralized Controller Cluster; Message Dispatching Cluster; Message Processing Cluster; Load Balancing; Fail-over.

I. INTRODUCTION

Recent advances in computing technologies have led to new business opportunities with improved business agility, high productivity, and reduced operational cost. However, while devices and applications on the network are advancing rapidly, the underlying networking infrastructure (architecture and the protocol stack specifically) has been evolving at a much slower pace. This disparity causes difficulties and ineffectiveness to support applications emerging from new business opportunities, because the traditional network control and management environment is not well-suited for new innovation, especially the complex legacy network infrastructure which cannot meet the growing demands of the application needs. There is a growing effort, spearheaded by the Open Networking Foundation (ONF) to create an open and programmable networking environment through software-defined networking (SDN) [1]. OpenFlow [2] is such an industry effort of ONF, in which the network control plane is logically centralized and decoupled from the data forwarding plane. In SDN, OpenFlow (OF) is used as the protocol between the SDN controllers and the data switches

to manage and dictate the networking behavior of OF compliant switches.

In this paradigm, application developers, enterprises, and carriers can gain unprecedented programmability to control how data flows in the underlying data networks to best support the applications and readily adapt to their changing business needs. The openness and the flexibility in SDN provide new capabilities for offering better Quality of Service, and reducing the need to purchase specialized and expensive networking equipments. However, in SDN, network controller forms a logically centralized network control apparatus. If the processing power of the controller remains constant, flow setup time delay can grow in proportion to the packet-in traffics to the controller. This situation needs to be carefully addressed in SDN, especially when it scales up with an increasing number of OF switches in the network [2]. As such, there is a critical need towards a scalable centralized controller design and implementation, such that it can scale well with the size of the network and the growth of the traffics. To address this critical technical issue in SDN, new techniques and designs are needed to overcome the limitations in existing solutions. In addition, more challenging stress tests and studies on the controller scalability are needed because of the central role of the controller(s) in SDN.

SOX [9] is designed and implemented based on model driven technologies for extensibility and consistency. This allows SOX to take advantage of the recent advances in distributed computing, to enhance the controller generality, scalability, reliability, and interoperability. In addition to be a generalized SDN controller which can support multiple evolving OF standards, e.g. OF1.0, OF1.2, OF1.3, etc., it also provides some key transformation capabilities to interworking with existing networks with BGP/PCE and existing transporting protocols such as MPLS.

In this paper, we present an extension to SOX, and describe a centralized SDN-OF controller cluster architecture to scale up the network management power with the size of the data network (e.g. department, campus, branch data center or even a WAN, etc.). The remaining part of this paper is organized as follows. In Section II, we briefly summarize some design and architectural principles that is applied in the proposed approach. In Section III, we describe the design and architecture of the proposed scalable controller cluster and its main functional components. In Section IV, we present scalability studies and performance results. Findings of this paper and some future R&D directions are summarized in Section V.

II. RELATED DESIGN AND ARCHITECTURE

HyperFlow [3] is a distributed event-based control plane for OpenFlow, in which each controller directly manages the switches connected to it and indirectly programs or queries the other controller through proprietary communications. However, its scalability, performance, and robustness needs to be further studied and evaluated on large test-beds with more stressful conditions. ONIX [4] provides a platform in which a network control plane can be implemented as a distributed system, and it provides a general API for interaction among data, control, and management planes. It recommends that, as a good “best practice”, some consistency among distributed controllers could be sacrificed for high performance requirements. However, it does not address how to leverage the advanced high performance computing techniques to improve the performance of the centralized controller in order to minimize such inconsistency. *ElastiCon* [5] focuses on dynamically shifting workload to allow the controllers operating at a pre-specified load window, while *Kandoo* [6] and *DIFANE* [7] try to limit the overhead of processing frequent events on the control plane. Moreover, *DIFANE* would keep all traffics in the data plane under control by selectively directing packets through intermediate switches with specific rules. Finally, *OpenDayLight* [8] currently allows a switch to be connected to one or more controllers, but only a single master controller can process new flow requests which is not for large networks in the current release. In addition, its management applications only operate on the controller(s) through fixed IP addresses that are assigned randomly without explicitly taking the load balancing needs into consideration.

Following ONIX [4], some important principles for the design and implementation of SDN controller cluster are described as follows:

A. Performance First

As the size of the network with SDN-Openflow enabled devices increases, the amount of new flow requests between the switches and the controllers grow accordingly. A scalable controller cluster is needed to ensure that the short new flow setup latencies are maintained for all switches in the network as the packet-in traffics from the switches to the centralized controller scale up. As such, the performance is always the most important consideration in our controller design.

B. Weak Consistency if Necessary

As the network scaling up, keeping strong consistency of network-wide states and related information will consume lots of computing and storage resources for the centralized multi-controllers in SDN. Similar to [4], in our approach, the controller cluster would trade some consistency for high performance, since most of the instantaneous inconsistencies would have negligible effects to the controller.

C. Light Weight Framework

The framework of the controller (cluster) should be as light weight as possible for large-scale centrally controlled networks, so that it can ensure quick response to all control traffics destined towards the controller cluster and achieve high performance at relatively low cost.

D. High Scalability:

In order to meet the growing control traffic demands for large-scale networks, the controller cluster should be able to scale seamlessly and autonomously, as long as the aggregated I/O capacities are not fully consumed.

E. Robust

A controller cluster should provide highest possible reliability, availability, and resilience to all kinds of potential network errors, including failures of devices, software, and more importantly malicious attacks to the centralized controller and its communication channel(s). Only with this stringent robustness, can it protect the network against significant loss to the business opportunities and avoid catastrophic consequences due to the possible single point of failure of the controller cluster. In our approach, the controller cluster with multiple servers runs in the equal mode, which effectively provides automatic fast failover and load-balancing.

III. OVERVIEW OF THE PROPOSED CENTRALIZED CONTROLLER CLUSTER APPROACH

As discussed in the previous section, the new centralized SDN/Openflow control paradigm calls for robust and scalable controllers, especially for large networks with ever increasing and varying traffic volume. In this paper we propose an innovative multi-controller cluster and management mechanism to address the key issues in building a powerful and reliable centralized SDN-OF controller system that offers:

- Automatic load balancing and failover.
- Reduced communication overhead (synchronization and coordination between controllers).
- Scalability and extensibility.

The proposed cluster architecture is depicted in Figure 1, while [9] provides a detailed description on other core components such as the Network Information Base (NIB), Topology Management, Routing, Host Management, etc. The controller cluster consists of two major components: the message dispatcher (MD) or a MD Cluster (MDC), and the message processing cluster (MPC).

The main function of MD is to establish TCP channel connection from switches to a MD server node in the cluster in order to receive various packet-in(s) (PI) into some policy-based and prioritized input queues depending on the type of packets (for example, normal PI messages, status updates, or control messages, etc.) and their processing priority.

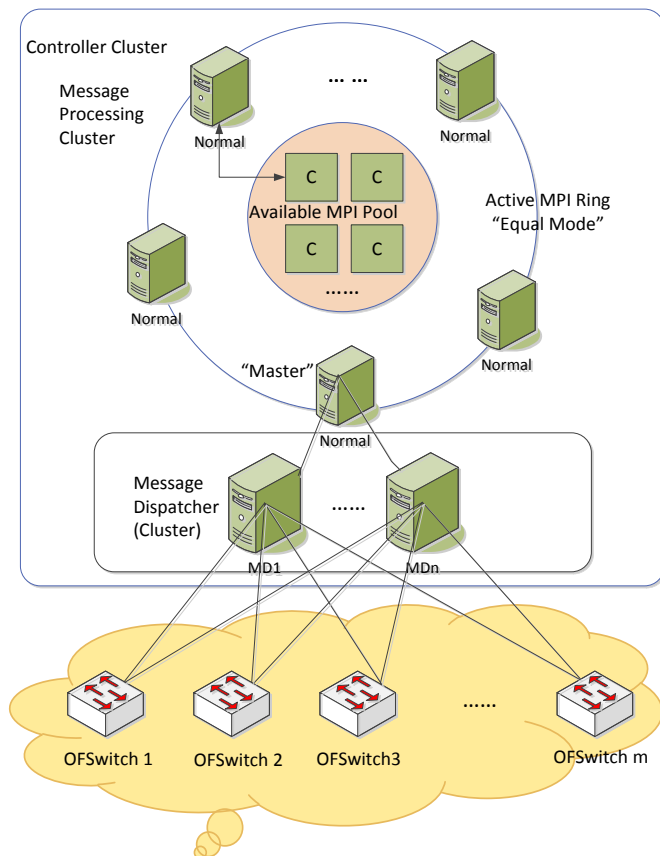


Figure 1. Diagram of controller cluster architecture

The MPC is responsible for processing the received packets, finding the proper path by invoking routing and resource allocation optimization algorithms, transforming the optimized routing decisions into flow table entries, and sending them back to the MD while the MD will eventually transport these flow entries back to all the switches en-route.

The following two subsections will discuss MD/MDC, and the MPC respectively.

A. Message Distribution Cluster (MDC)

A single message dispatcher server could control all switches and distributes all the messages to the MPC for a reasonably sized network. However its performance, especially its scalability could be limited by the total throughput of a single server, and also it lacks the required high availability and reliability for centrally managing a network of significant size, e.g. those networks with 1000 or more switching or routing nodes.

In our proposed controller cluster, we extend the single message dispatcher to a Message Distribution Cluster (MDC) with multiple dispatcher server nodes as illustrated in Figure 1. An OF1.3 enabled switch can be controlled by multiple controllers and support such features with a Message Distribution Cluster.

When the controller cluster is initiated, each switch is configured to establish a connection with all MDC servers. Then the MDC servers distribute the handshake messages received from the switches to the Message Processing

Cluster MPC, while the “instantaneous MP master” (to be descried later) in MPC can record and maintain all the connected switches information after processing the handshake messages. In general, the MPC coordinates the management of the switches via consistent hashing to designate a MDC server for each switch. Policies can be configured such that each server within the MDC works in equal mode, while each server control (responsible for receiving, processing, and replying asynchronous messages) a number of designated switches. Currently, some rudimentary techniques are used for partitioning the network into “domains”, mostly utilizing the network topology information, such as “neighboring” and hop-counts to measure how close a switch is to all other switches, and then grouping the “neighboring” nodes into the same domain. If the MP master detects some MD servers leave the cluster for any reason, it would reassign those switches originally controlled by those left MDC without overloading any remaining active servers in the cluster. As illustrated, the inherent ring structure of the active Message Dispatcher Instances (MDIs) and those available and ready to serve MDIs in the MDI pool effectively enable the use of the enhanced consistent hashing in order to make sure all the servers in MDC are utilized in a balanced fashion. Therefore the MDC could achieve the required high performance, adaptability, and scalability even when servers are added or removed from the cluster for any reason. The detail of the process is depicted in Figure2. For simplification, the internal active MDI ring and the available pool are not depicted herein.




Figure 2. Managing Switches in Dispatcher Server(s)

The operation workflow has the following main steps:

- Step1: Switch(s) initially connects with all the dispatcher servers in the MDC.
- Step2: Dispatcher servers send the connected switch messages to the MPC, where the instantaneous “master” MPI could process them.

- Step3: The “master” MPI will decide which dispatcher server is assigned to control this switch.
- Step4: Dispatchers only reply to switches that they are responsible, in either master/slave or the equal mode.

B. Message Processing cluster (MPC)

The main function of the message processing cluster (MPC) is to receive and process switch messages from the MDC, send proper response messages (such as flow entries for routing the new packet-in) back to the MDC, while the MDC will dispatch those replies back to the relevant switch(s).

MPC can be a stand-alone server, or a cluster of physical servers on its own, each such server can provide multiple message processing instances (MPIs) at either the thread or process level [9]. Again for the sake of high performance and reliability, a specialized cluster architecture is used in our approach.

As depicted in Figure 1, the active MPIs within the MPC are organized into a ring structure, where each MPI would take the “master” role in turn for a short amount of time, while it will process the control and other status related messages within a certain specified time slot or until there are no more such messages to process, then it will proceed to become a normal MPI. For example, the “instantaneous master” is responsible for handling OF switching events, such as session connection, topology discovery, or status updates. Then it will quickly pull some packet-ins from the prioritized message queue and immediately pass the “master” role to the next MPI on the ring. After the previous “master” transitions itself into a regular “equal” mode MPI, it will start to process the packet-in message, and generate flow entries based on the routing and forwarding strategies and policies using the key packet-in attributes.

As network load changes, more or less MPIs may be needed, and if the traffic suddenly surges, the proposed approach should be able to respond instantly to add more MPIs. Therefore, a separate pool of MPIs are initialized, and updated with the latest information from the shared NIB, and any MPIs in the pool are ready to be pulled into the active ring to process packet-ins or other control messages. In [9, 10], more details are presented on how the NIB, and the multiple priority queues are architected in the controller cluster. In this way, MPC can achieve high availability and automated load balancing with the desired high performance. In addition, it can significantly reduce the management work load, because all MPIs are identical in structure and can be cloned from a uniform infrastructure.

C. Packet-In Message Processing

We categorize messages handled by the MPC into two types: The control messages that cover all Openflow protocol messages, along with various events originated from network status or policy changes, and the normal packet-in (PI) messages that include forwarding-request packet (data flow, host-to-host request data flow etc.).

1) Control Messages

The control message is mainly used in the control layer, and they also provide critical information needed for the data forwarding layer (normal PI messages). As depicted in Figure 4, only the “master” controller will process those control messages, therefore the cost to achieve much desired consistency across all the MPIs in the cluster is minimized.

Typical control messages include:

- Protocol interactive messages (handshake messages, port status, etc.) between controller and switches.
- Some special control messages wrapped in PI Messages, such as LLDP link discovery messages.
- Static messages (basic configuration, policy configuration, etc.).
- Events due to network status changes.
- Messages or other data generated dynamically through the northbound APIs.

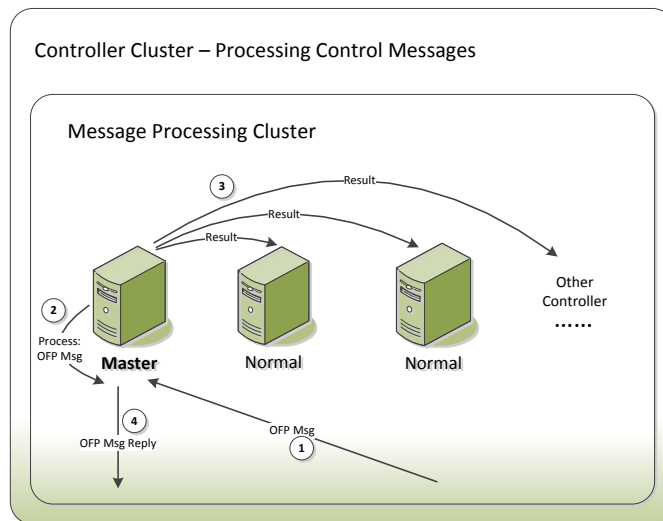


Figure 3. Processing Control Messages

Some of the control messages have strong dependencies with each other, and the order of processing them become critical to maintain the proper consistency. This is one of the main reasons why the current instantaneous “master” approach is taken, in order to process all control messages with one MPI at any given time. Therefore, not only it can enforce the proper processing order, but it also makes the updates to the other MPIs and the shared central NIB consistent at minimum costs. In case that the “master” MPI becomes dysfunctional, the next MPI on the ring would take over the “master” role. When this occurs, some control messages or their processing results may get lost, that could result in some temporary inconsistency in the network. However, the enforced short maximum time interval to transfer the “master” role to the next available MPI would limit the consequence of the inconsistency to an acceptable level. Furthermore, after the controller completes topology discovery and feature learning, there will be much more reading operations than writing in a SDN network, while only the failed writing operation could lead to such

inconsistency. The proposed approach properly follows what is recommended in [4], and innovatively provided a solution that greatly improves the processing performance with affordable costs, while allowing certain acceptable weak consistencies at certain duration-controlled time intervals.

2) Equal Mode for Normal Packet-In Messages

For normal PI messages, the equal mode is adopted, in which all servers (threads or processes) share the processing loads while providing backups to each other. As discussed earlier, the “master” MPI will pull a packet-in message from the normal PI message queues, transfer the control to the next MPI on the ring, and then start to process the new packet-in, as all other active MPIs will do.

Controllers in SDN are to process new flow routing requests by collaborating with upper layer management applications, such as routing and resource allocation, transforming the resulting decisions into flow entries, and then distributing them to all OF-switches along the designated path(s).

In the proposed approach, each MPI uses the local cached data, fetched or synchronized from the central NIB to perform the above tasks.

IV. ADAPTIVE RESOURCE CONTROL MODEL IN MPC

In order to maximize the performance of the contemporary powerful multi-core CPUs and available I/O throughput, an adaptive resource control model is established. Today’s abundant mobile internet traffics are dynamic and bursty, complicated with different QoS requirements. They necessitate finer-grained control with enough pre-provisioned resources to make sure that the perceived performance is acceptable even in worst situations. The providers, however, need to properly control the use of those resources for their increased utilization, and at the same time, it is aimed to reduce the network management complexity and operating cost. For example, if it can effectively apply SDN based centralized control with adaptive and optimized routing and resource allocation [12], they can offer to dynamically reroute the traffics, enable or disable the use of certain routes, and increase or decrease capacities on some routes. As traffic patterns in general would be quite different between working hours and off hours, such dynamic routing decisions would not only help achieve their business objectives, but also enable a greener computing and data networking paradigm.

With the proposed approach, as the new packet-in requests varies, the number of concurrent threads in each MPI, and also the number of servers allocated to the MPC can be increased or decreased dynamically. However, after numerous experiments and comparison studies, it was recognized that increasing the number of threads or servers may not necessarily improve the actual single MPI or the entire MPC processing capability, especially when we have to deal with a large number of PIs from a large scale network. When the number of concurrent threads in a MPI or servers in MPC increases beyond a certain limit, the newly added thread or server can consume additional precious data bandwidth, CPU processing power, and

memory resources, leading to the increased risks of data locking, exceeding the limit of the total available I/O capacity, and eventually, the overall performance degradation.

The use of the Multi-dimensional DHT (MDDHT) [11], along with the MDC and MPC in the proposed approach, can learn different traffic patterns and the related parameters to discover the hidden relationships between the traffic loads and the proper number of MPIs/servers needed for efficient network traffic control under various operating conditions. Such parameters include the number of packet-in messages that a thread or a server can process in a second, given its CPU/memory processing capabilities, the I/O capacity, and the bandwidth at the port and the aggregated switch level, etc.

V. PERFORMANCE AND SCALABILITY CONSIDERATION

Our main objective is to design and implement a powerful controller cluster to address the critical performance issue for the centralized controllers in a large SDN-OF network, while providing the desired fault tolerance and load balancing capabilities. As discussed before, slightly sacrifice made in the consistency with some acceptable packet losses to exchange for significant gains in performance [4], should be an excellent overall trade-off. In addition, such a compromise would render the proposed approach much more cost-effective.

In an earlier design of the proposed approach, direct connections, through some centralized and prioritized packet queues, were established between the switch and the MPC, without a dedicated MD or MDC. But soon we found that such a “centralized” and “single I/O” mechanism would limit the throughput of the controller/MP cluster and become a bottleneck, causing the total throughput of the cluster to decrease and the latency of processing each individual request to increase, no matter how powerful the MPC is.

Later MDC was introduced between the controllers/MPC and switches. Even though the direct connection model could provide faster response time for new flow requests if the overall I/O limit is not reached, but it could be degraded at an explosive rate soon after the sustained I/O is approaching the system upper limit. With multiple servers in the MDC, where each server brings in its own portion of I/O capability, the overall throughput of the proposed approach can be improved almost linearly until reaching the limit imposed by the capability of the MPC. On the other hand, it is found that accelerating TCP/IP communication performance by DPDK technology [13] and high-performance TCP/IP protocol stack can significantly improve the performance of both the MDC and the MPC.

VI. SCALABILITY AND PERFORMANCE ANALYSIS

In this section, we review how we conducted performance testing experiments and present related comparison results. We focus on how the proposed approach could improve the performance, scalability, and reliability. In particular, we study how fast such a centralized controller cluster could process new PI messages

in a large-scale communication network, and investigate on how certain known issues could be addressed to avoid potential performance bottleneck.

A. Experimental Methodology and System Setup

We designed and implemented the proposed scalable centralized controller cluster prototype with one Message Dispatcher Instance with one MD server and several Message Processing Instances (MPIs) running on two to four servers, and we compared the throughputs achieved by such a cluster with different number of MP servers. For a fair comparison, we first evaluated our centralized controller cluster with only one MPI in the MPC, and after that, we moved on to evaluate the throughput of the cluster with multiple MPIs in MPC.

As shown in Figure 6, we constructed a network topology with two OpenFlow switches, one Message Dispatcher Instance (MDI), and several Message Processing Instances (MPI). They were run on separate 2.4GHz 64bit Intel Xeon E5620 server machines with 8GB RAM to avoid interferences, where all data links were 10Gbps. It is also important to note that for each processed packet-in (PI) message, the centralized controller cluster would reply with one packet-out (PO) message and two flow-mod (FM) messages. In general, N separate FM messages would need to be created and distributed if there are N nodes en-route for a PI. Therefore, in our experimental configuration, there were three reply messages for each PI with a 1:3 (or N+1 in general) ratio between the receiving and replying endpoints in the centralized controller cluster.




Figure 6. Experiment setup

To our knowledge, currently published controller performance results are based mainly on using only ONE switch. Consequently, the ratio of the number of PI messages and the number of PO plus FM messages are only 1/2. In addition, the controllers and switches are mostly only OF 1.0 compliant, supporting only one flow table. In our experiments, the OF switches were all compliant to the newer and more sophisticated OF 1.3 standard, and the

controller and the switches actually supported a flow pipeline with nine flow tables optimized for typical network applications, such as firewall, ACL, layer 2 or layer 3 transporting, VLAN, etc. Therefore, what being tested and presented here are more practical as the older version of OF 1.0 cannot support many critical network functions for existing data services.

We injected traffic flows from one host to another based on a fixed but configurable rate. Each flow was generated as a single 64 Byte UDP packet. Consequently, in our experiments, each PI message was effectively 84 Bytes, while PO messages were 80 Bytes each, and FM messages were 136 Bytes each respectively.

B. Experiments with one MD and 1-5 MPIs in MPC

We computed the rate of packet-out (PO) and flow-mod (FM) messages that would be created and replied from the controller cluster. As depicted in Figure 7, we initially studied and compared the centralized controller cluster’s scalability and performance with varying number (1 to 5) of MPIs.

Message Processing Instances	Packet-in		Packet-out		Flow-mod		Packet-out + Flow-mod	
	Packet Rate (kpps)	Packet Bandwidth (Mbps)	Packet Rate (kpps)	Packet Bandwidth (Mbps)	Packet Rate (kpps)	Packet Bandwidth (Mbps)	Packet Rate (kpps)	Packet Bandwidth (Mbps)
	1 Instance	560	376	590	358	1120	1200	1680
2 Instances	1120	752	1150	716	2240	2400	3360	3116
3 Instances	1660	1115	1690	1062	3320	3600	4980	4662
4 Instances	2240	1182	1760	1126	3520	3800	5280	4926
5 Instances	2800	1182	1760	1126	3520	3800	5280	4926

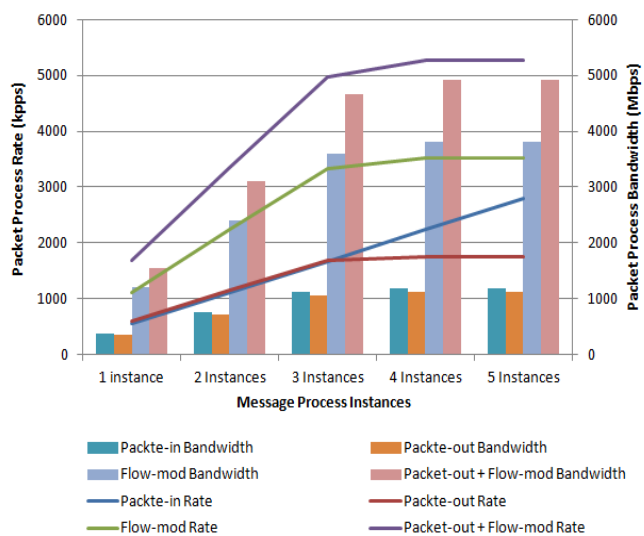


Figure 7. Centralized Controller Cluster throughput Vs. the number of Message Processing Instances: 1 MD Server

As can be seen from Figure 7, when the number of MPIs was three or less, the throughput of the centralized controller cluster could scale linearly with the number of MPIs, where each MPI could process 560K PI messages per second. But as we continued to increase the number of MPIs, the throughput of the centralized controller cluster stopped increasing linearly and eventually stagnated.

Further drill-down analysis revealed that with four MPIs, the centralized controller cluster could process 1.7 million PI messages per second; as a result, 1.7 million PO messages and 3.4 million FM messages would need to be generated and replied to the 2 switches in the experimental network per second. As such, the MDI would consume about (1.7 million POs * 80 Bytes/PO * 8 Bits/Byte + 3.4 million FMs * 136 Bytes/FM * 8 Bits/Byte), or requiring about a total of 4.8 Gbps bandwidth in one direction. That has almost exceeded the maximum TCP replying capability using the standard Linux socket API.

As the number of Message Processing Instances continues to increase, I/O performance eventually become the main bottleneck that impacts the scalability and the throughput of the centralized controller cluster.

C. Effects of Increasing the number of MDs in MDC

In order to overcome the I/O limit imposed by one MD server in the above experiment and fully utilize the processing capabilities of the MPC, two MD servers were used in a new experiment. Each MD server was assigned to be responsible for one switch just to simplify the matter. In practice, such as in the distributed SOX [10], some mechanism is needed to partition a large network so that each MD server could be assigned to be responsible for a subnet.

Figure 8 presented the result. As it can be seen, the total throughput now scales linearly as the number of MPIs increases. With the proposed approach, the MPC can be scaled easily to increase its processing capacity with additional MPIs in MPC, and now with the additional MD servers in the MDC to overcome the I/O limit, the centralized controller cluster can manage much larger network with adequate I/O bandwidth such that the overall system can scale up linearly.

Message	Packet-in ^o		Packet-out ^o		Flow-mod ^o		Packet-out+Flow-mod ^o	
	Packet ^o	Packet ^o	Packet ^o	Packet ^o	Packet ^o	Packet ^o	Packet ^o	Packet ^o
Instances ^o	Rate ^o	Bandwidth ^o	Rate ^o	Bandwidth ^o	Rate ^o	Bandwidth ^o	Rate ^o	Bandwidth ^o
	(kbps) ^o	(Mbps) ^o	(kbps) ^o	(Mbps) ^o	(kbps) ^o	(Mbps) ^o	(kbps) ^o	(Mbps) ^o
1 instance ^o	560 ^o	376 ^o	590 ^o	358 ^o	1120 ^o	1200 ^o	1680 ^o	1558 ^o
2 instance ^o	1120 ^o	752 ^o	1150 ^o	716 ^o	2240 ^o	2400 ^o	3360 ^o	3116 ^o
3 instance ^o	1680 ^o	1128 ^o	1740 ^o	1074 ^o	3360 ^o	3600 ^o	5040 ^o	4674 ^o
4 instance ^o	2240 ^o	1504 ^o	2330 ^o	1432 ^o	4480 ^o	4800 ^o	6720 ^o	6232 ^o
5 instance ^o	2800 ^o	1880 ^o	2920 ^o	1790 ^o	5600 ^o	6000 ^o	8400 ^o	7790 ^o

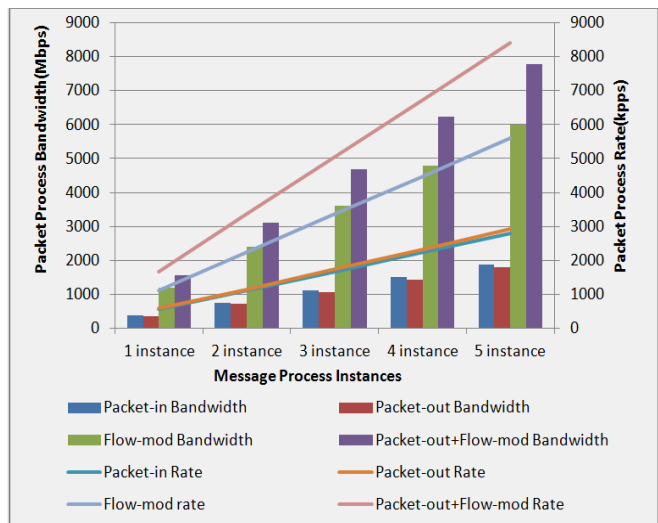


Figure 8. Centralized Controller Cluster throughput Vs. the number of Message Processing Instances (2 MD Servers)

D. Other Considerations

Furthermore, with this centralized controller cluster, one MD instance/server can be enough to match the PI message processing capacity of the deployed MPC. If more PI messages need to be processed, distributed controller clusters should be used, while each cluster manages a smaller domain or subnet of the large scale network with enough I/O bandwidth between the controller cluster and all switches it controls as in DSOX [10].

As higher speed (such as 40 Gbps or even 100 Gbps) NIC cards become available, obviously both processing power of MDC and MPC can be significantly increased. As it is relatively easy to increase the number of servers in the MPC to achieve much higher PI message processing rate, more MDIs (or servers) in the MDC should be added to further balance the load of the much increased traffics between the centralized controller cluster and the switches.

VII. CONCLUSION AND FUTURE DIRECTIONS

This paper presented a scalable centralized controller clustering approach for SDN with Openflow. In our approach, the controller instances (threads, and processes on a single or multiple physical servers) would handle the control and normal data packets differently, in which the packet-in message processing was carried out by different controller instances in the cluster running in equal mode. It is aimed to achieve maximum performance with slightly sacrificed but acceptable and temporary inconsistency. The controller cluster in our approach demonstrated excellent scalability until the system I/O limit was reached. A self-learning and adapting mechanism was also incorporated in our approach to further optimize the number of controller instances under varying operating conditions to adequately handle frequent changes in network states and traffic flows. In addition, the described approach provides automatic load balancing and fast fail-over at almost at no extra cost.

Experimental studies were conducted, and the test results with varying number of physical servers (or processing instances) were presented and studied. It clearly demonstrated the feasibility and advantages of the proposed approach.

Efficient control and management of large scale networks based on SDN is an active research area. Many issues remain to be investigated, such as the balance between the proper level of consistency and the high performance in a distributed cluster environment, the interoperability of SDN with existing IP networks, dynamic network resource utilization, etc.

REFERENCES

- [1] Open Networking Foundation. Software-Defined networking: The new norm for networks. ONF White Paper, 2012.
- [2] "OpenFlow Switch Specification 1.3.0", Open Networking foundation.
- [3] A Tootoonchian, Y Ganjali, "HyperFlow -- A Distributed Control Plane for OpenFlow", Proceeding of the 2010 INM conference, https://www.usenix.org/legacy/event/inmwren10/tech/full_papers/Tootoonchian.pdf
- [4] T Koponen, M Casado, N Gude, J Stribling, L Poutievski, M Zhu, R Ramanathan, Y Iwata, H Inoue, T Hama and S Shenker. "Onix: a distributed control platform for large-scale production networks", Proceedings of the 9th USENIX OSDI conference, Vancouver, 2010. http://static.usenix.org/events/osdi10/tech/full_papers/Koponen.pdf
- [5] A Dixit, F Hao, S Mukherjee, T Lakshman, R Kompella, "Towards an Elastic Distributed SDN Controller". Proc. of the ACM SIGCOMM Workshop on HotSDN. pp. 7-12, Hong Kong, August 2013.
- [6] S Yeganeh, Y Ganjali. "Kandoo: a framework for efficient and scalable offloading of control applications", Proc. of the ACM SIGCOMM Workshop on HotSDN, pp. 19-24, Helsinki, Finland, August 2012
- [7] M Yu, J Rexford, M Freedman, J Wang, "Scalable flow-based networking with DIFANE", Proc. of the ACM SIGCOMM. pp. 351-362, 2010
- [8] OpenDaylight, <http://www.opendaylight.org/>, [accessed: 2014-12-18]
- [9] M. Luo, Y Tian, Q Li, J Wang, W Chou. "SOX –A Generalized and Extensible Smart Network Openflow Controller(X)", The First SDN World Summit, Germany, October 2012. http://www.layer123.com/download&doc=Huawei-SOX_WP_V1.0
- [10] Q.Li, K. Lin, M. Luo, etc., "DSOX: Architecture and Design," Technical Report, Huawei Shannon Lab, May 2014.
- [11] M. Luo, X. Wu, Y. Zeng, J. Li, K. Lin, B. Man, and W. Chou, "Multi-dimensional In-Memory Distributed Hashing Mechanism for Fast Network Information Processing in SDN", accepted by the 9th International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS-2015), Blumenau, Brazil, July 2015
- [12] M. Luo, Y. Zeng, J. Li, W. Chou, "An Adaptive Multi-path Computation Framework for Centrally Controlled Networks", accepted by Journal of Computer Networks, Elsevier, February 2015
- [13] Intel® DPDK: Data Plane Development Kit. <http://dpdk.org/>, [accessed: 2014-12-18]

Management Information Systems in Health Sector: Evidence of Mandatory Use

Ioannis Mitropoulos

Dept. of Business Administration,
Technological Education institute of Western Greece
Patras, Greece
e-mail: mitro@teiwest.gr

Abstract — The purpose of this study is to investigate the factors that negatively impact the development of health information systems and the documentation of the theoretical model for the Mandatory Use of Software Engineering. The factors were investigated through a structured questionnaire on 120 users of the system of the General Hospital of Lamia in Greece, who had recent experience in the introduction of an integrated information technology system. The results of the analysis shows that targeted interventions of the administration of the Hospitals to inform, about expected net benefit of users, contribute to the acceptance and therefore to the functional and comprehensive introduction of information systems in health, highlighting the actual substance as tools in modern decision making.

Keywords-management information systems; health care; hospital; Greece.

I. INTRODUCTION

The rapid increase in health expenditure creates problems of sustainability in National Health Systems worldwide. Decision making in health aims to search and select the best solutions that define the best use of limited resources, the most efficient and qualitative confrontation of the health needs of the population and the maximization of the social benefit. Management Information Systems (MIS) in health sector is the critical tool for rational decision making and integral documentation in any effort of planning / programming implementation. The success of the implementation of management information systems depends on the attitude of users towards them especially under an environment that its used is mandatory, such as in hospital services that its use is imposed. Understanding the reasons why people accept or reject the information and communication technologies has proven decisive factor for their survival.

The best known bibliographic models-theories that have been developed to better understand the factors that contribute most to the user acceptance of the new technology are: the theory of reasoned action proposed by Ajzen and Fishbein [1], the model of acceptance of technology [2-3], the unified theory of acceptance and use of technology [4] and the integrated research model of Wixom and Todd [5].

We apply the above theories in hospital sector to show that the lack of user acceptance is the key barrier to the success of new information systems. We address the factors that explain why users accept or reject information systems and how their acceptance is affected by the mandatory use.

The rest of this paper is organized as follows. Section II describes the theoretical model and the research questions. Section III describes the statistical inference analysis and presents the results. Section IV addresses the conclusions.

II. PRESENTATION OF THE RESEARCH TOOL

The introduction of the Information System in the General Hospital of Lamia in Greece was implemented in a coordinated way and replaced the majority of previous procedures. It involved employees of all specialties (administrative, nursing, medical and paramedical staff) who are Integrated Information System users.

Therefore, we can consider the environment of the hospital as a mandatory environment and thus useful to investigate it through the model discussed above for mandatory environments (Model for Mandatory Use of Software Technologies - MMUST) of Koh, Prybutok, Ryan and Wu [6].

Based on the above theoretical model, the research hypotheses are as follows:

- H1: Higher quality information is associated with a higher level of satisfaction on the information.
- H2: Higher level of satisfaction on the information is associated with a higher level of expected performance.
- H3: Higher level of expected performance is associated with more positive attitude while using the system.
- H4: The social influence has a positive direct impact on the expected performance.
- H5: The positive attitude is associated with extensive use of the system.
- H6: The attitude of the user of the system is positively correlated with the user satisfaction by the system.
- H7: The use of the system is positively correlated with net benefits.
- H8: The overall satisfaction on the system is positively correlated with the net benefits.

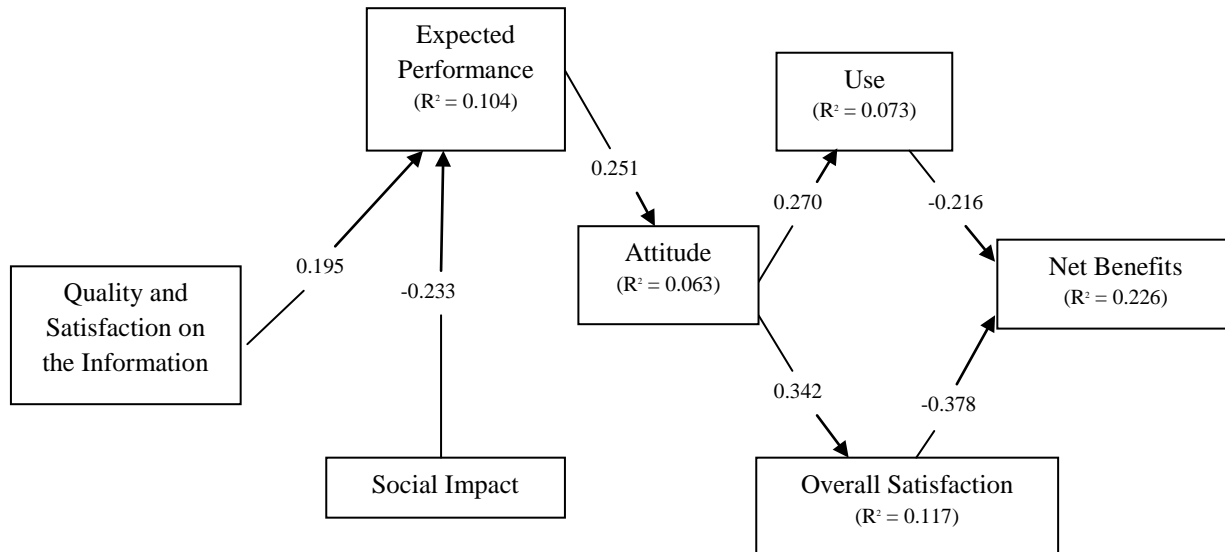


Figure 1. MIS acceptance model correlation and R² results

Based on the MMUST model, we use a structured questionnaire with the following dimensions: Information Quality; Satisfaction on the Information; Attitude; Expected Performance; Social Influence; Use; Overall Satisfaction; and Net Benefits. The responses were measured according to a 5-point Likert scale ranging from 'strongly disagree' to 'strongly agree'. The statistical analysis carried out was a series of controls to determine whether our sample match the proposed model that investigates the acceptance of information systems for mandatory environments.

III. RESULTS

Statistical analyses were conducted using the Statistical Package for the Social Sciences (SPSS 18). Exploratory factor analysis using principal axis factoring with oblique rotation was conducted to explore the factor structure of the MMUST.

Therefore, in terms of the factors-variables created are:

- Factor 1 - Expected Performance (PE).
- Factor 2 - Quality and Satisfaction on the Information (IQIS).
- Factor 3 - Net Benefits (NB).
- Factor 4 - Use (US).
- Factor 5 - Overall Satisfaction (OS).
- Factor 6 - Social Impact (SI).
- Factor 7 - Attitude (AT)

Considering the above, significant correlations exist between factors in the MMUST model. For example factor 1 "Expected Performance (PE)" has positive correlation coefficient (0.195) with Factor 2 "Quality and Satisfaction on the Information (IQIS)". On the other hand factor 1 has a coefficient of determination $R^2=0.104$ (ie, explained variance), when regression with factors 1 and 6 is performed. Figure 1 shows the analytic results.

The results showed that the MMUST model is partially confirmed, for the interpretation of user acceptance. It was found that the level of customization that shows the sample

of user responses is quite low. Regarding the factors "Information Quality" and "Meeting on the Information", the factorial analysis revealed the existence of a single factor, which we called "Quality and Satisfaction on the Information" instead of the two separate factors that appeared in the original model. This indicates a match and absolute relevance of these two concepts in the minds of users.

IV. CONCLUSIONS

Exploring the model proved that the determinants of acceptance is the "attitude" (as the central key variable) and "the view of the Net Benefits" (as the ultimate objective). The strengthening of these points should be therefore the focus and the interest of Administrations. Also, 'the contribution of social influence' is important, which importance increases in mandatory environments and is a factor that can greatly be enhanced by management.

These findings provide a useful tool to extract certain conclusions concerning the operations of the hospital administration to increase acceptance and therefore the use and performance of Information Systems.

ACKNOWLEDGEMENTS

This research has been co-funded by the European Union (European Social Fund) and Greek national resources under the framework of the 'Archimedes III' project of the 'Education and Lifelong Learning' Operational Programme.

REFERENCES

- [1] I. Ajzen, and M. Fishbein, "Understanding attitudes and predicting social behavior," Prentice-Hall, 1980.
- [2] F. Davis, "Perceived usefulness, perceived ease of use and user acceptance of information technology," MIS Quarterly, vol. 13, no. 3, pp. 319-340, 1989.
- [3] S. Y. Yousafzai, G. R. Foxall, and J. G. Pallister, "Technology acceptance: A meta-analysis of the TAM: Part

- 1," *Journal of Modelling in Management*, vol. 2, no. 3, pp. 251-280, 2007.
- [4] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User Acceptance of Information Technology: Toward a Unified View," *MIS Quarterly*, vol. 27, no. 3, pp. 425-478, 2003.
- [5] B. H. Wixom, and P. A. Todd, "A theoretical integration of user satisfaction and technology acceptance," *Information Systems Research*, vol 16, no.1, pp. 85-102, 2005.
- [6] E. C. Koh, R. V. Prybutok, D. R. Ryan, and Y. Wu "A Model for Mandatory Use of Software Technologies: An Integrative Approach by Applying Multiple Levels of Abstraction of Informing Science," *Informing Science: the International Journal of an Emerging Transdiscipline*, vol. 13, pp. 177-203 2010.

Designing Security Policies for Complex SCADA Systems Protection

Djamel Khadraoui and Christophe Feltus

Luxembourg Institute of Science and Technology, 5 Avenue des Hauts-Fourneaux,
L-4362 Esch-sur-Alzette, Luxembourg
email: firstname.name@list.lu

Abstract—The management and protection of these SCADA systems must constantly evolve towards integrated decision making and policy driven by cyber security requirements. The current research stream in this domain aims, accordingly, to foster the smartness of the field equipment which exist through the generic concept of SCADA management and operation. Those components are governed by policies which depend on the components roles, as well as on the evolution of the crisis which also confer to the latter the latitude to react based on their own perception of the crisis evolution. Their latitude is calculated based on the component smartness and is strongly determined by, and depending on, the cyber safety of the component environment. Existing work related to crisis management tends to consider that components evolve and are organized in systems but as far as we know, no systemic solution exists which integrates all of the above requirements. This paper proposes an innovative version of ArchiMate® for the SCADA components modelling purpose to enrich their collaborations and, more particularly, the description of their behavior endorsed in the cyber-policy. Our work has been illustrated in the frame of a critical infrastructure in the field of petroleum supply and storage networks.

Keywords- ArchiMate; metamodel; SCADA; multi-components system; trust; petroleum supply chains: critical infrastructure.

I. INTRODUCTION

Up to now, components represented at the business layers [1][2][7][8] have been considered human actors playing business roles. However, rising security requirements for the management of heterogeneous and distributed architectures calls for a rethinking of distribution of the security procedures in both: human and software autonomous entities. Although having been handled by human employees for years, the management of complex systems, nowadays, needs to be shared with intelligent software items, often perceived being more adapted to act in critical situations.

This statement is enforced by the characteristic ability of the component to act autonomously in open, distributed and heterogeneous environments, in connection or not with an upper authority. Acknowledging this situation, we are forced to admit that SCADA [4][39] components are no longer to be considered only as basic isolated solution deployed to support business activities, but that they are part of crisis reaction strategy [29]. Since then, acquiring an innovative enterprise architecture framework to represent the behaviors of such components appears fully justified in view of the arising cyber protection principles and required by the

practitioners, especially the ones engaged in the management of those critical infrastructures security protection.

In this paper, we propose to explore ArchiMate® and to redesign its structure in order to fit with component software actors' specificities and domain constraints. The main focus concerns the design and the consideration of the policies that are centric concepts related to the activation of component's behaviors. All along the modeling of the SCADA system and the definition of the policies according to these models, we are going to illustrate the theory with a case study related to the petroleum supply chain, and more specially the specific functions of Crude Oil Supply and Crude Oil Storage and Distribution. This extended case is introduced in Section 2. In Section 3, we will review the SCADA components metamodel and the SCADA layers for crisis management and we will model the concept of policy [22][24] that represents the engine of the component modeling framework in Section 4. Section 5 provides related works and Section 6 concludes the paper.

II. RELATED WORKS

Literatures explain methodologies to model Multi-Agent System (MAS) and their environments as a one layer model and give complete solutions or frameworks. Gaia [1] is a framework for the development of agent architectures based on a lifecycle approach (requirements, analysis, conceptualization and implementation). AUML [6] and MAS-ML [2] are extensions of the UML language for the modelling of MAS but do no longer exist following the release by the OMG of UML 2.0 [11][12] supporting MAS. Prometheus [7] defines a metamodel of the application layer and allows generating organizational diagrams, roles diagrams, classes' diagrams, sequences diagrams and so forth. It permits to generate codes but does not provide links between diagrams and therefore makes it difficult to use for alignment purposes or with other languages (e.g., MOF [3], DSML4MAS [5]). CARBA [15] provides a dynamic architecture for MAS similar to the middleware CORBA based on the role played by the agent. Globally, we observe that these solutions aim at modelling the application layer of MAS. CARBA goes one step further introduces the concept of Interface and Service. This approach is closed to the solution based on ArchiMate® that we design in our proposal but offers less modelling features. As we have noticed that agent systems are organized in a way close to the enterprises system, our proposal analyses how an

enterprise architecture model may be slightly reworked and adapted for MAS. Therefore, we exploit ArchiMate®, which has the following advantages to be supported by The Open Group. It has a large community and proposes a uniform structure to model enterprise architecture. Another advantage of ArchiMate® is that it uses referenced existing modelling languages like UML.

III. COMPLEX OIL DISTRIBUTION TEST MODELING

To represent the modelling of SCADA components metamodel and policy generation, we are going to illustrate this paper with the reference case study presented in [24]. The *crude oil distribution* [22] presented at the Section 2 includes both the oil supply and product distribution SCADA systems. Interconnection amongst Remote Terminal Unit (RTU) [31] of those SCADA is achieved using MTU [30]. The acronym MTU stands for Master Terminal Unit and its main purpose is to accept the different inputs from the remotely connected devices and to transmit these inputs over the rest of the network. Using the ArchiMate® for SCADA system theory introduced in previous Sections, the SCADA RTU of the crude oil distribution SCADA from the distributed plants may be modeled as illustrated on Figure 1.

As illustrated in this figure, both layers of the RTU are represented, the COS SCADA RTU Networks Organization (RTU-COSNO) and the COSSCADA RTU Networks Application (RTU-COSNA). At the COSNO layer, the crude oil network SCADA is composed of Crude oil portfolio that is assigned to *Call for IN* (aka *Organizational alert IN*), of application RTU monitoring services (e.g., *Moni SEGUA* and *Moni SEBAT* (Figure 1)), of RT information that impacts the generation of RTU behavioral policies. On the other side, the SEGUA and the SEBAT RTU (for instance) are represented as *actor* of the RTU organization and are composed of RTU *network console* dedicated to the SCADA management [38]. Both later are associated to the artifact modelled by the orange box that correspond to a collaboration between both SCADA functions, the crude oil supply and the product storage and product distribution. At the COSNA layer, four RTU/technical layer are modelled, respectively the SEGUA, SEBAT, RPBS and REVAP. The structure of this RTU/Technical layer is naturally always quite the same and is composed of the *technical monitoring service* (corresponding to the core of the RTU such as commonly addressed by the literature [21]), and of the interface named “in-[RTU/network location]”, which aims at connection the monitoring service with the RTU application itself. As illustrated at the level of the SEBAT model, the RTU application is potentially connected with the others RTU’s applications artefacts (cf. *two-ring* symbol). As summary, in this Section we have presented a metamodel for SCADA systems. This metamodel allows representing all the components of the SCADA following three layers: the

organization, the application and the technical layers. Those models offer the advantages to easily figure out the structure of the concepts and their interconnections and thereby, to easily capture the interconnection between the components within a SCADA and among two or many SCADA’s or SCADA functions. Given those advantages, the next Section explains how management policies may be designed and defined according to instance of this SCADA metamodel. Concretely, the usage of the metamodel has been illustrated trough a crude oil supply and distribution plan SCADA and connections have been depicted among the crude oil supply and the crude oil storage and distribution function of the SCADA system.

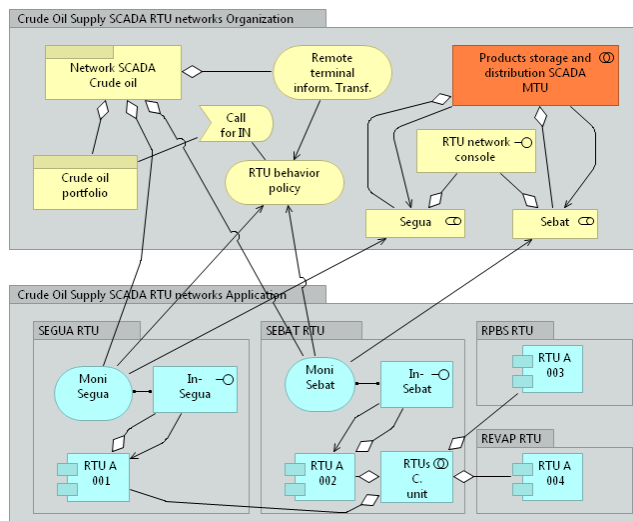


Figure 1. Crude oil SCADA MTU-RTU instance.

IV. SCADA POLICY MANAGEMENT

Based on ArchiMate® SCADA metamodel presented in Section 3 and illustrated by the *Crude Oil Supply SCADA*, this Section introduces the artefact of policy model by ArchiMate. Two types of Policy are depicted: the *Cognitive* and the *Permissive* [12].

A. Policy family

At the *Organizational Layer*, *Policy* can be represented as an *UML Use Case* [11] where concepts of *Roles* represent the *Actors* which have *Responsibilities* in the *Use Case*, and the *Collaboration* concepts show the connections between them. Concepts of *Products*, *Value* and *Organizational Service* provide the *Goal* of the *Use Case*. *Pre-* and *Post-conditions* model the context of the *Use Case* and are symbolized in the metamodel by the *Event* concept (pre-condition) and the *Organizational Object* (pre/post-condition). In the *Application Layer*, *Application Policy* is defined as the realization of *Responsibilities* by the *Application Domain* in a configuration of the *Data Domain*. UML provides support

for modelling the behavior performed by the *Application Domain as Sequence Diagram*. Configuration of the *Data Domain* can be expressed as *Pre-conditions* of the *Sequence Diagram* and symbolized by the execution of a test-method on the lifeline of the diagram. The metamodel designed in Section 3 has allowed providing the SCADA operators and managers with a holistic and integrated view of the SCADA architecture building blocks. In practice and to have policy extracted according to the metamodel concepts interconnections, this SCADA metamodel firstly needs to be instantiated for each architecture components. This step is achieved by shaping the component according to the three abstractions typically advocated by the enterprise architecture paradigm [11][12] and [13]. This allows discovering the building artefacts of the components as well as the connections amongst the components artefacts. An example of this instantiation is represented in Figure 1. The representation of each component implies paramount outcomes for the SCADA [31] operator since it confers to the latter a global functional insight of each component irrespective of any implementation or vendors' influence. The unitary SCADA [28] component models are then used in the second step to picture out the global structure of the SCADA architecture and of the connections, in terms of *policies*, amongst the components of the architecture. Previous works [40] highlights the two families of policies recovered in SCADA [29]: *Permissive policies* and *Cognitive policies*. *Cognitive Policies (CP)* [12] represent policies which govern the behavior of one artefact of the component architecture. This policy specifies the rule that the *Responsible* artefact needs to follow to execute a defined activity in a specific context. This rule is dictated by the artefact which exists in the same component or in another one. The artefact which generates the policy is the *Master* and the one, which execute it is the *Slave*. The *Cognitive Policy* morphology is articulated on the following set of attributes (perceived by [13]): Master artefact, Slave artefact, Master component, Slave component, Behaving rule, Trigger item, Usage context, Priority extension.

The application schema of a CP, as presented in Figure 1, obeys the two following controls: (1) the communication path is from a *Master* structural concept to a *Slave* behavioral concept or (2) the communication path is from a *Master* behavioral artefact to another *Slave* behavioral artefacts. Figure 2. They represent policies which govern the knowledge acquisition rules from the *Master* to the *Slave* artefact [14]. This knowledge acquisition traditionally takes the form of SCADA states data accessed or provided in order to provide the *Responsible* with the access (of *in*, *out*, *in_out* types [16]) to successive *Cognitive Policies* in case of occurring *events*. The *Permissive Policies* morphology is articulated on the following set of attributes [(perceived by [15]): Master artefact, Slave artefact, Master component, Slave component, Permission rules, Pre-permission conditions, Master permission cardinality, Slave permission cardinality, and Cognitive constraints - sustained

by *Cognitive Policy*. The application schema of a CP, as highlighted in Figure 1, obeys the two following controls: (1) the communication path is from a *Master* structural artefact to a *Slave* informational artefact or (2) the communication path is from a *Master* behavioral artefact to a *Slave* informational artefact.

B. Policy identification method

Designing automatic management strategy requires a rigorous two phase's policy elaboration mechanism, respectively the policy scheme identification and the policy scheme formalization.

1) Policy scheme identification step

The first step is itself structured in three phases. The first one aims at identifying the structure of the CI architecture in terms of unitary modules (components), including their three layers of abstraction build upon the SCADA [27] metamodel (i.e., organization, application, and technical). The second phase aims at identifying the external parameters of the CI (Critical Infrastructure) such as potential threat probes and indicators that may impact the CI normal functioning (flood, hijacking, etc.), the physical environment, and/or the contractual SLA (service level agreement). The third phase aims at identifying the reaction policies which may be of two types: *Cognitive* (artefact of a CI component which needs information from succeeding artefacts –Figure 2) or *Permissive* (artefact of a CI component which needs permission upon the succeeding lower layer artefacts – Red connections on Figure 1). Both types of policies are explained in [35][36][37][20].

2) Policy scheme formalization step

After policies being identified, the second step of the method aims at formalizing policy scheme using a three phases approach. The first one aims at depicting the *Master-Slave communication* artefacts (organization-organization, organization-technical, and technical-technical), the second aims at identifying the *cognitive* and *permissive behaviour* based on the automatic reaction strategy, and the last one aims at formalizing the policies accordingly. This latter is function of the policy type and is achieved, on one hand, with the inter-artefacts knowledge requirement, external probes and monitoring tools in case of *cognitive* policy and with the reaction strategy with the requirement of access to artefacts in case of *permissive policy*.

3) Inter Critical Infrastructures Study Case

This second part of the case study aims at defining *cognitive* and *permissive* security policies supported by the MTU-RTU model from Figure 1. In [32], authors argue that SCADA system network is *different from general network environment due to its operational environment in national infrastructure*. Therefore, in such a context, the SCADA system needs important *broadcast capability*, which must be highly protected. Among these protection mechanisms are the *key management schemes* [32][33][34] that also have to

support the multicasting messages protection. Figure 1 illustrates the modelling of *permissive* and *cognitive* policies related to the Key Management Exchange, such as expressed by [32] among the MTU dedicated to the *crude oil supply* function and the RTU from this function and from the *storage and distribution* function. This field has already been tackled by many researches such as [33] [34]. [32] has been preferred for this case illustration provided that it reduces consistently the number of keys to be stored and provides multicasting and broadcasting communication for efficient and stable operation of SCADA systems. Hence, the policies dedicated to the management of this broadcasting will be defined in the following.

Three constraints related to the key management broadcasting mechanism related to the SCADA architecture have been defined by [18][19][31] and need to be considered along the modelling of the policies: (1) the computational capacity limit which may be represented as an artefact of a type data object at the application layer of the MTU, (2) the low data transmission rate which is also a concept related to the MTU by means of a data object, and (3) the real-time processing that needs to be consider to prevent data processing delay and which may be represented as a data object from the RTUs structures. From Figure 2, we observe the following list of policies: Firstly at the organization layer: the MTU Management policy (1), and secondly at the application layer: the crude oil supply policy /MTU S1 (2) and /RTU S1 (3) and the crude oil storage /RTU St01 (4). (1) is existing at the organizational layer and is realized by (2) at the technical layer [17]. This first family of policies (1) accesses the key exchange value that represents the real encryption parameter introduced by the SCADA operator through the dedicated interface (aka MTU screen). The later aims at supporting the key management service which is represented by the key management unit artefact. It has the right of a type in, out, in/out on the key set MTU, key set S1 and key set st01 data objects (Table I).

TABLE I. PERMISSIVE POLICY FOR ATTRIBUTES' NAME AND ATTRIBUTES' ID

Attribute Name	Attribute's ID
Master artefact	Organizational service
Salve artefact	Data objects
Master component	Key management unit
Slave component	key set MTU, key set S1 and key set st01
Permission rules	In/Out/In-Out
Permission conditions	\exists of set of Master-Slave Associations
Master permission cardinality	1
Slave permission cardinality	1..n
Cognitive constraints	Key exchange values

This policy is a *permissive* policy provided that it gives an authorization. The second family of policies depicted through the RTU-MTU model concerns the application layer policies named MTU S1, RTU S1 and RTU St01. These policies are directly assigned and dictate the expected behavior of the application function (in this case, the selection of the encryption ID and system). These policies correspond to *Cognitive*. They express that 1 of the MTU S1, RTU S1 or RTU st01 policy (master artefact) may Select key Encryption ID, May enforce Key Encryption ID and Algorithm [32] related to the application MTU S1, RTU S1 and RTU st01 (slave artefact) if there exist at least one permission of a type Comp.-capa.-Limit, trans.-rate, real-ti.-proc. To process the above *Cognitive* policies, the MTU S1, RTU S1 and RTU St01 policies required to collect information related to the key by directly accessing the respective key set data object artefact, to know: Key Set MTU, Key Set S1 and Key Set St01. This collection of information is possible if the appropriate *permissive* policies are defined and deployed in the SCADA. For the sake of clarity, the later have not been represented in the MTU-RTU model (Table II).

TABLE II. COGNITIVE POLICY ATTRIBUTES' NAME AND ATTRIBUTES' ID

Attribute Name	Attribute ID
Master artefact	Application service
Slave artefact	Application
Master component	Policy MTU S1, Policy RTU S1, Policy RTU St01
Slave component	MTU S1, RTU S1, RTU St01
Permission rules	Select key Encryption ID - Enforce Key Encryption ID and Algorithm
Permission conditions	Comp.-capa.-Limit, trans.-rate, real-ti.-proc
Master permission cardinality	1
Slave permission cardinality	1
Cognitive constraints	\exists of Technical MTU S1, RTU S1, STU St01.

V. EVALUATION

Although the MTU S1 and RTU S1 are SCADA artefacts from the same SCADA (*crude oil supply* SCADA), RTU St01 is an artefact from another function, i.e.: *crude oil storage and distribution*. The later consists in an alternative SCADA system. Using the *ArchiMate*® metamodel for modelling SCADA policies of a type *cognitive* or *permissive* at both the organizational and the technical layers has allowed representing heterogeneous SCADA policies from two different SCADA using the same language (i.e.: *ArchiMate*® for SCADA systems).

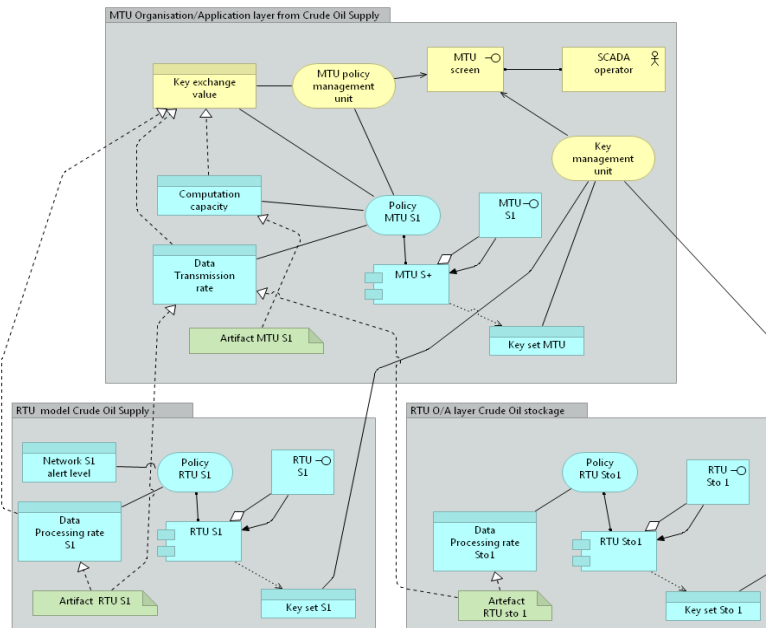


Figure 2. MTU-RTU Key distribution case.

VI. CONCLUSIONS AND FUTURES WORKS

The paper proposes an integrated approach for modelling the SCADA based on the enterprise architecture modelling language and more specially *ArchiMate*® which has been particularly tailored for SCADA systems [23][24][25][26]. Based on a dedicated metamodel, the paper has demonstrated how technical, application and organization policies could be designed and metamodeled, especially regarding the policy management for interconnected SCADA systems for two of its functions. All along the modelling of the SCADA model and the definition of the policies according to these models, we have illustrated the theory with a business case study related to the petroleum supply chain, and more specially the specific functions of *crude oil supply* and *crude oil storage and distribution*. The main future works consists in elaborating a concrete prototype to sustain the metamodel usage and the deployment in real usage settings. The metamodel drawback concerns the lack of a dedicated specialization of the *ArchiMate* language. This extension of the framework would highly enrich the decision making mechanism in CI.

REFERENCES

[1] F. Zambonelli, N. R. Jennings, and M. Wooldridge, 2003, "Developing multicomponent systems: The Gaia methodology". *ACM Trans. Softw. Eng. Methodol.* 12, 3 (July 2003), 317-370.
 [2] V. Torres da Silva, R. Choren, and C. J. P. de Lucena, 2004, "A UML Based Approach for Modeling and Implementing Multi-Component Systems". In

Proceedings of the Third AAMAS, Vol. 2. IEEE Computer Society, Washington, DC, USA, 914-921.
 [3] J. J. Gomez-Sanz, J. Pavon, and F. Garijo, 2002, "Metamodels for building multi-component systems". Proceedings of ACM symposium on Applied computing (SAC '02). ACM, New York, , USA, 37-41.
 [4] G. Beydoun, C. Gonzalez-Perez, G. and Low, B. Henderson-Sellers, 2005, "Synthesis of a generic MAS metamodel". *SIGSOFT Softw. Eng. Notes* 30, 4, 1-5.
 [5] AUML (Component UML), <http://www.auml.org/> [accessed: 2015-03-12]
 [6] G. Guemkam, C. Feltus, P. Schmitt, C. Bonhomme, D. Khadraoui, and Z. Guessoum, 2011, "Reputation Based Dynamic Responsibility to Agent Assignment for Critical Infrastructure". In Proceedings of the 2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology - Volume 02 (WI-IAT '11), Vol. 2.
 [7] C. Feltus, E. Dubois, E. Proper, I. Band, and M. Petit, 2012, "Enhancing the ArchiMate® standard with a responsibility modeling language for access rights management". In Proceedings of the Fifth International Conference on Security of Information and Networks (SIN '12). ACM, New York, NY, USA, 12-19.
 [8] Daneels, Axel, and Wayne Salter, "What is SCADA." International Conference on Accelerator and Large Experimental Physics Control Systems. 1999.
 [9] C. Feltus, M. Ouedraogo, and D. Khadraoui, "Towards Cyber-Security Protection of Critical Infrastructures by Generating Security Policy for SCADA Systems", The 1st International Conference on Information and Communication Technologies for Disaster Management (ICT-DM'2014), 24-25/3/2014, Algiers, Algeria.

- [10] Khadraoui, D., and Feltus, C., "Critical Infrastructures Governance, Exploring SCADA Cybernetics through Architected Policy Semantic," Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on , pp.4766-4771.
- [11] Blangenois, J., Guemkam, G., Feltus, C., and Khadraoui, D., "Organizational Security Architecture for Critical Infrastructure," Availability, Reliability and Security (ARES), 2013 Eighth International Conference on , vol., no., pp.316,323, 2-6 Sept. 2013
- [12] C. Feltus, D. Khadraoui, and J. Aubert, "A Security Decision-Reaction Architecture for Heterogeneous Distributed Network". 2012 Seventh Int. Conference on Availability, Reliability and Security. IEEE.
- [13] J. Sabater, and C. Sierra, "Review on computational trust and reputation models", Artificial Intelligence Review, vol. 24, no. 1, pp. 33–60
- [14] W. Jiao, and Z. Shi; "A dynamic architecture for multi-agent systems", Technology of Object-Oriented Languages and Systems, 1999. TOOLS 31. pp.253-260.
- [15] G. Eason, B. Noble, and I.N. Sneddon, "On certain integrals of Lipschitz-Hankel type involving products of Bessel functions," Phil. Trans. Roy. Soc. London, vol. A247, pp. 529-551, April 1955.
- [16] J. C. Maxwell, "A Treatise on Electricity and Magnetism", vol. 2. Oxford: Clarendon, 1892.
- [17] Davidson, Euan M., et al. "Applying multi-agent system technology in practice: automated management and analysis of SCADA and digital fault recorder data." Power Systems, IEEE Transactions on 21.2 (2006).
- [18] C. Feltus, D. Khadraoui, B. de Rémont, and A. Rifaut, "Business Governance based Policy Regulation for Security Incident Response", International Conference on Risks and Security of Internet and Systems 2-5/7/2007, Marrakech, Morocco.
- [19] C. Feltus, "Conceptual Trusted Incident-Reaction Architecture", The 6th International Network Conference 2010 (INC2010), June 2010, Heidelberg, Germany
- [20] Y. Yorozu, M. Hirano, K. Oka, and Y. Tagawa, "Electron spectroscopy studies on magneto-optical media and plastic substrate interface," IEEE Transl. J. Magn. Japan, vol. 2, pp. 740-741, August, p. 301, 1982.
- [21] C. Feltus, "Preliminary Literature Review of Policy Engineering Methods - Toward Responsibility Concept", International Conference on Information & Communication Technologies: from Theory to Applications (IEEE ICTTA2008), Damascus, Syria.
- [22] M.S. Neiro, and J. M. Pinto, "A general modeling framework for the operational planning of petroleum supply chains", Computers & Chemical Engineering, Volume 28, Issues 6–7, 2004, Pages 871–896.
- [23] C. Feltus, M. Petit, and M. Sloman, "Enhancement of Business IT Alignment by Including Responsibility Components in RBAC", 5th International Workshop on Business/IT Alignment and Interoperability (BUSITAL 2010), 2010, Hammamet, Tunisia.
- [24] G. Guemkam, J. Blangenois, C. Feltus, and D. Khadraoui, "Metamodel for Reputation based Agents System - Case Study for Electrical Distribution SCADA Design", 6th ACM International Conference on Security of Information and Networks (ACM SIN 2013), November 2013, Aksaray, Turkey.
- [25] Feltus, C. and Petit, M., "Building a Responsibility Model Including Accountability, Capability and Commitment", Availability, Reliability and Security, 2009. ARES '09. International Conference on , vol., no., pp.412,419, 16-19 March 2009.5
- [26] Gateau, B., Khadraoui, D., and Feltus, C., "Multi-agents system service based platform in telecommunication security incident reaction," Information Infrastructure Symposium, 2009. GIIS '09. Global , vol., no., pp.1,6, 23-26 June 2009
- [27] Patel, S. C., Bhatt, G. D., and Graham, J. H. (2009), "Improving the cyber security of SCADA communication networks". Communications of the ACM, 52(7), 139-142.
- [28] Bailey, D., and Wright, E, (2003), "Practical SCADA for industry". Newnes.
- [29] Donghyun C.I, Hakman K., Dongho W., and Seungjoo K., "Advanced Key-Management Architecture for Secure SCADA Communications," Power Delivery, IEEE Transactions, vol.24, no.3, pp.1154,1163, 2009
- [30] Beaver, C., Gallup, D., Neumann, W. and Torgerson, M. (2002), "Key management for SCADA," Technical report, Sandia.
- [31] R. Dawson, C. Boyd, E. Dawson, J. Manuel, and G. Nieto, "SKMA A Key Management Architecture for SCADA Systems," In Proc. Fourth Australasian Information Security Workshop, Vol. 54, pp. 138-192, 2006.
- [32] C. Feltus, and D. Khadraoui, "On Designing Automatic Reaction Strategy for Critical Infrastructure SCADA System", 6th ACM International Conference on Security of Information and Networks (ACM SIN 2013), 26-28/11/2013, Aksaray, Turkey.
- [33] <http://pubs.opengroup.org/architecture/archimate2-doc/>
- [34] Prometheus Methodology. <http://www.cs.rmit.edu.au/agents/SAC2/methodology.html>
- [35] Chan, M. L. (1991, April), "Interrelation of distribution automation and demand-side management". In Rural Electric Power Conference, 1991. Papers Presented at the 35th Annual Conference (pp. B1-1). IEEE.
- [36] Kato, K., and Fudeh, H. R. (1992), "Performance simulation of distributed energy management systems. Power Systems", IEEE Transactions on, 7(2), 820-827.
- [37] Choi, D., Kim, H., Won, D., and Kim, S, 2009, Advanced key-management architecture for secure SCADA communications. Power Delivery, IEEE Transactions on, 24(3), 1154-1163.

Impact of Late-Arrivals on MPI Collective Operations

Christoph Niethammer, Dmitry Khabi, Huan Zhou, Vladimir Marjanovic and José Gracica

High Performance Computing Center, University of Stuttgart

Stuttgart, Germany

Email: {niethammer, khabi}@h1rs.de

Abstract—Collective operations strongly affect the performance of many MPI applications, as they involve large numbers, or frequently all, of the processes communicating with each other. One critical issue for the performance of collective operations is load imbalance, which causes processes to enter collective operations at different times. The influence of such late-arrivals is not well understood at the moment. Earlier work showed that even small system noise can have a tremendous effect on the collective performance. Thus, although algorithms are optimized for large process counts, they do not seem to tolerate noise or consider delay of involved processes and even a small perturbation from a single process can already have a negative effect on the overall collective execution. In this work, we show a first detailed study about the effect of late arrivals onto the collective performance in MPI. For the evaluation a new, specialized benchmark was designed and a new metric, which we call delay overlap benefit, was used. Our results show, that there is already some potential tolerance to late arrivals - but there is also a lot of room for future optimizations.

Keywords—collectives; late-arrivals; benchmarking; MPI

I. INTRODUCTION

Collective operations strongly affect the performance of many Message Passing Interface (MPI) applications, as they involve large numbers, usually all, of the processes communicating with each other. One critical issue for the performance of collective operations is load imbalance, which causes processes to enter collective operations at different times. The influence of such delayed processes is not well understood at the moment. Earlier work showed that even small system noise can have a tremendous effect on the collective performance [1] [2]. So, though algorithms are optimized for large process counts [3], they do not seem to tolerate noise or consider delay of involved processes and thus even a small perturbation from a single process can already have a negative effect on the overall collective execution.

The MPI 3.0 standard introduced non-blocking collective operations which give the opportunity to speed up applications by allowing overlap of communication with computation [4], reducing the synchronisation costs of delayed processes as well as the effects of system noise. Many MPI programs are written using non-blocking point-to-point communication operations and application developers are familiar with managing this process using request and status objects. Extending this to

include collectives allows programmers to straightforwardly improve application scalability.

In contrast to the already existing blocking collectives, the non-blocking counterparts require the MPI implementations to progress the communication task in parallel to computations. This is a non-trivial task, even if the network hardware provides support for offloading network operations from the CPU, e.g., message buffers may have to be refilled for large messages or more complex collective operations need multiple communication steps. The Cray XE6 and XC30 platforms feature a special “asynchronous process engine” for this, which uses spare hyperthreads (XC30) or dedicated CPU cores (XE6) for the required operations [5].

This work analyses and emphasizes the effect of late arrivals on collective operation in MPI for large number of processes. Therefore, a benchmark and metric for evaluation and detection of effects caused by late arrivals are introduced. The obtained results point to potential for improving performance by solving the issue of late arrivals.

This work is structured as follows. Section II, describes the testing methodology and the micro benchmark suite, which we designed specifically to study the impact of late arrivals, i.e. delay, on collective performance. At the begin of section III, we define a metric to quantify the amount of tolerated delay. Then results for different, application relevant collective operations are presented and evaluated on basis of absolute times as well as the delay overlap metric.

II. METHODOLOGY AND BENCHMARK DESCRIPTION

To study collective operations with respect to late arrivals, a micro benchmark suite was designed. This requires the use of a global clock, which is chosen to be the one of process with MPI rank 0. For this purpose the micro benchmark suite determines the clock offsets between process zero and all other processes. Based on the global time, the benchmark performs the following tasks:

- Measures start and end times of all involved MPI processes.
- Determines earliest start and latest end time over all involved MPI processes.

The design of the benchmark suite allows for easy extendibility and addition of new benchmarks. The following MPI collective operations are included at the moment: Barrier, Ibarrier, Bcast, Ibcast, Reduce, Ireduce, Allreduce, Iallreduce, Alltoall and

laltoall. Within this work, results for blocking and non-blocking barrier, allreduce and altoall operations are reported. Each benchmark is run with different number of processes and different data sizes. Each benchmark is run initially for several times before the real measurement is performed, to warm up the network, CPUs, etc. Then, the times for the real benchmark runs are recorded.

A. Clock offset determination

The local clocks of different processors report different times as they are not perfectly synchronized. They may even run at slightly different speeds [6] [7], which we do not take into account. This simplification is acceptable because the benchmark runs only for a relatively short time and a verification shows, that there is no significant change in the time differences over it's runtime. To compare the measured times, the error between the clocks has to be taken into account [8]. For the collective benchmarks, we consider the clock offset σ defined as the constant difference between the locally measured time t and the time measured at the remote processor t' at the time point when the benchmark is started

$$t' = t + \sigma . \quad (1)$$

A modified ping pong experiment is used to determine the clock offset following Cristian's algorithm [9]: A root process sends a request to another process, which answers with his current local time. We improve the accuracy by adding another timer allowing to determine the timer delay, which is the time required to obtain the current time itself. From this experiment the ping pong latency λ_p and timer delay Δ , are obtained, see Figure 1.




Figure 1. Modified ping pong experiment to determine ping pong latency λ_p , timer delay Δ and clock offset on the basis of the remote time t'_r , which is assumed to be taken at the mid of the ping pong.

To obtain the clock offset σ between local clock t and remote clock t' defined in (1), the time t'_r , which is assumed to be at the mid of the ping pong, is measured. The clock offset is then given by

$$\sigma = t'_r - t_0 - (\lambda_p + \Delta)/2 , \quad (2)$$

where the timer delay Δ is obtained via

$$\Delta = t_2 - t_1 . \quad (3)$$

To verify the correctness and to obtain an estimate for the error in the obtained clock offsets intra node times can be

compared, which should not vary much. As can be seen in Table I, the clock offsets between rank 0 and all processes residing on one node are the same with a standard deviation of not more than $\pm 2 \mu\text{s}$ in 100 measurements. In contrast, the clocks of different nodes vary by more than 10 ms between each other.

Table I. DETERMINED AVERAGE CLOCK OFFSET $\bar{\sigma}$ AND STANDARD DEVIATION σ_σ FOR A BENCHMARK RUN WITH 12 PROCESSES AND 4 PROCESSES PER NODE BASED ON A SET OF 100 MEASUREMENTS. (RESULTS OBTAINED ON HERMIT SYSTEM AT HLRS, SEE SECTION III)

rank	$\bar{\sigma}$ [s]	σ_σ [s]
0	+0.000000	0.000000
1	+0.000000	0.000000
2	+0.000000	0.000000
3	+0.000000	0.000000
4	-0.017258	0.000002
5	-0.017258	0.000001
6	-0.017258	0.000001
7	-0.017258	0.000002
8	-0.011140	0.000002
9	-0.011140	0.000002
10	-0.011140	0.000002
11	-0.011140	0.000002

B. Initial synchronization

A synchronization of all processes is done at the beginning of each benchmark run using a barrier. The synchronization is not perfect as can be seen in Figure 2 and the processes finish the barrier at slightly different times. The time difference between the processes at the exit of the barrier is in the order of $4 \mu\text{s}$ for 32 processes and the observed exit time pattern may be the result of a tree algorithm [10]. But so far, there is no better way of synchronization. Measuring the time differences and trying to improve the sync using delays with an accuracy of $1 \mu\text{s}$ for faster processes, resulted in even worse synchronization.




Figure 2. Vampirtrace image of synchronization barrier (red) before the benchmark is executed on Cray XE6 (32 PEs on 2 nodes). The processes leave the barrier in a time shifted front as can be seen by the timer calls enclosing the benchmarking loop (blue). Results are collected afterwards with an MPI_Gather operation (green).

For each measurement the process id (integer) as well as its start and end time (double) are stored, which requires $S = 20N$ Bytes of storage. The stored times are times corrected on the basis of the clock offset determined before. If not mentioned explicitly global times for the collective operations are reported, which is the time between the start time of the first process entering and the end time of the last process finishing the collective.

C. Delaying of single process

Load imbalances in programs cause some processes to enter collectives later than the rest. To study the influence of such late-arrivals on the overall collective time, the benchmark suite allows to delay one processes by a given amount of time, see Figure 3. The delay is implemented based on the POSIX `gettimeofday` function, providing a microsecond accuracy.




Figure 3. Processes are except one synchronized at time t_a and enter the collective. The one delayed process enters the collective at time $t_b = t_a + \delta$.

III. RESULTS

The influence of different delay times and the number of processes is studied. Blocking collectives and their non-blocking counterparts are compared side by side. Beside the actual collective times the benefit b of internal overlap of the delay with communication within the collectives itself will be examined:

$$b = \frac{t_0 + \delta - t_\delta}{t_\delta}, \quad (4)$$

with t_0 being the collective time for no delay and t_δ the collective time for delay δ . A positive benefit is found when there is overlap potential within the collective operation. A negative value means, that the delay even results in additional cost compared to a synchronized collective which is started after waiting delay time.

In the following, results for different collective operations on the Hermit system at HLRS are reported. Hermit is a Cray XE6 system with 3552 dual-socket compute nodes and a total of 113 664 cores, which are connected via the Gemini 3D Torus network. The native Cray MPI implementation optimized for this system in combination with the GNU compiler was used for all tests.

All benchmarks were run during normal operation mode of the system so that other jobs on the system influenced the process placement and network usage. Benchmark runs were performed up to a maximum of 16 384 processes and were grouped into jobs with the same processor count. We report the found minimum values for the global times within 100 measurements. This is responsible for some outlying data points, even if multiple measurements were performed to reduce this effect. Obtaining the accurate minimum time for an operation under workload conditions is not always possible—especially for the longer benchmark runs using more processes, which get easily disturbed by other jobs.

For all measurements the MPI process with rank 0 was delayed. Most tree based algorithms—usually using rank 0 as tree root—should be badly affected by this choice, if they do not switch over using another process as the tree root.

A. Barrier

The first collective studied is the barrier. As the barrier is used for synchronization within the benchmark suite, the understanding of this operation is essential. While the time for `MPI_Barrier` is measured straightforward, the time for `MPI_Ibarrier` includes the time for the corresponding `MPI_Wait`.

A wide variety of different barrier algorithms exists [10]. Depending on the algorithm and the hardware support used within the implementation, different algorithms may profit differently. On the one hand, for example the Central Counter barrier may hide the delay of a late arrival easily by concept, or the Binomial Spanning Tree Barrier could intelligently assign the delayed process to a node, which is involved in later communication steps. On the other hand, for example the Dissemination Barrier requires a ring like communication in each step—which will not tolerate a late arrival.

The results in Figure 4 show a nearly logarithmic scaling of the blocking and non-blocking barrier operation up to approximately 2048 processes. For higher process counts, the behaviour seems to have a linear scaling. But we note here that a single cabinet of the Hermit system has 96 nodes with a total of 3072 cores. Jobs exceeding this number of processes are more likely to be spread around the system and therefore affected by network contention caused by other applications. So, finding the minimum time for the barrier operation with our benchmark may not have provided the correct result in this case.

The delay benefit as defined in (4) of the `MPI_Barrier` and `MPI_Ibarrier` for different delay times, where the delayed rank was always rank 0, is shown in Figure 5. As the benefit is mostly positive the implemented blocking and non-blocking barrier algorithm already seem to tolerate smaller delays. The non-blocking version `MPI_Ibarrier` seems to perform slightly better than the blocking variant here. Figure 5 shows an change in behaviour at 1024 processes: While at the beginning smaller delays have a higher overlap benefit, for more processes a larger benefit can be seen for longer delays. It is unclear if at this point an algorithm switch occurs within the MPI implementation.

B. Allreduce

An important collective to aggregate data of multiple processes into a single value is the allreduce operation. It may be used to determine, e.g., global energies in molecular simulations, time step lengths in finite element based programs or residues in linear solvers. While the time for `MPI_Allreduce` is measured straightforward, the time for `MPI_Iallreduce` includes the time for the corresponding `MPI_Wait`.

Again, the influence of delaying the process with rank 0 for different number of processes is studied. Results for 8B messages and a delay of 50 μ s are presented in Figure 6. We see perfect logarithmic scaling up to 1024 processes, adding less than 5 μ s when doubling the number of processes. For larger process counts the scaling is worse and adds up to 100 μ s when doubling the number of processes. The behaviour for larger

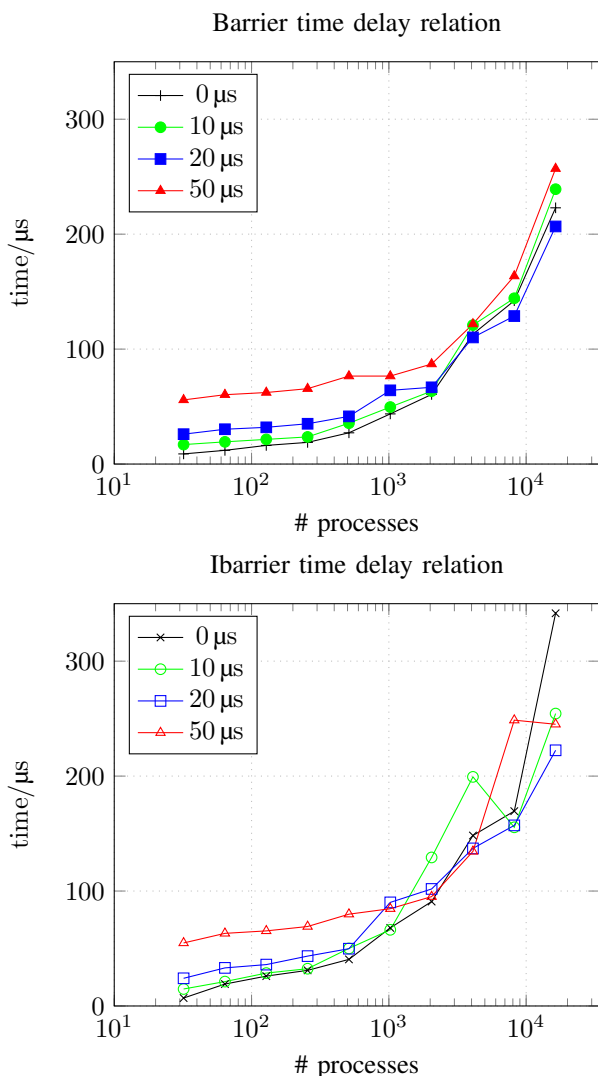


Figure 4. MPI_Barrier and MPI_Ibarrier global times for different delay times.

message sizes is similar. It is unclear how the synchronization barrier influences the behaviour, as we showed earlier that the processes do not exit from it perfectly at the same time and it does not scale well for larger process counts according to our benchmark results, too, see Figures 2 and 4.

We have to mention a data outlier for the non-delayed Allreduce/Iallreduce benchmark runs with 4096 processes—which were grouped within one job. The job collecting these data was likely disturbed by other jobs and seems not to have been able to find an accurate value for the minimum collective time.

The delay benefit of the blocking and non-blocking allreduce operations presented in Figure 7 shows slight overlap for smaller number of processes. For more than 1024 processors the delay has a negative effect onto the overall performance. The message size does not have an influence on the delay




Figure 5. Delay benefit of the MPI_Barrier and MPI_Ibarrier as defined in (4) for different delay times.




Figure 6. MPI_Allreduce (circles) and MPI_Iallreduce (squares) global times for 8B message size and a delay time of 50 μs (blue) together with perfectly synchronized reference data (black).

benefit for the chosen values. The peak for 4096 processes is caused by too high values for the perfectly synchronized collectives time t_0 .

C. Alltoall

The alltoall operation is another important collective pattern used in many parallel codes to distribute data in an application. It is the most time consuming collective operation but it may benefit a lot from intelligent algorithms, taking into account




Figure 7. Delay benefit of the allreduce collective for different message sizes at a delay time of 50 μs).

delayed processes.

The same measurements as that for the allreduce operation were performed. Results in Figure 8 show a nearly perfect linear scaling for the alltoall algorithm up to the maximum of 16384 processes used during the benchmarks. The message size has a strong influence on the execution time of the alltoall collective but does not affect the overall scaling behaviour.




Figure 8. MPI_Alltoall (circles) and MPI_Ialltoall (squares) global times for 8 B message size and a delay time of 50 μs (blue) together with perfectly synchronized reference data (black).

The results for the delay benefit for the alltoall collective, presented in Figure 9, show zero effect for small messages and an inconclusive behaviour for larger messages which may be caused by the fact, that our benchmark does not find the minimum time as already mentioned before. So we find slight decreases as well as huge gains in performance.




Figure 9. Delay benefit for the alltoall collective for different message sizes at a delay time of 50 μs.

IV. CONCLUSION AND OUTLOOK

In this paper, we have evaluated the impact of late arrivals, i.e. a delayed MPI process, on the performance of the collective operations MPI_(I)Barrier, MPI_(I)Allreduce and MPI_(I)Alltoall on Cray XE6. The results show that blocking and non-blocking collective barriers can tolerate small delays, i.e. hide a part of the load imbalance within an application. The collectives MPI_(I)Allreduce tolerate small delays for up to 1024 processes but is badly affected for larger processes counts. The MPI_(I)Alltoall operations tolerate small delays well for up to 1024 processes and the delays have no negative effects for large processes counts. The alltoall operation can profit a lot in some cases for larger message sizes, while we see no negative effects for small messages.

We have shown that the overlap availability of non-blocking collectives and benefit of the overlapping depends on the type of the collective operations, size of the communicator and the amount of data to be communicated.

This work shows that the state of the art implementation of the relatively new MPI 3.0 non-blocking collective specification in Cray MPI is mostly head up or better than their blocking counterparts. We expect new algorithms and hardware with better overlapping capabilities and communication offloading support in the future. Our preliminary work in this area shows already some potential to hide small delays of single processes for barrier, allreduce and alltoall operations. The techniques

for overlapping communication may also improve collective operations in the case of system noise.

For the future studies about other important collectives like bcast are planed as well as detailed analysis of delaying other processes than rank 0. Studies are planed to evaluate other MPI library implementations. Here open source implementations can provide insights into the algorithms as well as the cross over points between them for different message sizes and process counts, allowing better understanding of the results.

ACKNOWLEDGEMENT

This work has been supported by the CRESTA project that has received funding from the European Community's Seventh Framework Programme (ICT-2011.9.13) under Grant Agreement no. 287703. This work made use of computational resources provided by the High Performance Computing Center Stuttgart (Hermit). We thank P. Manninen (Cray Finland) and R. W. Nash (EPCC) for valuable discussions and assistance.

REFERENCES

- [1] T. Hoefler, T. Schneider, and A. Lumsdaine, "The Effect of Network Noise on Large-Scale Collective Communications," *Parallel Processing Letters*, Dec. 2009, pp. 573–593.
- [2] K. B. Ferreira, P. G. Bridges, R. Brightwell, and K. T. Pedretti, "The impact of system design parameters on application noise sensitivity," *Cluster Computing*, vol. 16, no. 1, 2013, pp. 117–129.
- [3] R. Thakur and R. Rabenseifner, "Optimization of collective communication operations in mpich," *International Journal of High Performance Computing Applications*, vol. 19, 2005, pp. 49–66.
- [4] T. Hoefler, A. Lumsdaine, and W. Rehm, "Implementation and performance analysis of non-blocking collective operations for mpi," in *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*, ser. SC '07. New York, NY, USA: ACM, 2007, pp. 52:1–52:10.
- [5] H. Pritchard, D. Roweth, D. Henseler, and P. Cassella, "Leveraging the Cray Linux Environment Core Specialization Feature to Realize MPI, Asynchronous Progress on Cray XE Systems," in *Proc. Cray User Group*, 2012.
- [6] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," *Commun. ACM*, vol. 21, no. 7, Jul. 1978, pp. 558–565.
- [7] S. J. Murdoch, "Hot or not: Revealing hidden services by their clock skew," in *In 13th ACM Conference on Computer and Communications Security (CCS 2006)*. ACM Press, 2006, pp. 27–36.
- [8] D. Becker, R. Rabenseifner, and F. Wolf, "Implications of non-constant clock drifts for the timestamps of concurrent events," in *Cluster Computing, 2008 IEEE International Conference on*, Sept 2008, pp. 59–68.
- [9] F. Cristian, "Probabilistic clock synchronization," *Distributed Computing*, vol. 3, no. 3, 1989, pp. 146–158.
- [10] T. Hoefler, T. Mehlan, F. Mietke, and W. Rehm, "A Survey of Barrier Algorithms for Coarse Grained Supercomputers," *Chemnitzer Informatik Berichte*, vol. 04, no. 03, Dec. 2004.

Remote Control and Monitoring of Smart Home Facilities via Smartphone with Wi-Fly

Omar Ghabar

School of Computing and Engineering
University of Huddersfield
Huddersfield, UK
e-mail: omar.ghabar@hud.ac.uk

Joan Lu

School of Computing and Engineering
University of Huddersfield
Huddersfield, UK
e-mail: j.lu@hud.ac.uk

Abstract— Due to the widespread ownership of smartphone devices, the application of mobile technologies to enhance the monitoring and control of smart home facilities has attracted much academic attention. This study indicates that tools already in the possession of the end user can be a significant part of the specific context-aware system in the smart home. The behaviour of the system in the context of existing systems will reflect the intention of the client. This model system offers a diverse architectural concept for Wireless Sensor Actuator Mobile Computing in a Smart Home (WiSAMCinSH) and consists of sensors and actuators in various communication channels, with different capacities, paradigms, costs and degree of communication reliability. This paper focuses on the utilization of end users' smartphone applications to control home devices, and to enable monitoring of the context-aware environment in the smart home to fulfil the needs of the ageing population. It investigates the application of an iPhone to supervise smart home monitoring and control electrical devices, and through this approach, after initial setup of the mobile application, a user can control devices in the smart home from different locations and over various distances.

Keywords—*smartphone; remote control; smart home; ad-hoc; context-aware.*

I. INTRODUCTION

Nowadays, smartphones are becoming increasingly successful in the area of context-aware services, such as smart home technology, computer-assisted homecare and healthcare, and other applications for social and safety purposes. Such services appear to recognise users' contexts, for example their location, current activities, identity and surrounding environment [1].

Additionally, there has been rapid growth in the field of remote control home automation, which enables smartphones to be used to control electronic devices in the household. An increasing number of home appliances, such as heating, TVs, lights, washing machines, A/V players and many other products can be embedded with computer technology and controlled by smartphones [2]. There are many forms of accessible media which can be used for wireless communication by remote control, for example, internet services, SMS with GSM, mobile phones and multi agent systems (MAS). These systems allow different actions to be performed, such as switching a device on or off,

controlling the level of sound or temperature and other functions.

Such wireless networks of sensors and actuators [3][4] have become an area of active research. Sensors and actuators can be used to improve the interaction between the client and the physical environment. It is thought that by 2020 [5], there will be between 50 and 100 billion appliances connected to the internet. Researchers are trying to promote changes to the mode by which each object is activated [6]. In addition to sensor technology, the mobile phone has become increasingly powerful in the field of computing as well as that of communication.

This paper defines and explains the design of an iPhone smart home system, using the Wi-Fly module, for monitoring safety in the home environment as well as for controlling electronic appliances in the same house. In effect, therefore, the prototype system is responsible for both monitoring and control, so this research offers a mixture of these two functions. The model is highly flexible and can be applied to several areas of study. This study describes the smart home architecture that can remotely monitor home devices and also provide real-time monitoring of home safety status through the implementation of a smartphone. It focuses on a system which ensures a high level of service for the intelligent home, including activity delay when the resident leaves the building. Thus, the proposed system provides a personal device for monitoring and control of different circumstances through the iPhone interfaces which are used to process context-aware information. The main benefit of the iPhone, compared to the use of a PC, is that it can be used from different locations. Various context-aware data, to be monitored using Wi-Fly, were tested in the laboratory using the iPhone.

There are very few studies offering a complete explanation of the actuators, sensors and robotics built into context-aware system design, as this field includes a wide variety of technologies and also requires several application situations which might be difficult to provide and assess. Analysis of specific problems is still at the theoretical stage, and few critical studies have yet offered a clear and extensive description relating to quality of context [7].

There has also been little previous research and development work carried out into aspects of the smart home, and even the number of users has not been clearly determined. Most research has concentrated on feasibility

studies and the use of interactive communication technologies with context-aware applications [8].

This prototype design attempts to offer a solution for the challenges relating to wireless sensor/actuator mobile computing in a smart home. In designing and implementing the architecture of the technologies for such a system, it is therefore important to incorporate the sensors and actuator devices for computational issues which support the occupants' interaction with the physical environment. In the course of this study, a general context-awareness model is investigated, as well as technological solutions that have the capability of sensing, monitoring and controlling automation in the smart home.

The remaining sections of the paper are structured as follows. Section II describes related studies with respect to wireless sensor/actuator and mobile computing in the smart home. Section III presents the system architecture for the proposed model, comprising Wi-Fly adapter, sensors, actuators, microcontroller unit, end user and database server. Section IV outlines the infrastructure for wireless adapter network implementation, Section V defines techniques to set up application control for both the smartphone and the embedded system, and Section VI offers the smartphone interfaces. Section VII considers and evaluates the results and finally, Section VIII provides the conclusion and suggestions for future work.

II. RELATED WORK

Wireless technologies in the smart home have been studied by many researchers. These studies have used different types of wireless technology, such as Wi-Fi, IEEE 802.15.4, Bluetooth, GSM, ZigBee and Konnex RF (KNX-RF), [9][10][11][12], and much progress is currently being made on this issue. One of the most popular examples is the use of wireless sensors and actuators [13], but different kinds of technology and home automation devices have been integrated with different wireless protocols for ease of use in control and data gathering.

Studies by Hoof et al. [14] show that trends in smart home technology are based on the technical potential of their function, but are seldom focused on the needs and requirements of residents. Therefore, new techniques must take into account the attitudes of occupants and earn their approval of intelligent home equipment, particularly in the context of provision for the aging and disabled population. In general, smart home technologies in the residential environment require five things to make them 'smart', which are wireless technology, home automation, remote control, intelligent control and a database.

The most important features of this technology are the actuators and sensors. An actuator is used to create an action according to the data collected by sensors, which are used to decide the action. Currently, developments in technologies permit inhabitants to control home appliances without difficulty by operating devices such as a remote control, smartphone or PC through a mere click. The smart home prototype is typically based in a house, with a convenient

wireless communication system, security and energy saving power source to reduce the input of the user.

Smart home automation is an important subject; it refers to the automatic control of equipment and devices in the smart home environment. Many studies have proposed designs for the system and even smart home products [15]. In the system of intelligent home equipment and appliances, human control is generally easy. A home automation system typically includes an internal network, and practical and intelligent rules and equipment for specific use in the home network. Systems and equipment can be controlled automatically and used to provide information on the home automation environment. Another facility is that changing the state or mode of a device can change or trigger other devices in the smart home environment [16].

An article by Alheraish describes the design of a remote control model entitled Mobile to Machine (M2M), based on the GSM network. The system consists of an embedded system, sensor, relay and a mobile to control home brightness, monitoring and security issues. Some advantages were observed during the experimental stage of this model; for example, the system has the ability to use different techniques such as SMS and GPRS, and it is appropriate for simple applications using ADC or DAC. On the other hand, it also has some drawbacks for which the designer needs to find suitable solutions. It will be difficult to make the model M2M viable before the solutions to these problem are found [17].

Another study by Acker and Massoth describes a system design solution for monitoring and controlling smart home technologies, and also for secure access, via smartphone. The model system involves a variety of techniques and equipment to control home devices, such as a BlackBerry smartphone, a sensors box (Allnet ALL 4000) and BlackBerry Enterprise Solution. The appliances connected included TV and Hi-Fi, and Java Development Environment (JDE) software was used. The results which they report include three different settings of access time and show respectable results in terms of usability. However, the drawback was that the BlackBerry smartphone software would need to be regularly updated from BES in order to maintain a high level of usability [18].

Das et al.'s [19] study makes specific reference to the circumstances of intelligent home security using smartphones. They were able to remotely contact and control electronic security devices and integrate other services. This study managed home automation and security (HASEC) through the design and implementation of smartphone devices' processing systems, demonstrating that mobile technology can be used to provide the necessary security for products as well as operations related to control. The HASEC system permits different types of mobile phone to access the full Web browser.

Within the smart home setting, one very important development would be to make context-aware applications available which could assist older people, by making them more comfortable and making the activities of daily life easier. To succeed in this goal, the current study seeks to

incorporate various advances in technology into the smart home prototype.

III. SYSTEM DESIGN AND IMPLEMENTATION

Figure 1 illustrates the comprehensive system architecture from the viewpoint of communication, electronic devices and data management. These aspects are described in the following sections.

A. Sensors Design

There are different types of sensors and actuators involved in the prototype system. They can also be used to obtain a context using systems embedded in the smart home environment. These sensors include an obstacle avoidance sensor; smoke sensor; human detection sensor; brightness sensor; humidity sensor; ultrasonic distance sensor; and sound sensors (see Figure 2 for the different types of sensors used in the system design). The main aim of this function is to transfigure real physical variables into digital variables, which are then processed in the computing system. The acquisition of this type of context will need features such as a wireless adapter, sensors, an embedded system and the end user client. More detail can be found in [20].

B. Wi-Fly (RN-370)

The wireless adapter is powered by an external AC to 5V DC and two AAA batteries, which can run up to eight hours on full charge. The device is connected using only a RS-232 serial port interface of the DB9. It receives information from sensors over the MCU, and then sends this information to end devices, such as an iPhone or PC. The data is also transferred over a reliable TCP/IP socket using an ad-hoc or infrastructure network.

C. Microcontroller Unit (Logic Converter):

The STC89Cxx series MCU is an important part of the system design, as it is an 8-bit single-chip microcontroller. This chip is compatible with outlying device communication that monitors resident activity and motion by gathering data from sensors, interfaced with the Wi-Fly wireless transceiver and embedded with 64K bytes flash memory to store data, which is shared with In-system Programming (ISP). It also has an In-System Application (ISA) to assist the operator.

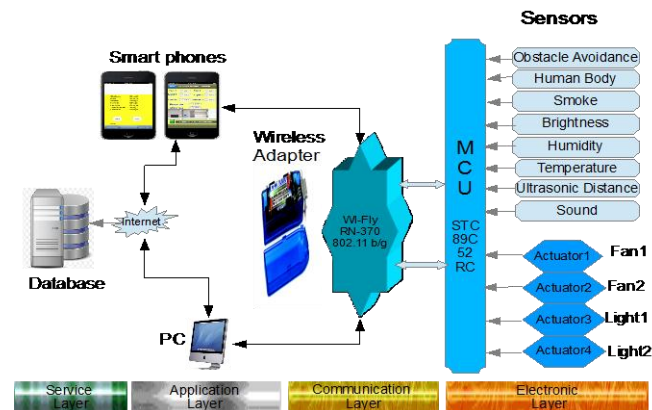


Figure 1. Framework layers in the prototype system.

These communication requirements are based on a microcontroller which needs to be adapted with a Universal Asynchronous Receiver Transmitter (UART) [18]. This is an appliance that transmits and receives raw information from the choice of MCU to support low level programming language C.

D. Relay (Actuators) Function

An actuator (SRD-05VDC-SL-C) is an electrical advice used to control electronic appliances by opening or closing the switches of other circuits. It has an electromagnetic coil to operate the switch in two positions, ON or OFF. It is activated by signal or electric current that flows through the coil, as a result of which the magnetic field attracts an armature. Once the current to the loop is interrupted, the armature is about twice as strong as the magnetic force to return it to its relaxed situation. This actuator is designed to be operated with DC voltage, and the diode on the coil is often installed to disable the magnetic field from disappearing. Most actuators are designed to trigger as quickly as possible to achieve the actions. Figure 3 shows the relay device used in this research and its connection with other electrical and electronic devices in the lab.



Figure 2. Smart home sensors design.




Figure 3. The MCU, Wi-Fly RN-370 and four actuators connected to LEDs and fans.

IV. WIRELESS CONNECTION OF THE ACTUATORS AND SENSORS

The prototype system in the smart home area is described by the diagram in Figure 4. The connection and communications are variable to enable low cost electronics to be employed for the design and implementation of the system. The Wi-Fi network is employed to make the connection between the electronic devices, as it is widely available in most home areas. The sensors and actuators are connected by wireless infrastructure to collect the data from the sensors and perform the control actions. The wireless adapter module RN-370 supports ad-hoc and infrastructure networking modes. The user interacts with the system using her/his smartphone. The requirements for the wireless prototype are given in Table I. The computing of the context data is done by the computer service and the decision-making algorithms.

A. Ad-hoc Configuration

To configure the ad-hoc mode, the Wi-Fly adapter needs to be set up for the ad-hoc network within the hardware or software instructions. To permit the ad-hoc network to establish communication with a smartphone through hardware, switch 1 must be turned ON. There are different kinds of connection available to create Wi-Fi applications, for example, remote environment sensors; control and monitoring of appliances; diagnostics; and mobile phone connections, such as GPS. Therefore, Wi-Fly configurations can be set as peer-to-peer (ad-hoc) networking to an iPhone, by using the IP address (169.254.1.1) and subnet mask (255.255.0.0), with the need to set the Service Set Identifier (SSID) and Wi Fly-GSX 21 in a wireless mobile computing network [21].

B. Infrastructure Configuration

To configure the infrastructure mode using software to start communication and associate the network, it is necessary to set up two commands for the wireless adapter, (1) to start an automation connection with a server and (2) to establish automated tracking of a connection from the remote host [21]. These setups need the Chortle series access point (1100) to connect the infrastructure network with the database server. When wireless communication for the local area network is set up using a Wi-Fly adapter, the device will automatically connect to the network. For context update and request to the network, the following commands are used:

- (\$\$\$); type 3 dollar signs to see message with CMD.
- Scan; scan available wireless network.
- Set wlan ssid chortle; ask to join the network.
- Set ip proto; turn HTTP mode=0x10 and TCP mode=0x02.
- Set ip host 10.0.0.11; web server DSN
- Set ip remote 80; set web server port.
- Set sys autoconn 1; automatically connect.
- Set com remote GET\$/work2/allensors.php? Sample server application.
- Set uart mode 2; automatically trigger mode.
- Save and reboot; save configuration and restart.

TABLE I. ELECTRONIC APPLIANCE ENGAGEMENT

Appliance Employed	Appliance Models	Functionality
Smartphone	iPhone 4	Receive the information gathered from sensors and control the actuator
Actuators	Relays	Switch on or off
Sensors	Temperature, Sound, Brightness, Smoke, Humidity, Distance, Human Body and Obstacle Avoidance	Gathering data from home environment
Wireless adapter	Wi-Fly RN-370	Transfers information written or read to the successive interface
Embedded system	MCU-STC89C52RC	Collect data from sensors and pre-process raw data
Wireless access point	Chortle (Cisco Aironet 1100).	Wireless network.
Home automation	Fan and Light	Brightness and cooling




Figure 4. Wireless communications over ad-hoc and infrastructure networks.

V. APPLICATION CONTROL FOR BOTH MICROCONTROLLER UNIT AND IPHONE

A. Function to Control Microcontroller Unit

The device can be used to control actions by, for example, sending a message via the mobile or keyboard including specific characters such as letters and numbers. The software program should decode the character message to be compatible with other details regarding the electronic appliances to be controlled. The action control can be accomplished through the output of the microcontroller unit after running the software program and passing the compiler through the test program in Debugger (Keil uVision3) for the debugger instruction. At the same time, the program will also

run in STC-ISP.exe, after selecting the hex file and the port from the PC to be ready for sending the character or action over wired and wireless technology. The following function may be used to control home devices:

```
If (OnOff=='0') RELAYL1 = 0x00;
If (OnOff=='a') RELAYL1 = 0xFF;
```

B. Function to Control iPhone

The iPhone 4 SDK (version IOS 6.1.3), and Apple iMac version 10.7.5 run with Xcode version 4.6.2, were used in this experiment. Objective-C and Cocoa Touch were the programs most implemented in this part of the project, with Cocoa Touch used by Objective-C offering the structure for iPhone application. One of the most important applications for the iPhone is the connection of the CFSocket function to create a TCP socket, as well as to set up the communication over the ad-hoc network. For a low level structure, the first step is to build the socket context, the second is to create the service socket as a TCP and set a callback, and the third is to set a local address which is used to connect the socket to the server.

The final step, when the interface for remote home control is linked up, is to offer options for action. For example, the user may decide whether the residence needs to switch on or off electronic devices, such as lights, fan, heater or speakers, or all the appliances together. For the first step of remote control, the code used involves simple 'IF' and 'ELSE' statements to regulate this. When the IF statement is defined, UIwatViewController 'on' status is set to character '0', and when the ELSE statement is defined, UIwatViewController 'off' status is set to character 'a'. The same function can be used for the remaining devices, e.g.:

```
- (IBAction) switchLight1 ValueChanged:(id)sender
{
    NSString *typecode;
    If (watViewController.switchLight1.on)
        typecode = @"0";
    else
        typecode = @"a";
    [alert release];
}
```

VI. MOBILE APPLICATION INTERFACES

A mobile application in a smart home system is a critical interface between the consumer and the computer system. It implements residents' control commands and gathers context-aware information using HTTP protocol, as well as SPLIT and POST functions.

Three of the selected prototype systems (iOS) for the apparatus are as follows. In Figure 5, which is a sample data interface, (a) is a login screen interface, (b) is a context monitor screen interface and (c) displays the status of the device and the control interface. The system design uses parts of the embedded system, sensors/actuators, and a smartphone located in the smart home environment to control home electrical devices using home automation within a context-aware computation service prototype. This

section describes the smartphone (iOS) interface used to offer the implementation of context-aware services using the functions in the model system.

The sensing data is first collected by the MCU, and then sent to the smartphone via the Wi-Fly transceiver, as shown in Figure 6. The figure shows the last observations of the experiment exploring the gathering of information from sensors using the Wi-Fly serial adapter and iPhone.

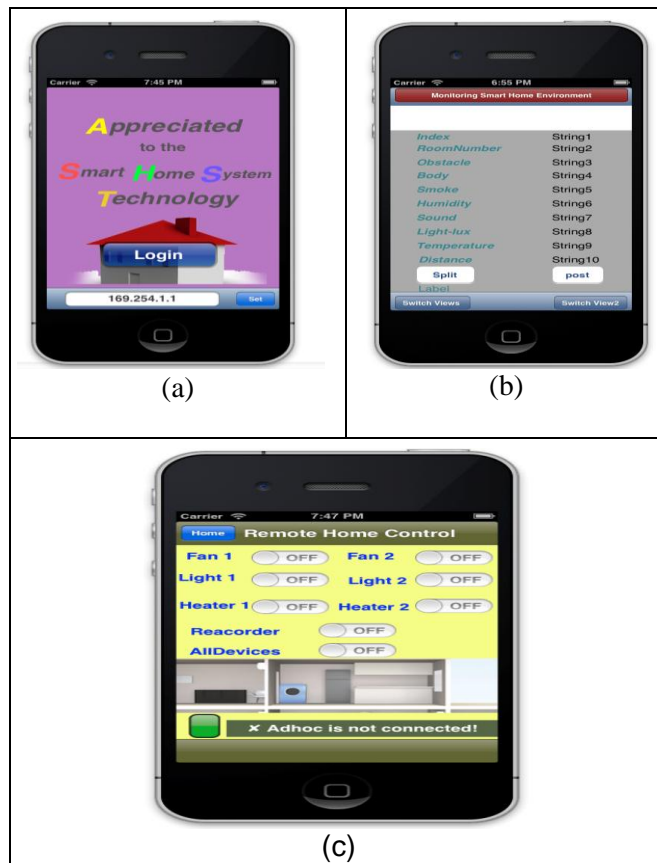


Figure 5. Smartphone (iOS) interface for smart home application.



Figure 6. Values collected from seven sensors by iPhone.

VII. TESTING AND EVALUATION

A. Monitoring Environment

In order to investigate the performance of the monitoring and control model within a real time environment situation, the iPhone was used to examine and test access time over various distances. The experiment concentrated on the time period required to connect the Wi-Fly access with the smartphone to remotely gather data from the lab environment and to remotely control the prototype automation system. In this test, the iPhone mobile used was linked to the CFSocket network via ad-hoc wireless communication.

The access period was measured five times, and each time the test was initiated when the smartphone holder moved into another location. The distance of measurement ranged from 1 metre to 25 metres, and the Wi-Fly function on the iPhone was updated to start the collection of information from the sensors. Once the prototype sensors were connected with the TCP/IP socket, an indication of whether the network was connected or not was given by receiving the notification of “Connection is successful, please login” or “Connection is not successful, please try again”. The remote control data gathering would begin when a successful connection was established.

The results in Figure 7 show a statistical examination of the system design used to monitor the smart home environment and the outcome of requests regarding security events. The experiment was repeated five times for each distance. The performance of all service providers already in Wi-Fly networks were tested and showed good stability in terms of the implementation of a wireless link. Comparisons can be made between the performance of this prototype system, using the results for time response in monitoring the environment and controlling home facilities, with other results for the same situation but using different technologies and methods [14].

Figure 8 shows an evaluation of the results of home environment monitoring using eight sensors and different locations, times and IDs. The average time response using the Wi-Fly 802.11 b/g serial adapter and iPhone smartphone was 5.406 u/s at a distance of 0cm, which increased to 5.838 u/s at a distance of 25m. The signal strength then began to decrease at a distance of 27m.




Figure 7. Time taken to monitor environment by touch screen.




Figure 8. Average time taken to monitor environment by touch screen.

B. Device Control

For the home automation function, on the other hand, the performance of the prototype system was able to reach a distance of up to 30m. All the devices could be controlled from 5m to 30m, but after that the system only gained a response from some of the devices, i.e. LED1 and LED2.

Table II. TEST RESULTS FOR HOME AUTOMATION

Distance	5m	10m	15m	20m	25m	30m
All Devices	✓	✓	✓	✓	✓	x
Fan1	✓	✓	✓	✓	✓	x
Fan2	✓	✓	✓	✓	✓	x
LED1	✓	✓	✓	✓	✓	✓
LED2	✓	✓	✓	✓	✓	✓

(Note: ✓ represents acceptable action, x not acceptable action)

Table II illustrates the control of four electronic appliances, comprising two fans and two LED lights, using a mobile phone, the Apple iOS 6. It was observed that most of the devices responded to the smartphone function from distances of between 5m and 25m, as shown by the ✓ symbol. However, when the distance increased to 30m, some devices were unable to respond to the action command, as shown by the symbol x. The test focused mainly on functionality and distance.

VIII. CONCLUSION AND FUTURE WORK

A. Conclusion

This study presents an important part of a prototype system for remote control monitoring of smart home facilities using smartphones and low price commercial sensor equipment. The system architecture presented is a wireless system that may be used to integrate the smartphone of the user with the end devices in the smart home. It is also hypothesised that this system could be used to enhance the comfort and security of an elderly or disabled person living alone. In addition, the different features of the prototype, such as functions elaborated by smart home technologies, have been described, as listed below.

- 1) A model for the incorporation of sensors, actuators, wireless adapter (WI-Fly RN 370) and computer

services for remote control in the smart home environment has been theoretically and practically demonstrated.

- 2) Home automation has been part of the research, and therefore a number of appliances have been involved in the prototype, such as an actuator to control fans and LED lights using a smartphone.
- 3) The use of ad-hoc and wireless network infrastructures in this context has been explained.
- 4) The uses of smartphone applications to achieve smart home automation between the end user and computer service have been described.
- 5) The mobile programming was run with the iPhone SDK, using the TCP/IP socket to communicate between Wi-Fly and iPhone in order to collect information from sensors and to control home facilities; the effectiveness of this has been tested over a range of distances.

This paper has focused on the performance of a prototype system combining context-aware services with the remote control of home appliances. The results show that both monitoring and control of home facilities can be achieved very successfully using a smartphone; information was gathered from all sensors with a very short time lapse of between 5.406 u/s over a distance of 0cm and 5.838 u/s over 25m. Moreover, the smartphone was able to remotely control all the devices over a distance of 5m to 20m, though the fans failed to respond to the action command over 25m. Nevertheless, there are still some difficulties which require further analysis and discussion.

B. Future Work

Monitoring and automation control of the smart home is an important research application, and may be especially valuable for elderly people who live alone. The target is clearly to provide a good service and environment by employing advanced technology and methods.

The goals for future work must be case-based reasoning, taking into consideration the needs and requirements of the elderly or disabled end user, and decision-making rules regarding the information required from various sensors and how this is used to control electrical devices.

The precise nature of the set of rules, and the influencing factors within various sets of context-aware information, may be calculated and determined using methods such as a Bayesian network. Another future piece of work should be to use computing algorithms to investigate the quality of the context which has been created using information from various sensors.

REFERENCES

- [1] A. Sashima, Y. Inoue, T. Ikeda, T. Yamashita, and K. Kurumatani, "CONSORTS-S: A mobile sensing platform for context-aware services," in *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, International Conference on, 2008, pp. 417-422.
- [2] J. Nichols, et al., "Generating remote control interfaces for complex appliances," in *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology*, 2002, pp. 161-170.
- [3] F. Xia, Y.-C. Tian, Y. Li, and Y. Sung, "Wireless sensor/actuator network design for mobile control applications," *Sensors*, vol. 7, 2007, pp. 2157-2173.
- [4] S. Yan, C. Qiang, and L. Hong, "Smart home based on wireless sensor-actuator networks," *China Communications*, vol. 8, 2011, pp. 102-109.
- [5] W. Lumpkins, "The internet of things meets cloud computing [Standards Corner]," *Consumer Electronics Magazine, IEEE*, vol. 2, 2013, pp. 47-51.
- [6] L. Trappeniers, M. A. Feki, F. Kawsar, and M. Boussard, "The internet of things: the next technological revolution," *Computer*, vol. 46, 2013, pp. 0024-25.
- [7] P. Bellavista, A. Corradi, M. Fanelli, and L. Foschini, "A survey of context data distribution for mobile ubiquitous systems," *ACM Computing Surveys (CSUR)*, vol. 44, 2012, p. 24.
- [8] D. Ding, R. A. Cooper, P. F. Pasquina, and L. Fici-Pasquina, "Sensor technology for smart homes," *Maturitas*, vol. 69, 2011, pp. 131-136.
- [9] N. Langhammer and R. Kays, "Evaluation of wireless smart metering technologies in realistic indoor channels," *Wireless Communication Systems (ISWCS)*, 2011 8th International Symposium on, 2011, pp. 387-391.
- [10] T. Han, B. Han, L. Zhang, X. Zhang, and D. Yang, "Coexistence study for WiFi and ZigBee under smart home scenarios," *Proc. 3rd IEEE International Conference on Network Infrastructure and Digital Content*, 2012, pp. 669-674.
- [11] S. O. Al Mehairi, H. Barada, and M. A. Qutayri, "Integration of technologies for smart home application," ed, 2007, pp. 241-246.
- [12] Y. Tajika, T. Saito, K. Teramoto, N. Oosaka, and M. Isshiki, "Networked home appliance system using Bluetooth technology integrating appliance control/monitoring with Internet service," *Consumer Electronics, IEEE Transactions on*, vol. 49, 2003, pp. 1043-1048.
- [13] J. Zhang, G. Song, H. Wang, and T. Meng, "Design of a wireless sensor network based monitoring system for home automation," ed, 2011, pp. 57-60.
- [14] J. Van Hoof, H. Kort, P. Rutten, and M. Duijnste, "Ageing-in-place with the use of ambient intelligence technology: Perspectives of older users," *International Journal of Medical Informatics*, vol. 80, 2011, pp. 310-331.
- [15] C. Kühnel, T. Westermann, F. Hemmert, S. Kratz, A. Müller, and S. Möller, "I'm home: Defining and evaluating a gesture set for smart-home control," *International Journal of Human-Computer Studies*, vol. 69, 2011, pp. 693-704.
- [16] Y.-W. Kao and S.-M. Yuan, "User-configurable semantic home automation," *Computer Standards & Interfaces*, vol. 34, 2012, pp. 171-188.
- [17] A. Alheraish, "Design and implementation of home automation system," *Consumer Electronics, IEEE Transactions on*, vol. 50, 2004, pp. 1087-1092.
- [18] R. Acker and M. Massoth, "Secure ubiquitous house and facility control solution," in *Internet and Web Applications and Services (ICIW)*, 2010 Fifth International Conference on, 2010, pp. 262-267.
- [19] S. R. Das, S. Chita, N. Peterson, B. Shirazi, and M. Bhadkamkar, "Home automation and security for mobile devices," in *Pervasive Computing and Communications Workshops (PERCOM Workshops)*, 2011 IEEE International Conference on, 2011, pp. 141-146.
- [20] O. Ghabar and J. Lu, "The Designing and Implementation of a Smart Home System with Wireless Sensor/Actuator and Smartphone," *INFOCOMP 2014, The Fourth International*

Conference on Advanced Communications and Computation,
2014, pp. 56-64.

- [21] R. Networks. (2010, WiFly Serial Adapter RN-370 & RN-374. Install Guide and User Manual 1.00, pp.1-21. Available: www.rovingnetworks.com

Adaptive Noise Reduction in Ultrasonic Images

Somkait Udomhunsakul

Engineering and Architecture Faculty
Telecommunication and Electronics Engineering Department
Rajamangala University of Technology Suvarnabhumi
Nonthaburi, Thailand 11000
Email: sudomhun@hotmail.com

Abstract—Ultrasonic image is an imaging technique that is commonly used for medical diagnostics. Unfortunately, the quality of ultrasonic images is limited, mainly due to speckle noise. Speckle noise reduction is one of the most important processes in enhancing the quality of ultrasonic images. In this paper, an adaptive noise reduction technique in wavelet domains for ultrasonic images is studied. First, a logarithmic transformation is performed on an original ultrasonic image in order to convert a multiplicative noise to an additive one. Next, Stationary Wavelet Transform is used to decompose the image resulting from the first step into four subbands. Then, an adaptive Wiener filter is applied to all detailed subbands in order to suppress additive noises in these subbands. Subsequently, the reconstructed image is derived by performing an inverse Stationary Wavelet Transform on those resulting subbands and following by an exponential transformation. The performance of the studied algorithm is evaluated objectively and subjectively on several ultrasonic images and it is compared against several well-known methods, such as Median filter, Wiener filter, Discrete Wavelet Transform based on soft thresholding, and Discrete Wavelet Transform along with Wiener filter. The results clearly demonstrate the superior performance of the studied method in terms of signal to mse ratio (S/mse), edge preservation (β) values as well as perceptible image quality.

Keywords—Stationary Wavelet Transform; Multiplicative Noise Reduction; Wiener Filter; Ultrasonic Images.

I. INTRODUCTION

Ultrasound imaging is predominant and plays an important role in medical diagnosis because it is a noninvasive, nonradioactive, real-time and inexpensive modality [1]. However, ultrasonic images usually suffer from three component kinds of noises. The first arises from the electronics of the detection system. For instance, the signal intensity of the backscattered ultrasound signals is affected by the operating frequency of the transducer: the higher the frequency, the greater the tissue attenuation, which therefore produces a lower signal-to-noise ratio (SNR). The second source, speckle, corresponds to coherent wave interference in tissue. It is well known to be signal-dependent in ultrasound imaging systems. The final term, clutter, is applied to signals arising from side lobes, multipath reverberation, and tissue motion that add noise to the ultrasound images [1].

Over the years, speckle noise suppression has been widely studied and considered. When filtering random noise from an image, there are two main issues to be considered: how much noise has been removed, and how well edges are preserved without blurring. Traditionally, there are several simple techniques for noise suppression, such as a moving average filter

and Gaussian filter. Being merely low-pass filters, they can effectively suppress noise, but they fail to preserve many useful details [2]. For speckle noise reduction techniques, some of the well-known filters include Lee filter, Kuan filter, median filter, and homomorphic Wiener filters [3]–[5]. These filters can effectively suppress speckle noise, but they fail to sufficiently preserve the edges. In the past decade, there has been considerable interest in using Wavelet transform as a powerful tool for recovering signal from noisy data. This method is generally referred to as a wavelet shrinkage technique. In 1995, D. L. Donoho presented a soft threshold method for denoising in one dimensional signal [6]. S. Chang, B. Yu and M. Vetterli introduced an adaptive wavelet threshold for image denoising and compression. They proposed a new shrinkage method, BaeyShrink [7], which also outperformed Donoho and Johnstone's Sureshrink [8]. Furthermore, other authors proposed probabilistic methods for speckle noise reduction in the wavelet domain [9]–[12]. Recently, A. K. Gupta and D. Sain have proposed a speckle reduction technique using a logarithmic threshold contourlet [13]. The method proposed by C. Barcelos and L. Vieira used an adaptive edge-controlled variation function to detect and reduce speckle noise [14]. Another proposed approach uses adaptive block-based singular value decomposition for speckle noise suppression [15].

In this paper, an adaptive noise suppression technique for ultrasonic images is proposed. The studied method is a preprocessing step for speckle noise reduction, before applying a feature extraction process [16]. First, a logarithmic transformation is applied to an original image in order to convert the multiplicative noise into additive noise. Next, Stationary Wavelet Transform (SWT) is used to decompose the transformed image resulted from the first step into four subbands. SWT is a wavelet transform algorithm that do not decimated but instead padding the filters with zeros [17]. That is, all the subband images would have the same size as the original images. Therefore, it has several advantages, as compared to Discrete Wavelet Transform (DWT). First, the transformation is translation-invariance. In addition, there is no information loss in each subband, since there is no downsampling process, unlike DWT. Then, an adaptive Wiener filter is applied to all the detailed subbands. An adaptive Wiener filter is a well-known filtering technique that has been applied not only to reduce stationary noise in noisy images but also to suppress blocking artifacts [18]. As a result, a reconstructed images derived from the studied method would be smoother, as compared to other filtering techniques, e.g., a block-based SVD based approach [15]. Finally, an inverse

SWT is computed and applied to the exponential transformation to reconstruct the denoised image. Then, in order to evaluate the performance of the studied method, the quality of reconstructed images derived from the studied method is compared against other existing approaches, such as median filter, Wiener filter, Discrete Wavelet Transform (DWT), based on soft thresholding, and DWT coupled with Wiener filter.

The rest of this paper is organized as follows. In Section II, a studied method for speckle noise reduction is described in details. Then, the quantitative image quality measurements and experimental results are given in Section III and IV respectively. Finally, the conclusion remarks are provided in Section V.

II. STUDIED METHOD

Similar to homomorphic Wiener filtering, the studied method could be used to reduce a speckle noise in medical images, which is done in the SWT domain. The block diagram of studied method is illustrated in Figure 1. Details are as follows:

- Take a logarithmic transformation to the original image (f), which yields image result (g).
- Perform a 2-D SWT on the log transformed image in order to decompose the transformed image g into four subbands (LL, LH, HL and HH).
- Perform a 2-D adaptive Wiener filter only in the detailed subbands (LH, HL and HH), window size of 7x7 is chosen which yields the image result (\hat{Y}).
- Apply the inverse 2-D SWT which yields a denoised image (\hat{g}).
- Take the exponential transformation of the denoised image to get the reconstructed image (\hat{f}).

The two main components of this method: 2-D stationary wavelet transform and adaptive wiener filter, are described below.

A. 2-D stationary wavelet transform

Unlike the conventional Discrete Wavelet Transformer (DWT), the two dimensional Stationary Wavelet Transformer (2-D SWT) is based on the idea of no decimation, which means the SWT is translation-invariant [19]. It applies the DWT and omits both down-sampling in the forward and up-sampling in the inverse transformation. 2-D SWT can be implemented by first applying the DWT along the rows of an image, and then applying it on the column of an image. Therefore, a transformed image is decomposed into four subbands, which are the same size as the original image. The LL band contains the approximation coefficients, the LH band contains the horizontal details, the HL band contains the vertical details and the HH band contains the diagonal details. Without translation-invariance, slight shifts in the input signal will produce variations in the wavelet coefficients that might introduce artifacts into the noise reduction process. This property is good for noise removal because the noise is usually spread over a small number of neighbouring pixels. The 2-D SWT decomposition scheme is illustrated in Figure 2.

B. Adaptive wiener filter

Two dimensional Wiener filter is a minimum mean-square error filter [20]. It is a nonlinear spatial filtering that moves a window or kernel over each pixel in the image, computes and replaces the central pixel values under the window. It uses a collection of window sizes to estimate the noise power from the local image mean (μ) and standard deviation (σ). The output of 2-D Wiener filter [21] [22] is defined by:

$$\hat{Y}(x_i, y_j) = \mu + \frac{\sigma^2 - v^2}{\sigma^2} \times (Y(x_i, y_j) - \mu) \quad (1)$$

where μ and σ^2 represents the local mean and standard deviation obtained from the noisy image window respectively. Y is the noisy pixel and \hat{Y} is the filtered pixel. Also, v is the noise variance, estimated from the average of all the local estimated variances in the image. Note that, only an odd number should be used as the size of the kernel. If the size is too large, important feature will be lost. On the other hand, if the size is too small, noise reduction may not yield good results. In general, a kernel size of 3x3 and 7x7 provides good results [13].

III. QUANTITATIVE QUALITY MEASURES

To quantify the achieved noise reduction ability performance, there are two main issues to be considered, which are how much noise has been removed, and how well edges are preserved without blurring. In the past decade, there have been many quantitative quality measurements proposed. In this research study, three image quality measurements: Mean Square Error (MSE), Signal to MSE ratio (S/mse), and edge preservation (β) are computed using original and reconstructed image data [23] [24].

A. Mean Square Error (MSE)

$$MSE = \frac{1}{mn} \sum_{i,j=1}^{m,n} (\hat{S}_{i,j} - S_{i,j})^2 \quad (2)$$

Where n and m are image dimension. \hat{S} and S are referred to reconstructed and original images, respectively. The higher MSE value denotes the lower image quality.

B. Signal to MSE ratio(S/mse)

To evaluate speckle noise reduction, a Signal to MSE ratio (S/mse) is used, instead of the standard signal to noise ration. It is defined as below:

$$\frac{S}{MSE} = 10 \log_{10} \frac{\frac{1}{mn} \sum_{i,j=1}^{m,n} (S_{i,j}^2)}{MSE} \quad (3)$$

The lower S/mse means the lower image quality.

C. Edge preservation (β)

To consider the performance of edge preservation, a parameter β based on a correlated operation between original and reconstructed images is used and given by:

$$\beta = \frac{\Gamma(\Delta s - \bar{\Delta s}, \Delta \hat{s} - \bar{\Delta \hat{s}})}{\sqrt{\Gamma(\Delta s - \bar{\Delta s}, \Delta s - \bar{\Delta s})\Gamma(\Delta \hat{s} - \bar{\Delta \hat{s}}, \Delta \hat{s} - \bar{\Delta \hat{s}})}} \quad (4)$$

where $\bar{\Delta s}$ and $\bar{\Delta \hat{s}}$ are the mean values in the region of interest(ROI) $s_{i,j}$ and $\hat{s}_{i,j}$ respectively. Also, Δs and $\Delta \hat{s}$ represent the high pass filtered operation of s and \hat{s} , respectively,




Figure 1. Block diagram for speckle noise reduction in SWT domain [16].




Figure 2. 2-D Stationary Wavelet Transform Decomposition Scheme.

obtained from a 3x3 pixel standard approximation of Laplacian operator where:

$$\Gamma(s_1, s_2) = \sum_{i,j=1}^{m,n} s_{1(i,j)} s_{2(i,j)} \quad (5)$$

Here, the larger value of β means the better feature preservation ability of the reconstructed image.

IV. EXPERIMENTS

In this part of the experiment, to validate the performance of the studied method, various cholecystitis ultrasonic images are used, as shown in Figure 3. The image size is 256x256. A number of experiments were conducted and compared with other traditional methods, which were 2-D median filter (7x7), 2-D adaptive Wiener filter (7x7), DWT with soft thresholding, and DWT along with Wiener filter. The experiments reported in this section have been tested using MATLAB 10.0 - R2010b (64 bit). All the wavelet-based techniques used Daubechies 4 wavelet basis, with one level of DWT and SWT decomposition. In fact, noise is generally spread over in detailed subbands, due to the components of highpass wavelet filters. Therefore, the

TABLE I. EXPERIMENTAL RESULTS OBTAINED BY VARIOUS NOISE REDUCTION TECHNIQUES AT NOISE VARIANCE OF 0.08.

Images	Methods	Mask	S/MSE	β
Cholecystitis (a)	Median filter	7x7	15.7360	0.1333
	2-D Wiener filter	7x7	16.1618	0.2575
	DWT with soft thresholding	-	16.4526	0.2457
	DWT and Wiener filter	7x7	16.6765	0.2999
	SWT and Wiener filter	7x7	17.2609	0.4224
Cholecystitis (b)	Median filter	7x7	12.3104	0.1451
	2-D Wiener filter	7x7	12.3741	0.2412
	DWT with soft thresholding	-	113.1566	0.2388
	DWT and Wiener filter	7x7	13.5087	0.2923
	SWT and Wiener filter	7x7	14.6532	0.4480
Cholecystitis (c)	Median filter	7x7	16.4115	0.1157
	2-D Wiener filter	7x7	16.8007	0.3557
	DWT with soft thresholding	-	17.0805	0.2670
	DWT and Wiener filter	7x7	17.3977	0.3237
	SWT and Wiener filter	7x7	17.8405	0.4297

2-D adaptive Wiener filter is applied only in detailed subbands. To quantify the achieved performance in terms of the ability of speckle noise reduction and edge preservation, three original cholecystitis ultrasonic images, as illustrated in Figure 3, are used. These images are first corrupted with noise at variance of 0.08 where the noisy images are depicted in Figure 4(a). Then, different speckle noise reduction techniques are applied on these noisy images. Then, the quality of reconstructed images are evaluated using S/mse and β . The numerical results, presented in Table I, show that the proposed approach outperforms other methods in terms of S/mse and β .

To visually compare with all other methods, the comparatives of various results are shown in Figure 4. As for the results, Figure 4(b) and Figure 4(c) are operated by a fixed window size 7x7 in spatial domain. The reconstructed images are overly smoothed and have artifacts around the object. On the other hand, the combination of Wiener filter and SWT outperforms DWT with soft thresholding and DWT along with Wiener filter, as shown in Figure 4(d), 4(e) and 4(f). It is seen that the studied method can efficiently reduce noise in the homogeneous area and simultaneously preserve the edges feature thereby resulting in a better reconstructed image in terms of visual perception.

Next, in order to evaluate performance of studied method more extensively, the realistic noisy liver and kidney ultrasonic




Figure 3. Cholecystitis ultrasound images.




Figure 4. Results of various speckle reduction methods: (a) noisy cholecystitis ultrasound image, (b) denoised image using 2-D median filter, (c) denoised image using 2-D Wiener filter, (d) denoised image using DWT with soft thresholding, (e) denoised image using DWT and Wiener filter, and (f) denoised image using SWT and Wiener filter.

images are also used. The original images and resulting images of different speckle noise reduction techniques are shown in Figure 5 and Figure 6 respectively. The results demonstrate that SWT along with Wiener filter outperforms other methods since it can effectively reduce speckle noises and simultaneously preserve edge features as well as important details of the ultrasonic images.

V. CONCLUSION AND FUTURE WORK

In this research, the main aim is to study and compare the different methods of speckle noise suppression in ultrasonic images. The studied method uses SWT to transform a logarithmic image and then applies an adaptive Wiener filter in each detailed subband. The advantage of multi-resolution analysis using SWT for speckle noise reduction is that it

can reduce noise while preserving the feature structure of the reconstructed image. From the preliminary results, the combination of the SWT and adaptive Wiener filter has better quantitative and qualitative performances, compared with existing methods. Future work, the use of different types of mother wavelets and other types of wavelet transform, such as Wavelet Packet Transform, will be investigated in order to get the best result. Moreover, subjective assessment has to be performed by ultrasonographers in order to visually ensure the reconstructed image quality. Consequently, the best technique could be developed and be implemented in hardware.

ACKNOWLEDGMENT

The author would like to thank Dr. Samuel P. Kozaitis, currently a professor in the Division of Electrical Engineering




Figure 5. Results of various speckle reduction methods: (a) noisy liver ultrasound image, (b) denoised image using DWT with soft thresholding, (c) denoised image using DWT and Wiener filter, and (d) denoised image using SWT and Wiener filter.




Figure 6. Results of various speckle reduction methods: (a) noisy kidney ultrasound image, (b) denoised image using DWT with soft thresholding, (c) denoised image using DWT and Wiener filter, and (d) denoised image using SWT and Wiener filter.

at Florida Institute of Technology, and would also like to offer special thanks to Dr. Napa Sae-Bae, presently a lecturer in the faculty of science and technology at Rajamangala University of Technology Suvarnabhumi, for their very helpful suggestions towards making this research complete.

REFERENCES

- [1] A. Webb and G. C. Kagadis, Introduction to biomedical imaging. Wiley Hoboken, 2003.
- [2] A. McAndrew, Introduction to digital image processing with MATLAB. Course Technology. ISBN-10 534400116, 2004.
- [3] J. S. Lee, "Digital image enhancement and noise filtering by use of local statistics," Pattern Analysis and Machine Intelligence, IEEE Transactions on, no. 2, 1980, pp. 165–168.
- [4] D. T. Kuan, A. A. Sawchuk, T. C. Strand, and P. Chavel, "Adaptive noise smoothing filter for images with signal-dependent noise," Pattern Analysis and Machine Intelligence, IEEE Transactions on, no. 2, 1985, pp. 165–177.
- [5] A. K. Jain, Fundamentals of digital image processing. Prentice-Hall Englewood Cliffs, 1989.
- [6] D. L. Donoho, "De-noising by soft-thresholding," Information Theory, IEEE Transactions on, vol. 41, no. 3, 1995, pp. 613–627.
- [7] S. G. Chang, B. Yu, and M. Vetterli, "Adaptive wavelet thresholding for image denoising and compression," Image Processing, IEEE Transactions on, vol. 9, no. 9, 2000, pp. 1532–1546.
- [8] D. L. Donoho and J. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," Biometrika, vol. 81, no. 3, 1994, pp. 425–455.
- [9] P. Kishore, A. Sastry, A. Kartheek, and S. H. Mahatha, "Block based thresholding in wavelet domain for denoising ultrasound medical images," in Signal Processing And Communication Engineering Systems (SPACES), 2015 International Conference on. IEEE, 2015, pp. 265–269.
- [10] J. Jin, Y. Liu, Q. Wang, and S. Yi, "Ultrasonic speckle reduction based on soft thresholding in quaternion wavelet domain," in Instrumentation and Measurement Technology Conference (I2MTC), 2012 IEEE International. IEEE, 2012, pp. 13–16.
- [11] C. Chen and N. Zhou, "A new wavelet hard threshold to process image with strong gaussian noise," in Advanced Computational Intelligence (ICACI), 2012 IEEE Fifth International Conference on. IEEE, 2012, pp. 558–561.
- [12] J. Scharcanski, C. R. Jung, and R. T. Clarke, "Adaptive image denoising using scale and space consistency," Image Processing, IEEE Transactions on, vol. 11, no. 9, 2002, pp. 1092–1101.
- [13] A. K. Gupta and D. Sain, "Speckle noise reduction using logarithmic threshold contourlet," in Green Computing, Communication and Conservation of Energy (ICGCE), 2013 International Conference on. IEEE, 2013, pp. 291–295.
- [14] C. A. Barcelos and L. E. Vieira, "Ultrasound speckle noise reduction via an adaptive edge-controlled variational method," in Systems, Man and Cybernetics (SMC), 2014 IEEE International Conference on. IEEE, 2014, pp. 145–151.
- [15] N. Sae-Bae and S. Udomhunsakul, "Despeckling algorithm on ultrasonic image using adaptive block-based singular value decomposition," in Photonics Asia 2007. International Society for Optics and Photonics, 2007, pp. 68 330J–68 330J.
- [16] S. Udomhunsakul and P. Wongsita, "Feature extraction in medical ultrasonic image," in 3rd Kuala Lumpur International Conference on Biomedical Engineering 2006, ser. IFMBE Proceedings, F. Ibrahim, N. Osman, J. Usman, and N. Kadri, Eds. Springer Berlin Heidelberg, 2007, vol. 15, pp. 267–270.
- [17] G. P. Nason and B. W. Silverman, "The stationary wavelet transform and some statistical applications," in Wavelets and statistics. Springer, 1995, pp. 281–299.
- [18] V. Nath and D. Hazarika, "Blocking artifacts suppression in wavelet transform domain using local wiener filtering," in Emerging Trends and Applications in Computer Science (NCETACS), 2012 3rd National Conference on, March 2012, pp. 93–97.
- [19] V. Raj and T. Venkateswarlu, "Denoising of medical images using undecimated wavelet transform," in Recent Advances in Intelligent Computational Systems (RAICS), 2011 IEEE, Sept 2011, pp. 483–488.
- [20] J. S. Lim, Two-dimensional signal and image processing. Englewood Cliffs, NJ, Prentice Hall, 1990.
- [21] D. Lai, N. Rao, C. H. Kuo, S. Bhatt, and V. Dogra, "An ultrasound image despeckling method using independent component analysis," in Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on. IEEE, 2009, pp. 658–661.
- [22] E. Ercelebi and S. Koc, "Lifting-based wavelet domain adaptive wiener filter for image enhancement," IEE Proceedings-Vision, Image and Signal Processing, vol. 153, no. 1, 2006, pp. 31–36.
- [23] A. Achim, A. Bezerianos, and P. Tsakalides, "Novel bayesian multiscale method for speckle removal in medical ultrasound images," Medical Imaging, IEEE Transactions on, vol. 20, no. 8, 2001, pp. 772–783.
- [24] F. Sattar, L. Floreby, G. Salomonsson, and B. Lovstrom, "Image enhancement based on a nonlinear multiscale method." IEEE transactions on image processing: a publication of the IEEE Signal Processing Society, vol. 6, no. 6, 1996, pp. 888–895.

Writer Identification Method using Inter and Intra-Stroke Information

Jungpil Shin, Zhaofeng Liu
 School of Computer Science and Engineering,
 The University of Aizu, Fukushima, Japan
 E-mail: jpsin@u-aizu.ac.jp

Abstract— Most research in the field of writer identification currently uses positional data to identify the writer. But there is no writer identification method that is suitable for practical use so far. So, we newly devise two algorithms to raise accuracy rate and find some parameters that have beneficial effects on writer identification. One algorithm is “Block Type Model” which is a method to analyze positional relation and shapes of Chinese characters. The other algorithm is “Hidden feature analysis” which is an algorithm that identifies the writer using multi-parameters with all the strokes. As a result, any eight Chinese characters were enough to achieve over 99.9% accuracy rate. Additionally, we discovered that there are personal attributes in the speed of each stroke. With the new algorithm, we raised accuracy rate and found out that parameters have beneficial effects on writer identification.

Keywords- *Writer Identification; Stroke Information*

I. INTRODUCTION

Hand-written characters contain unique characteristics of people. Recently, what is most desired in the field of handwriting analysis is a proposal of a method or algorithm with an accuracy which is higher than or equivalent to the fingerprint authentication system or the vein authentication system.

To identify the characteristics of the writer, many algorithms have been proposed, but their identification success rate is not at the level of practical use. The biggest reason is that the amount of information obtained from the position of dots and lines is not sufficient to identify the writer.

The methods to extract characteristics of the writer from the limited information can be divided into two main categories. One is known and described as an “Intra” type. This type verifies the shape of a line of characters or the position among radicals. The other is called as an “Inter” type, which verifies the shape of every part of characters (e.g., distance between each stroke). Using the two methods is expected to bring higher accuracy.

Edge-Based Directional Features, connected-component contours and other methods have been invented in previous works. Off-line means the samples were collected from papers or by some analog method. On-line means the samples were collected from tablets or by some digital method. Number of characters means how many characters were necessary for matching.

Many algorithms analyze the position information of lines or points, and extract individual features. However,

there are few methods that analyze the pen pressure, stroke speed and an inclination angle of a pen. Moreover, there are few databases which contain huge amount of location data. Therefore, we newly devised two algorithms. One is an “Inter” type algorithm which uses the relationship among block such as radicals. The other is an “Intra” type algorithm with position data, pen pressure, speed data, and an inclination angle of a pen, an angular difference between paper and pen. Then, we combine these two methods and do an experiment using the database that we have been creating for over five years. The database currently contains 110 kinds of Chinese characters, written by 48 writers of different nationalities, 10 times each, adding up to 52,800 characters in total. In this research, half of the characters were used to extract the characteristics of the writers, and the other half was used for an identification experiment. The results indicate that features extracted from eight Chinese characters were enough to identify the writer. In Section 2 we show the approaches for writer identification. In Section 3 we show the experiment and results. In Section 4 we show the discussions. In Section 5 we conclude the paper.

II. WRITER IDENTIFICATION APPROACHES

In this section, we present an algorithm for writer identification using the Block Type Model, Hidden Feature Analysis.

A. Initialization

To identify the writer who has written Chinese characters, we used a database that contains positional data, pen pressure, inclination angle of the pen, angular difference between paper and pen, and pen speed. When we collected character data, we had no idea how to standardize the size of characters. In this research, we have to standardize the size of characters for matching. First we computed the average size of all the characters in the database for Dynamic Time Warping (DTW) matching. Then, we used it as the standard size of characters. However, since there are vertically long and horizontally long types in Chinese characters, expansion and reduction cannot be simply performed in every direction by the same ratio. Therefore, we distinguished the characters of a vertically long or horizontally long type. We divided the Chinese characters into two or more blocks: a radical block and another block. In the case of a Chinese character with two or more radicals, we added the number of blocks as radical block.

TABLE I. RESULTS OF BLOCK TYPE MODEL (%)

No. of Char.	1	2	3	4	5	6	7	8	9	10
Block	62.94	80.02	88.52	91.95	94.83	96.41	97.48	98.15	98.62	98.96

We further classified them according to the direction of vectors between the two most distant blocks. If the vector pointed sideways, we posited it as a horizontally long type character. If the vector was vertical, we posited it as a vertically long type. In particular, we used the expressions below for making judgements.

TABLE II. HIDDEN FEATURE ANALYSIS WITHOUT WEIGHT

No. of char	Pressure	Theta	Phi
1	55.45	33.18	33.60
2	73.54	42.02	45.40
3	81.98	47.89	52.56
4	86.93	51.49	57.35
5	89.95	54.05	61.45
6	91.90	56.32	64.66
7	93.41	58.29	67.18
8	94.46	59.44	69.51
9	95.40	60.56	71.60
10	96.02	61.63	73.24

The relation between the lengthwise direction of the interior of characters can be an important factor for a vertically long type of character, and the relation between the crosswise direction of the interior of characters can be an important factor for a horizontally long type of character. So, we resized characters according to the type of character (horizontally long or vertically long). If the character was vertically long, we resized it with the ratio of the lengthwise direction of the standard. If the character was horizontally long, we resized it with the ratio of the crosswise direction of the standard; we did not change the vertical ratio of vertically long type characters or the horizontal ratio of horizontally long type characters in Figure 1.

B. Block Type Model

As previously noted, our research attempts to combine an inter-type algorithm, the Block Type Model, with an intra-type algorithm, Hidden Feature Analysis.

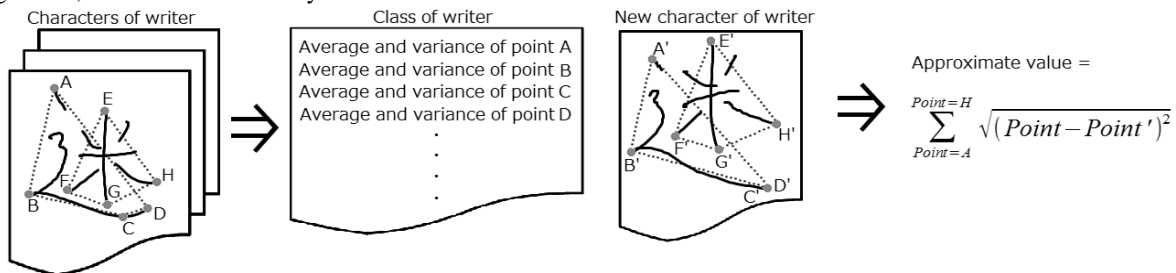


Figure 2. Graphic explanation about class




Figure 1. Graphic explanation about character types

1. A character is divided into some blocks.
2. The points that contain the maximum X-coordinate, the minimum X-coordinate, the maximum Y-coordinate, and the minimum Y-coordinate of each block are recorded.
3. The average of points of the same kind of character written by the same writers is computed. Then it is defined as a class of writer (e.g., Writer A's character B class)
4. The points of characters and points of classes are compared according to the Euclidean distance and weight (it should be defined as an approximate value). Then the class of the best approximation of the character is identified.
5. The class with the lowest approximate value is determined as the writer's class.

Weights are decided with the rule stated below.

Weight = Average of variance of each point / Variance of each point.

The approximate value is defined by the equation below number of characters.

$$A = E * W$$

where, A, E, and W is approximate value, Euclidean distance between each pair of points, and weight of each point, respectively.

In other words, higher variance means a lower approximate value. To put it another way, lower variance means a higher approximate value. A graphic explanation is given in Figure 2. In Figure 2, points A, B, C, D, and so on mean the points that have the maximum X-coordinate and Y-coordinate and minimum X-coordinate and Y-coordinate.

TABLE III. WEIGHTS OF HIDDEN FEATURE ANALYSIS.

	Average of accuracy of each parameter (%)	Weight
Pressure	85.91	1.18
Theta	52.49	0.72
Phi	59.67	0.82
Speed	93.68	1.28
Average	72.94	

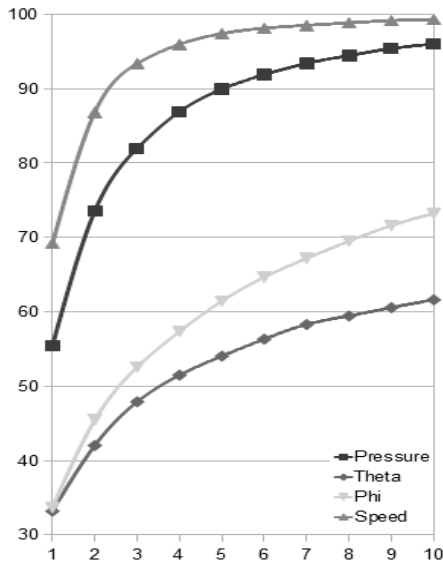


Figure 3. Results of Hidden feature analysis without weights (%).

We only show the results of the Block Type Model (Table 1), “No. of char.” means the number of Chinese characters used for matching.

This experiment was conducted using 110 kinds of Chinese characters written by 48 writers. One writer wrote 10 times each character. Namely, we collected 52800 characters in total. This experiment was conducted with 1 to 10 characters chosen randomly from the database, calculating the distance between the class and the character for each character. We used the sum total of the approximate values to identify the writer. Then, we repeated this 1000 times. We obtain the results of the average.

C. Hidden feature Analysis

This is an algorithm that identifies the writer using the average of pen pressure, the inclination angle of the pen, the angular difference between paper and pen, and the speed and range of these parameters in each stroke. The Hidden Feature Analysis was carried out as below.

1. Characters were divided into strokes.
2. The average of pen pressure, inclination angle of the pen, angular difference between paper and pen, and speed and range of these parameters in each stroke were

TABLE IV. RESULTS OF HIDDEN FEATURE ANALYSIS (%).

No. of char	1	2	3	4	5	6	7	8	9	10
4-Axis	87.79	95.1	97.19	98.25	98.85	99.24	99.45	99.63	99.73	99.82
4-Axis with weight	88.80	95.68	97.61	98.56	99.12	99.48	99.64	99.77	99.83	99.89

computed. Then, we defined the character as a class of writer (e.g., Writer A’s character B class)

3. We compare the written characters with class. (It was defined as an approximate value.)

Then, we identified the class of the best approximation to the character. (The class with the lowest approximate value was determined to be the writer’s class.)

Graphic explanation is given in Figure 1. The results of experiment with an individual parameter without its weight are found in Table 2.

In Figure 3, “pressure” and “speed” mean the pen pressure and speed of each stroke respectively. Theta is the inclination angle of the pen. “Phi” is the angular difference between paper and pen. The horizontal axis corresponds to the number of characters used for the analysis and the vertical axis corresponds to the correct rate.

The result shows that there is a difference in matching performance among the parameters. Then the average accuracy of each parameter is computed (Table 3).

We defined the average of the overall parameters as the standard and used the rule below to decide the weights.

$$\frac{\text{Average of accuracy of each parameter}}{\text{Average of accuracy of all parameter}}$$

The weight of each parameter is given in Table 3. In other words, we used the weight to improve the influence of the parameters that were performing well and also to lessen the influence of the parameters that were not performing so well.

The results of the combination of the Block Type Model and Hidden Feature Analysis are given later. Here we only show the results of the Hidden Feature Analysis (Table 5). “Hidden feature” means the accuracy rate without the weight. “Hidden feature with weight” means the accuracy with the weight.

This experiment was conducted with 1 to 10 characters which were chosen randomly from the database, calculating the distance between the class and the character for each character. We used the sum total of the approximate values to identify the writer. Then, we repeated this 1000 times. We obtained the results of the average.

III. MAIN EXPERIMENT AND RESULTS

We describe the results of two experiments: (1) an inter-type Block Type Model and (2) an intra-type Hidden Feature Analysis. Now, we combine these two methods and carry out the experiment. In particular, we summed up

the two methods' approximate values and appointed the class that had the lowest approximate value as the writer's class.

This experiment was conducted using 110 kinds of Chinese characters written by 48 writers. Each writer wrote each character 10 times. Namely, we collected 52800 characters in total.

Comprehensive results are given in Table 5. This experiment was conducted with 1 to 10 characters chosen randomly from the database, and the distance between the class and the character was calculated for each character. We used the sum total of the approximate values to identify the writer. Then, we repeated this 1000 times. We obtained the results of the average. Additionally, the ratio of approximate values of the two algorithms is set to 1:1.

There are some researches similar to our research [4] and [5]. Our database is smaller than that used in [5], but we can achieve a higher accuracy rate with a smaller number of characters. We can achieve 90% accuracy with one character, while [5] needed 50 words. Moreover, we can achieve a higher accuracy rate than [4] in the same situation. Their accuracy was 97.5%. But our maximum accuracy rate is 99.9%.

TABLE V. RESULTS OF COMBINATION OF BLOCK TYPE MODEL AND HIDDEN FEATURE ANALYSIS (%).

No. of char	1	2	3	4	5	6	7	8	9	10
Result of combination of Rhombus type model and 4-Axis analysis	90.53	96.71	98.52	99.24	99.58	99.81	99.86	99.91	99.94	99.98

IV. CONCLUSION

In this paper, we presented novel two algorithms to improve the accuracy rate and find some parameters that have beneficial effects on writer identification. The results indicate that four Chinese characters were enough to achieve an accuracy rate of 99% and eight Chinese characters were enough to achieve an accuracy rate of over 99.9%. Table 5 indicates that the accuracy rate of the two methods combined is higher than the results of the two individual approaches. This means that this method is quite useful for identifying the writer.

Our algorithm does not support one-to-many matching. This means that when we identify the writer, the written Chinese character must be the same kind as was written at the time of class generation. In the future, we hope to be able to generate a class that can support more Chinese characters with a small number of characters..

REFERENCES

[1] S. N. Srihari, S.-H. Cha, H. Arora, and S. Lee, "Individuality of Handwriting," *Journal of Forensic Sciences*, 47(4), pp. 1-17, July 2002.

[2] E. Zois, "Morphological Waveform Coding for Writer Identification," *Electronics Laboratory, Physics Department, University of Patras, Greece, March. No.33*, pp.385-398, 2000.

[3] C. Tomai, B. Zhang, and S. Srihari, "Discriminatory Power of Handwritten Words for Writer Recognition," *Intl. Conf. of Pattern Recognition (ICPR2004)*, Vol. 2, pp. 638-641, NY, 2004.

[4] Z. Long, W. Yunhong, and T. Tieniu, "Personal Handwriting Identification Based on PCA," *Proceedings of SPIE Second International Conference on Image and Graphics*, pp 766-771, 2002.

[5] B. Li, Z. Sun, and T. Tan, "Hierarchical Shape Primitive Features for Online Text-independent Writer Identification," *10th International Conference on Document Analysis and Recognition, ICDAR 2009*, pp. 986-990, Barcelona, Spain, 26-29 July 2009.

[6] Y. Yan, C. Chen, W. Deng and F. Yuan, "Chinese Handwriting Identification Based on Stable Spectral Feature of Texture Images," *International Journal of Intelligent Engineering and Systems*, Vol. 2, No. 1, pp.17-22, 2009.

[7] B. Zhang, S. Srihari, and S. Lee, "Individuality of Handwritten Characters," *Seventh Int'l Conf. Document Analysis and Recognition (ICDAR)*, pp. 1086, Aug. 2003.

[8] N. Mohammed, E. AlKassis, and D. Al-Muslih, "Off-Line Text-Independent Arabic Writer Identification using Contour-Based Features," *Mir@cl Lab, FSEGS, University of Sfax*, 2010.

[9] H. Said, T. Tan, and K Baker, "Personal identification based on handwriting," *Pattern Recognition*, Vol 33, No 1, pp. 149-160 January 2000.

[10] L. Schomaker and M. Bulacu, "Automatic Writer Identification Using Connected-Component Contours and Edge-Based Features of Uppercase Western Script," *IEEE, Transaction on pattern analysis and machine intelligence*, Vol. 26, No. 6 June. 2004.

[11] A. Bensefia, T. Paquet, and L. Heutte, "A writer identification and verification system," *Pattern Recognition Letters*, Vol 26, No. 13, pp. 2080-2092, October 2005.

[12] C. Siewkeng, T. Yonghaur, and V. Christian, "Online Text Independent Writer Identification Using Character Prototypes Distribution," *6th International Conference on Information, Communications & Signal Processing*, pp.1-5, 10-13, Dec. 2007, Utar, Malaysia.

[13] T. Matsuura, P.Thumwarin, and N. Kato, "On-live writer identification based on handwriting velocity and curvature of script," *23rd International Conf., on Image and Vision Computing*. pp. 192-197, 26-28 Nov. 2008, New Zealand..

[14] A. Schlappbach, M. Liwicki, and H. Bunke, "A Writer Identification System for On-line Whiteboard Data," *Pattern Recognition*, Vol 41, No 7, pp. 2381-2397, July 2008.

CHIC – Coupling Habitability, Interior and Crust

A new Code for Modeling the Thermal Evolution of Planets and Moons

Lena Noack, Attilio Rivoldini, Tim Van Hoolst

Department of Reference Systems and Planetology

Royal Observatory of Belgium (ROB)

Brussels, Belgium

email: lena.noack@oma.be, attilio.rivoldini@oma.be, tim.vanhoolst@oma.be

Abstract—We present a new numerical code CHIC (*Coupling Habitability, Interior and Crust*) for the simulation of the thermal evolution of terrestrial planets, with a focus on the numerical aspects of the code and its validation. The thermal evolution of the mantle is calculated either by solving the energy conservation equation supplemented by boundary-layer theory (1D parameterized thermal evolution model) or by solving the energy, mass, and momentum conservation equations (2D/3D convective thermal evolution). For the latter setting, the equations can be solved either in the Boussinesq or extended Boussinesq approximation. The code provides information on the temperature field, convective velocities and convective stresses in the mantle. Simulations can also be run in steady-state regime. The code provides a user updatable library of thermodynamic properties of iron and common mantle silicates as well as associated equations of state, which allow to compute material properties at high pressure and temperature. CHIC has been benchmarked with different convection codes, and compared to published interior-structure models and 1D parameterized models. The CHIC code handles surface volcanism, crustal development, and different regimes of surface mobilization like plate tectonics. It is therefore well suited for studying scenarios related to the habitability of terrestrial planets. CHIC is an advanced simulation code that can be applied to a diverse range of geodynamic problems and questions.

Keywords - *fluid dynamics; convection; numerical modeling; thermal evolution; planetology.*

I. INTRODUCTION

Numerical and parameterized models for convection and thermal evolution are essential tools to understand different geophysical processes in and between the interior of a planet and its atmosphere. These processes include for example the CO₂-cycle, the subduction cycle (delivering volatiles to the mantle), the release of volatiles by volcanic outgassing, the evolution of continents (stabilizing plate tectonics) and the possible maintenance of a magnetic field by strong cooling of the core. All those mechanisms may be important for the habitability of Earth, i.e., for its ability to host life, and may also play an important role for other planets. The magnetic field shields lighter volatiles in the atmosphere from erosion to space, whereas subduction of carbonates is an important phenomena that helps to regulate surface temperatures over geophysical timescales [1].

Several 2D and 3D convection codes have been developed over the past decades to investigate Earth-like planets. They typically concentrate solely on either the thermal evolution or do steady-state snapshots of the mantle and crust. Some models include the evolution of the core [2] or of the atmosphere [3][4]. Here, we describe a new code called CHIC that has been developed at the Royal Observatory of Belgium. The code is written in Fortran and is used to investigate different geophysical processes and feedback cycles on Earth-like planets. The planets are assumed to consist of several different spherical layers (shells). The lowermost shell represents the core and is overlain by a silicate shell (mantle and crust) and a potential water-ice layer. The uppermost shell represents the planet's atmosphere. All shells are thermally coupled, i.e., the heat flux and temperature are continuous at each interface between the different layers. The surface temperature is allowed to vary with time depending on the greenhouse gases in the atmosphere, or is taken constant if changes in the atmosphere are neglected.

The CHIC code is able to treat both 1D parameterized models (using the thermal boundary layer theory to determine the temperature evolution in a terrestrial or ocean planet) and 2D/3D models to investigate the detailed convection pattern in a silicate mantle or ice layer over time – both models have their advantages and disadvantages.

In a convection model (either modelling a 3D sphere or a 2D spherical annulus), lateral variations in the mantle can be investigated, including mantle plumes, local melt regions, and plate motions. A 1D model on the other hand assumes a laterally averaged profile for temperature and material properties like the density. As a result, simulations of, for example, the volcanic history of a terrestrial planet may differ between 1D and 2D/3D models.

1D thermal evolution models also have several advantages over 2D/3D convection models. The employed parameterization [5] is applicable to a large parameter space, including the simulation of both liquid and solid materials. Especially strongly convecting systems (e.g., liquid core or ocean) can be treated – which is generally unfeasible for planetary convection codes, as they will either produce numerical instabilities or require a high amount of computational power. 1D models, on the other hand, are very fast compared to convection models. Depending on the specific application, a 1D thermal evolution model runs in

the order of seconds or minutes, whereas 2D/3D models (that typically need a high resolution to avoid numerical errors) run for days or months.

CHIC couples a 1D parameterized module with a convection module in either 2D or 3D. The different modules can be applied as needed: for the core, either only changes in the core-mantle boundary (CMB) temperature are investigated, or a 1D model of the iron core including inner core freezing is applied (Figures 1(a) and 1(b)); the thermal state of the mantle and high-pressure ice layers are investigated either via a convection model or a 1D parameterized model; the atmosphere and a potential water ocean (Figures 1(c) and 1(d)) are investigated with a 1D module. CHIC is therefore a powerful tool for the investigation of the evolution of terrestrial or ocean planets - from interior to atmosphere - and their possible habitability.




Figure 1. Possible configurations that are investigated with CHIC: a) mantle with variable CMB temperature, b) mantle with core evolution and inner core freezing, c) mantle and core with an atmosphere, d) mantle and core with a deep ocean on top (we neglect here a possible atmosphere).

The outline of the paper is as follows. In Section II, we describe the different modules of CHIC, followed by the benchmark results in Section III to show the validity of the code in various applications. Finally, in Section IV, we summarize the possibilities of a coupled 1D - 2D/3D code and the planned future work.

II. MODEL

CHIC uses various modules for the different shells of a terrestrial or ocean planet. The basic modules used in CHIC are listed below. The density of the material and other physical properties are determined as described in Section II-A, and can be applied to both 1D and 2D/3D modules. The input file used for the simulations is similar for all modules, which simplifies comparison of the 1D model with the 2D mantle model.

A. Interior structure model and material properties

We developed an interior structure model to assess the radius of a terrestrial planet for given mass and composition, and to determine the depth-dependent pressure, density and other thermodynamic properties for an initial temperature profile. We assume, that the planet is differentiated into an iron core, a silicate mantle (containing olivine, perovskite and post-perovskite) and an ocean layer. We derive average mantle and core values for the thermodynamic properties (e.g., thermal conductivity) to be used in the thermal evolution models.

To obtain profiles for the depth-dependent pressure p and gravity g , we solve the Poisson equation (1) and the

hydrostatic pressure equation (2). The density ρ , the thermal expansion coefficient α and the heat capacity c_p at local conditions are obtained from equations of states of the relevant materials [6][7].

The gravitational acceleration $g(r)$ depends on the gravity value at the surface of the planet (calculated from its mass and planet radius), the radius r and the density profile $\rho(r)$. G is the gravitational constant.

$$dg/dr = 4\pi G\rho - 2g/r \quad (1)$$

The gravitational and differential pressure forces are assumed to satisfy hydrostatic equilibrium. The pressure p decreases with increasing radius depending on local gravity and density, yielding a pre-defined atmospheric pressure at the surface:

$$dp/dr = -g\rho \quad (2)$$

The mass $m(r)$ below a sphere of radius r increases with increasing radial coordinate until the pre-defined planet's mass M is reached.

$$dm/dr = 4\pi r^2\rho \quad (3)$$

Our interior structure model solves (1) to (3) by integration from the center of the planet outwards (starting from zero gravity and mass in the interior), thus yielding the radius of a planet for given composition (in terms of water, silicate and iron mass fraction) and mass.

B. Core evolution model

Our 1D core evolution module determines the variation of upper core temperature with time via the energy conservation equation

$$\rho_c c_{p,c} V_c \varepsilon_c dT_c/dt = -q_c A_c \quad (4)$$

where the index "c" denotes core values, V_c is the core volume and A_c the core surface area, ε_c is a constant relating the average core temperature to the CMB core temperature, t is the time and q_c is the heat flux from the core into the mantle

$$q_c = -k_m dT/dr|_{r=R_c} \quad (5)$$

We neglect radioactive heat sources in the core, as well as potential tidal heating effects. For 2D/3D convection models, the temperature gradient at the core-mantle boundary (CMB) is calculated over the two bottom shells of the mantle grid with laterally averaged temperatures; in the 1D parameterized model, the temperature drop over the lower thermal boundary layer of the mantle is used instead.

We either treat the core consisting of pure iron or a mixture of iron and lighter elements like sulfur. We can consider possible freezing of the core, when the core temperature falls below the melting temperature. This model, however, only works if the freezing of the core starts at the

core center (leading to a solid inner core as on Earth). This may not be the case for Mercury or Ganymede, where iron may solidify in the upper part of the core and sink down as (so-called) iron snow. We only model planets without the iron snow regime and adopt the model of [2], which determines latent heat released by iron solidification and gravitational energy produced by differentiation of the core into an inner and outer core. Both mechanisms lead to an increasing temperature at the CMB temperature and thus have an influence on the thermal evolution of the mantle, as well. For super-Earths (i.e., planets up to 10 Earth masses), we neglect lighter elements in the core, as the EOS for mixtures of iron and lighter elements have only been derived in a limited pressure range not suitable for super-Earths.

C. Mantle: 1D parameterized model

The 1D module to assess the thermal evolution of the mantle of a terrestrial planet is based on [5][8]. We refer to these references for full details on the model. The model determines the evolution of the upper mantle temperature T_m over time by considering that the loss of energy due to mantle cooling and heat flux out of the mantle is balanced by the heat flux into the mantle and the radioactive heat production in the mantle:

$$\rho_m c_{p,m} V_l \varepsilon_m dT_m/dt = -q_l A_l + q_c A_c + Q_m V_l \quad (6)$$

The index “m” denotes mantle values. V_l is the volume of the mantle from core to lithosphere (thus excluding the conductive lithosphere), and A_l is the area at the boundary between mantle and lithosphere. The constant ε_m relates the average mantle temperature and T_m . The mantle temperature decreases due to heat flux out of the mantle into the lithosphere q_l , increases due to inflowing heat flux from the core q_c and increases with heat released by radioactive heat sources Q_m .

We also consider possible melting events and crust formation over time, leading to additional terms in (6). For details on the crustal evolution, as well as the definition of the thermal boundary layers and calculation of the temperature in the lithosphere, we refer to [8].

Note, that the 1D parameterized model only considers the evolution of the temperature over time, and assumes effective convection. To understand the convection mechanism and its strength depending on mantle parameters and planet size (possibly triggering plate tectonics at the surface), a more sophisticated 2D/3D convection model is needed.

D. Mantle: 2D / 3D convection model

The CHIC code uses a finite volume (FV) field approach to solve the conservation equations of mass, momentum and energy. A finite grid is placed in the mantle, with shells from the CMB to the planet surface, and a predefined number of grid points per shell. We then define Voronoi cell volumes around each grid point and solve the system of equations on each cell volume considering the flux in and out of the cell and the energy production in the cell, see Figure 2.

The grid is either defined in Cartesian coordinates in a 2D or 3D box or in polar coordinates for a 2D cylindrical sphere (a cut through the planet at the equator representing the temperature profile of a cylinder with the 2D plane as a basis) or a 2D spherical annulus (an equatorial cut that mimics the temperature profile of a sphere in 3D, [9]). For the 2D models with spherical or cylindrical geometry, it is often useful to employ a regional sector of the 2D spherical model (as shown in Figure 4) to save computational power. In addition to the grid, randomly distributed particles, that move along the convective stream lines, are used to transport local information as for example the water content.




Figure 2. Geometries implemented in CHIC. Top: 2D Cartesian box and 3D Cartesian box. Bottom: 2D cylinder and 2D spherical annulus.

We solve the equation system for an incompressible medium with the Extended-Boussinesq approximation (EBA), which yields an adiabatic temperature increase with depth (see [5] for details on the model). For a dissipation number Di of zero, the formulation reduces to the Boussinesq approximation (BA).

In the EBA approximation, the non-dimensional conservation equations of energy, mass and momentum can be expressed as [10]

$$\partial T/\partial t + \mathbf{u} \cdot \nabla T + Di(T+T_0)u_r = \nabla^2 T + 0.5Di \eta \dot{\varepsilon}_{II}^2/Ra + H(7)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (8)$$

$$\nabla P - \nabla[\eta(\nabla \mathbf{u} + \nabla \mathbf{u}^T)] = Ra T \mathbf{e}_r \quad (9)$$

Here, T is temperature, T_0 surface temperature, t time, and Di the dissipation number. The convective pressure is denoted by P ; \mathbf{u} is the velocity and u_r the radial velocity, whereas \mathbf{e}_r is the radial unit vector. The Rayleigh number Ra is a measure for the convective vigor of the material, and H is the heat source (e.g., radioactive heat source). η is the viscosity and $\dot{\varepsilon}_{II}$ is the second invariant of the strain rate.

The equations (7)-(9) are written in a non-dimensional form, as is typically done for convection simulations [11]. The non-dimensionalization is obtained by dividing the dimensional value of each variable by a reference value as

given in [11]. The quantities given in Section III are also non-dimensional values.

The energy equation is solved with a second-order implicit Euler method. For the results reported here, we also apply an upwind scheme. To solve the conservation equation of mass and momentum, we either use a direct solver or a coupled mass and momentum solver. The direct solver uses one solver matrix for (8) and (9) and applies a penalty formulation following [12]. The iterative, coupled solver employs a SIMPLER pressure correction algorithm following [12][13]. In this paper, we apply the direct solver.

The resulting linear equations (for mass, momentum and energy equation) are solved iteratively using either the Fortran Pardiso solver [14] or a biconjugate gradient (BiCG) solver with an under-relaxation scheme. The BiCG solver is slower compared to the Pardiso solver but can be easily parallelized.

The equations above depend on the viscosity of the material η . The viscosity depends on several factors including the temperature, pressure, grain size, water content and strain rate of a material. In Earth's mantle, creep of minerals is typically described by dislocation creep (motion of dislocations through the crystal lattice) and diffusion creep (deformation of crystalline solids by the diffusion of vacancies through their crystal lattice). The latter is largely independent of the strain-rate, whereas dislocation creep does not depend on the grain size. In CHIC, the user can choose between a dislocation viscosity, a diffusion viscosity and a mix of both formulations. The smaller viscosity is the dominant viscosity for material motion.

The general equation that we use for the viscosity follows an Arrhenius law [15][16]

$$\eta = A \dot{\epsilon}_{II}^{1-n/n} d^{p/n} C_{OH}^{-r/n} \exp((E+pV)/(nRT)) \quad (10)$$

A is a material-dependent constant, n is the stress exponent, d is the grain size, C_{OH} is the concentration of water (for dry materials $r=0$), r is the water exponent, E the activation energy and V the activation volume. p is the pressure and R the gas constant. Note, that the pressure p is the hydrostatic pressure and not the convective pressure as in (9). The parameters for both diffusion and dislocation creep are taken from [15][16] for both wet and dry materials. The concentration of water C_{OH} is traced via particles in CHIC and does not only influence the viscosity, but also the local melt temperature, which is smaller in the presence of water than for dry materials [17].

Note, that even though the Arrhenius viscosity (10) is preferentially used for simulations of terrestrial planets, for benchmarks and basic convection simulations, often an approximated viscosity is used, the so-called Frank-Kamenetskii approximation (FKA), given by

$$\eta = A \exp(-\gamma_T T + \gamma_p z) \quad (11)$$

Here, γ_T and γ_p are either the logarithm of a pre-defined viscosity contrast with respect to temperature or pressure, respectively, or they are derived from the parameters in (10)

[18]. z is the non-dimensional depth (0 at the surface and 1 at the CMB). Note, that for the application to terrestrial planets (especially for plate tectonics planets), the FKA (11) is not suitable and the Arrhenius viscosity (10) should be applied [18].

E. Additional modules

CHIC can also be used for ocean planets and icy moons, where a silicate-iron shell is covered by a deep water or ice sphere. Another module treats the evolution of the atmosphere with time, where we consider outgassing of greenhouse gases CO_2 and H_2O . The model can be applied to planets with a Mars- or Venus-like atmosphere [3].

III. RESULTS AND DISCUSSION

To validate our code, we applied several benchmark tests for the convection module. To our knowledge, unlike for the mantle convection calculation, benchmark results for the 1D parameterized model have not been published. Therefore, we have validated our code by reproducing results of [8]. The results are very similar [19], but differ in detail because not all parameters used in the studies are known. The module has been integrated into the CHIC code and has been extended to include a regolith layer and compared to [20], yielding again comparable results.

We compared our 1D parameterized model to the convection module by applying the 2D spherical annulus. In Figure 3, we plot the thermal evolution of Mars determined for a Boussinesq approximation, a Newtonian viscosity law ($n=1$ in (10)) and fit the pre-factor such that we obtain a reference viscosity of 10^{20} Pas at 1600 K and 3 GPa. We use an activation energy of $E=300$ kJ/mol and an activation volume of $2.5 \text{ cm}^3/\text{mol}$. The initial mantle temperature is 2000 K and the CMB temperature is 2300 K, the surface temperature is set to 220 K. Heat sources are homogeneously distributed in the mantle and are taken Earth-like [5]. For the 2D model, we use a quarter sphere with a radial resolution of 80 shells.




Figure 3. Upper temperature, CMB temperature and lid thickness for a thermal evolution of Mars applying either the 2D spherical annulus (black curve) or the 1D parameterized model (red).

The mantle initially heats up due to radioactive decay, convection then leads to efficient cooling of the mantle. In the convection model, we define the lid over the depth where the conductive heat transport is more efficient than the convective heat transport. The lid thickness is then fitted by a third-order polynomial since oscillations occur. The obtained lid thickness is in the beginning larger than for the 1D model

(where we plot the total conductive layer thickness including both the lid and the upper thermal boundary layer), but shows a similar increase with time after 2Gyr. The different lid thicknesses at the beginning of the evolution can be explained by a delayed on-set of convection in the 2D model, which also leads to a slightly weaker mantle cooling at the beginning and hence a shift in the upper mantle temperature compared to the 1D model.

For the convection model, we applied several standard benchmarks published in the literature. The most basic benchmark has been developed for convection in a 2D Cartesian box [21] and uses either isoviscous convection or temperature- and pressure-dependent viscosity in the Boussinesq approximation. A free-slip boundary condition is applied to the walls of the box. The non-dimensional temperature at the surface of the box is set to 0 and at the bottom to 1. The simulations are run until steady-state is reached (i.e., temperature variations drop below a tolerance value of 10^{-10}).

In Table I, we compare our results (for a fixed resolution of 80(200)x80 cells, depending on the aspect ratio) to published results. Note, that in [21] different resolutions have been used, and we give the min and max values for resolutions of at least 33x33 cells. We list the three most important quantities: the root-mean-square (RMS) velocity, the maximum of the upper mantle temperature profile at the middle ($x=0.5\ell$, where ℓ is the length divided by height, i.e., the aspect ratio) and the surface Nusselt number, which is a measure of the ratio of convective to conductive heat transport at the surface of the box. For more information on the benchmark setup we refer the reader to [21]. CHIC yields results in good agreement with all cases published in [21], see Table I, and lies either in the range of published results or differs by less than 0.5%.

TABLE I. BENCHMARK COMPARISON OF CHIC (CH) TO [21] (BL).

	RMS velocity		Max temperature		Nusselt number	
	CH	BL	CH	BL	CH	BL
1a	42.92	42.74-42.87	0.425	0.421-0.427	4.920	4.864-4.896
1b	194.3	192.4-198.0	0.432	0.415-0.437	10.60	10.42-10.69
1c	835.1	823.7-842.5	0.440	0.431-0.446	21.81	21.08-22.07
2a	496.6	458.3-503.3	0.725	0.716-0.741	10.43	10.04-10.07
2b	183.1	166.7-193.1	0.390	0.385-0.403	7.271	6.806-7.409

C1: isoviscous material, $\ell=1$, a) $Ra=1e4$, b) $Ra=1e5$, c) $Ra=1e6$.
 C2: FKA (11), a) $Ra_{surf}=1e4$, $\gamma_T=\ln(1000)$, $\gamma_P=0$, $\ell=1$, b) $Ra_{surf}=1e4$, $\gamma_T=\ln(16384)$, $\gamma_P=\ln(64)$, $\ell=2.5$.

A similar benchmark has been published in [10] for the Extended Boussinesq approximation and compressible materials in a 2D Cartesian box. We compare the benchmark results using EBA (see Section IID) for a resolution of 80x80 cells to the results in [10], see Table II.

CHIC compares well to the published results for EBA with deviations of few percent at most. A similar benchmark for a 2D cylindrical shell including compressibility is in preparation for publication and includes 10 different codes, including CHIC. The accuracy of the CHIC code is comparable to the Cartesian simulations.

TABLE II. BENCHMARK COMPARISON OF CHIC (CH) TO [10] (KI).

	RMS velocity		Average temperature		Nusselt number	
	CH	KI	CH	KI	CH	KI
1a	38.38	38.39-38.50	0.4911	0.4909-0.4914	4.125	4.05-4.10
1b	61.43	61.35-61.50	0.4935	0.4937-0.4942	5.203	5.10-5.16
1c	111.7	111.5-111.7	0.4980	0.4987-0.4996	6.997	6.86-6.95
1d	174.0	173.5-174.2	0.5018	0.5033-0.5045	8.710	8.54-8.66
1e	269.2	267.9-269.5	0.5056	0.5083-0.5101	10.80	10.6-10.7
1f	474.2	465.0-466.8	0.5097	0.5161-0.5178	14.25	13.9-14.1
2a	33.70	33.76-33.92	0.4822	0.4816-0.4823	3.408	3.34-3.38
2b	54.14	54.17-54.41	0.4854	0.4851-0.4861	4.255	4.16-4.22
2c	98.36	98.80-98.86	0.4920	0.4926-0.4940	5.648	5.52-5.60
2d	152.3	152.3-153.4	0.4980	0.4997-0.5017	6.957	6.79-6.90
2e	232.8	231.0-232.8	0.5041	0.5080-0.5109	8.516	8.26-8.40
3a	23.78	23.95-24.24	0.4678	0.4658-0.4671	2.209	2.15-2.19
3b	38.77	39.07-39.50	0.4685	0.4664-0.4682	2.675	2.60-2.65
3c	70.22	70.85-71.66	0.4734	0.4719-0.4747	3.399	3.30-3.36
3d	106.4	107.1-108.2	0.4793	0.4789-0.4823	4.031	3.89-3.97
3e	152.3	147.0-148.0	0.4868	0.4893-0.4938	4.663	4.40-4.44

C1: $Di=0.25$, $\ell=1$, a) $Ra=1e4$, b) $Ra=2e4$, c) $Ra=5e4$, d) $Ra=1e5$, e) $Ra=2e5$, f) $Ra=5e5$.
 C2: $Di=0.5$, $\ell=1$, a) $Ra=1e4$, b) $Ra=2e4$, c) $Ra=5e4$, d) $Ra=1e5$, e) $Ra=2e5$.
 C3: $Di=1.0$, $\ell=1$, a) $Ra=1e4$, b) $Ra=2e4$, c) $Ra=5e4$, d) $Ra=1e5$, e) $Ra=2e5$.

A two-code benchmark for the different geometries (between CHIC and GAIA [22]) has been realized for a Boussinesq material [23]. For the 2D box, we apply a resolution of 80(160 for $\ell=2$)x80 cells, for the 3D box 20x20x20 cells and for the 2D shells we apply 80 shells in radial direction with 754, 377, 189, 440, 419 and 754 points per shell for the 6 considered cylindrical/spherical cases. Note, that we compare the 2D spherical annulus of CHIC to the case of 3D sphere of GAIA.

Both codes are well in agreement, with deviations below 5% apart from the 3D box (12.5% deviation), where we applied a lower resolution than in [23], see Table III. The plots in Figure 4 show the steady-state for all cases.

TABLE III. BENCHMARK COMPARISON OF CHIC (CH) TO [23] (NT).

Case	RMS velocity		Average temperature		Top Nusselt number	
	CH	NT	CH	NT	CH	NT
2D box, $\ell=1$, RBC ^a	55.73	53.06	0.6904	0.6872	2.001	1.956
2D box, $\ell=2$, RBC ^a	55.73	53.79	0.6904	0.6871	2.001	1.956
2D box, $\ell=1$, PBC ^b	56.88	54.62	0.7038	0.6993	2.069	2.014
3D box, $\ell=1$, RBC	65.43	57.21	0.6997	0.6927	2.200	2.363
2D full cylinder ^c	36.24	35.25	0.5744	0.5711	1.439	1.439
2D half cylinder	35.64	34.84	0.5751	0.5725	1.435	1.440
2D quarter cylinder	35.64	34.87	0.5751	0.5725	1.435	1.440
2D cylinder, CV ^d	16.80	17.11	0.4362	0.4377	0.967	0.995
2D cylinder, CR ^e	14.46	14.51	0.4046	0.4039	0.891	0.914
3D sphere ^f	15.56	16.19	0.3657	0.3374	0.782	0.744

^aWe apply a surface Rayleigh number of $Ra=10$ and a FKA (11) viscosity contrast of $1e5$.

^bRBC stands for reflective boundary condition at the side wall with free-slip boundary.

^cPBC stands for periodic boundary conditions.

^dThe sphere uses a radius ratio of 2, i.e., the core radius is half the planet radius.

^eCR means corrected volume such that ratio of core area divided by mantle volume is as in 3D.

^fWe use a 2D spherical annulus instead of a 3D sphere with 4 initial plumes instead of 6.




Figure 4. Convection patterns obtained with CHIC for different available geometries. See text and Table III for more details.

The 2D spherical annulus furthermore has been compared to the results in [9]. The non-dimensional radius of the core is 1.2222 and the planet radius is 2.2222. We use a resolution of 32 shells with 256 points on each shell. The CHIC results are well in agreement with the published steady-state results, with differences of not more than 6%, see Table IV.

TABLE IV. BENCHMARK COMPARISON OF CHIC (CH) TO [9] (HT).

	<i>Ra</i>	Average RMS velocity		Nusselt number	
		CH	HT	CH	HT
1	1e4	39.87	37.7	4.39	4.18
	1e5	174.1	~160	7.61	~7.39
	1e6	719.3	~640	16.53	~14.4
	<i>Ra / H</i>	Average RMS velocity		Average mantle temperature	
2	1e4 / 3.4	25.09	23.5	0.295	0.308
	1e5 / 6.6	~97	~78.5	~0.369	~0.349
	1e6 / 14	~340	~265	~0.385	~0.350

isoviscous material, case 1: bottom-heated convection, case 2: internally-heated convection.

For time-dependent simulations (indicated by “~”), larger deviations can appear between different codes (here up to 22%). For this reason typically only steady-state simulations are used in community benchmarks.




Figure 5. Simulation of the subduction process of an oceanic plate including dehydration of the subducted slab and related melting processes.

Recently, an increasing attention has been drawn to benchmarks for time-dependent simulations, for example for surface mobilization in a subduction zone (Figure 5, in preparation for publication), or for plastic deformation and episodic overturn [24].

IV. CONCLUSION AND FUTURE WORK

CHIC is a new, advanced numerical code developed at the Royal Observatory of Belgium, and can be applied to different geodynamic applications, including the simulation of the thermal evolution of terrestrial planets. Both 1D and 2D/3D geometries can be applied to the silicate mantle to investigate the thermal evolution (as well as convection and surface mobilization for the 2D models). The core, the ocean layer and the atmosphere are solved with the parameterized model. The thermal evolution of Mars’ mantle has been investigated with both the 1D model and the 2D spherical annulus model with comparable results. Furthermore, the code has been validated by comparison to published 1D parameterized models and benchmarks for the 2D/3D mantle convection model for all available geometries (Figure 4). CHIC is in good agreement with literature values. New benchmark projects with contributions from CHIC are currently in preparation for publication or in review. Readers interested in access to the code are asked to directly contact the authors of the paper.

ACKNOWLEDGMENT

L. Noack has been funded by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office through the Planet Topers alliance. A. Rivoldini was supported by the PRODEX program managed by the European Space Agency in collaboration with the Belgian Federal Science Policy Office. This work results within the collaboration of the COST Action TD 1308. We thank Clemens Heistracher, Nastasia Zimov, François Labbé, and Thomas Boiveau for their contributions to the project.

REFERENCES

- [1] P. van Thienen et al., “Water, Life, and Planetary Geodynamical Evolution”, *Space Sci. Rev.*, vol. 129, 2007, pp. 167-203.
- [2] G. Schubert, M. N. Ross, D. J. Stevenson, and T. Spohn, “Mercury’s thermal history and the generation of its magnetic field”, in *Mercury*, 1998, pp. 429-460.
- [3] L. Noack, D. Breuer, and T. Spohn, “Coupling the atmosphere with interior dynamics: Implications for the resurfacing of Venus”, *Icarus*, vol. 217, 2012, pp. 484-498.
- [4] C. Gillmann and P. Tackley, “Atmosphere/mantle coupling and feedbacks on Venus”, *JGR Planets*, vol. 119, 2014, pp. 1189-1217.
- [5] G. Schubert, D. L. Turcotte, and P. Olson, “Mantle Convection in the Earth and Planets”, Cambridge University Press, Cambridge, 2001.
- [6] L. Stixrude and C. Lithgow-Bertelloni, “Thermodynamics of mantle minerals--I. Physical properties”, *GJI*, vol. 162(2), 2005, pp. 610-632.
- [7] J. Bouchet, S. Mazevet, G. Morard, F. Guyot, and R. Musella, “Ab initio equation of state of iron up to 1500 GPa”, *Physical Review B*, vol. 87 - 094102, 2013, pp. 1-8.
- [8] A. Morschhauser, M. Grott, and D. Breuer, “Crustal recycling, mantle dehydration, and the thermal evolution of Mars”, *Icarus*, vol. 212(2), 2011, pp. 541-558.
- [9] J. W. Hernlund and P. Tackley, “Modeling mantle convection in the spherical annulus”, *PEPI*, vol. 171, 2008, pp. 48-54.
- [10] S. D. King et al., “A community benchmark for 2D cartesian compressible convection in the Earth’s mantle”, *GJI*, vol. 180(1), 2010, pp. 73-87.

- [11] U. R. Christensen, "Convection with pressure- and temperature-dependent non-Newtonian rheology", *Geophys. J. R. astr. Soc.*, vol. 77, 1984, pp. 343-384.
- [12] S. Zhong, D. A. Yuen and L. N. Moresi, "Numerical methods in mantle convection", *Treatise on Geophysics*, vol. 7, 2007, pp. 227-252.
- [13] S. V. Patanker, "A calculation procedure for two-dimensional elliptic situations", *Numerical Heat Transfer V*, 1981, pp. 409-425.
- [14] O. Schenk and K. Gärtner, "On fast factorization pivoting methods for symmetric indefinite systems", *Elec. Trans. Numer. Anal.*, vol. 23, 2006, pp. 158-179.
- [15] S.-i. Karato and P. Wu, "Rheology of the upper mantle: a synthesis", *Science*, vol. 260, 1993, pp. 771-778.
- [16] G. Hirth and D. Kohlstedt, "Rheology of the upper mantle and the mantle wedge: a view from the experimentalists", in Eiler, J. (Ed.), *Inside the Subduction Factory. Geophysical Monograph Series*, vol. 138, AGU, Washington, D.C., 2003, pp. 83-105.
- [17] R. Katz, M. Spiegelman, and C. Langmuir, "A new parameterization of hydrous mantle melting", *G³*, vol. 4(9) - 1073, 2003, pp. 1-19.
- [18] L. Noack and D. Breuer, "First- and second-order Frank-Kamenetskii approximation applied to temperature-, pressure- and stress-dependent rheology", *GJI*, vol. 195, 2013, pp. 27-46.
- [19] T. Boiveau, "Convection mantellique dans les planètes telluriques". Rapport de stage, unpublished.
- [20] N. Tosi, M. Grott, A.-C. Plesa, and D. Breuer, "Thermochemical evolution of Mercury's interior", *JGR Planets*, vol. 118, 2013, pp. 1-14.
- [21] B. Blankenbach et al., "A benchmark comparison for mantle convection codes", *GJI*, vol. 98, 1989, pp. 23-38.
- [22] C. Hüttig and K. Stemmer, "Finite volume discretization for dynamic viscosities on Voronoi grids", *PEPI*, vol. 171, 2008, pp. 137-146.
- [23] L. Noack and N. Tosi, "High-performance modelling in geodynamics", in *Integrated Information and Computing Systems for Natural, Spatial, and Social Sciences*, ed. Rückemann, C.-P., IGI Global, 2013, pp. 324-352.
- [24] N. Tosi et al., "A community benchmark for viscoplastic thermal convection in a 2-D square box", *G³*, in review.

Creation of Objects and Concordances for Knowledge Processing and Advanced Computing

Claus-Peter Rückemann

Westfälische Wilhelms-Universität Münster (WWU),
Leibniz Universität Hannover,
North-German Supercomputing Alliance (HLRN), Germany
Email: ruckema@uni-muenster.de

Abstract—This paper presents the summary of the main results from creating objects and concordances used for advanced processing and computing based on knowledge resources. Today big data collections and resources sadly combine one or more deficits of being unclassified, unstructured, isolated, and weakly developed on the one hand and in consequence only accessible with insufficient simplistic means on the other hand. The goal of this research is to create new classification features, structures, and components, which can be flexibly used with multi-disciplinary, multi-lingual long-term knowledge resources. The focus is to develop new extended facilities for universal long-term knowledge resources beyond the simple and isolated use of knowledge and data.

Keywords—*Creation of Objects and Concordances; Universal Decimal Classification; Sustainability; Knowledge Processing; Advanced Computing.*

I. INTRODUCTION

This paper summarises the results from the development of advanced object features for long-term knowledge resources, which can be used for universal documentation and consequent purposes. For the consequent purposes like knowledge discovery, further creation and development of knowledge resources, visualisation, and education we can implement suitable, data-centred structures being usable with dynamical computing, knowledge processing, advanced computing.

Up to now the world of increasingly big data is limited to data and data collections, which are rapidly growing in quantity, mostly even growing in storage requirements instead of knowledge only. Besides the data being unstructured, isolated, and often inconsistent in content and form it is also missing quality and essential features for conceptual knowledge like classification. One consequence is that in the last decades the means for accessing and handling data have not changed a lot regarding the content, context, and a next generation of features and quality.

The focus of this research aims on the sustainable creation of long-term facilities for integrated documentation and applicability. The facilities provide advanced features for processing of knowledge as well as for flexible computing. Therefore, the creation and long-term care for suitable knowledge objects is a central issue. The knowledge resources have to be able to document any knowledge and data, e.g., factual and procedural knowledge, which require vertical as well as horizontal scalability for individual and subsequent use.

In order to cope with the deficits this architecture can integrate structured and unstructured data, support universal classification and concordances and it can enable advanced knowledge processing, like parallel processing and dynamical visualisation. This paper presents the up-to-date research results from the creation of long-term knowledge resources' components, structures, and workflows for advanced processing and computing – and in the end most important on the long run, fostering the investments in the development of the data itself. Therefore, the major contributions of this research are the content and context on the one hand and on the other hand the new practical insights on improving the state-of-the-art of long-term documentation and application of universal knowledge.

This paper is organised as follows. Section II introduces the state-of-the-art and the motivation for creating an overall system. Sections III and IV present the data-centric architecture, the implementation, and resources, especially the creation of objects and concordances. Section V shows the implemented features for knowledge processing and computing and Section VI provides an evaluation. Section VII presents the main results, summarises the lessons learned, conclusions and future work.

II. STATE-OF-THE-ART AND MOTIVATION

Many developments have contributed to the state-of-the-art on knowledge processing and discovery. Also, many developments have provided new technical means to cope with the new developments in computing architectures and services. In the context of knowledge processing, discovery, and search we are often faced with “prominent” examples like Internet search engines, library search engines, specialised expert systems, and maybe social media systems. If at all then there are some interfaces, e.g., for automated requests or web-service creation.

The algorithms applicable to this kind of ‘art’ are very limited and often not sufficient in delivering a reasonable quality for requested results or for following advanced goals. Therefore, various concepts, developments, and approaches have been created, addressing different aspects and special purposes. Nevertheless, these tools, classifications, and algorithms only try to handle the symptoms of the state of the data.

The approaches cover in-depth classifications and handling, e.g., library classification specialised on geological publications [1], handling historical geographic resources, especially in library context [2], and international patent classification [3].

The algorithms touch processing and automation, e.g., statistical models for online text classification [4] and automation with a classification [5]. The discussions and analyses range from research aspects to reliability and non-disciplinary approaches, e.g., classification as a research tool [6], reliability of diagnostic classification [7], and the Universal Decimal Classification (UDC) [8] as a non-disciplinary [9] universal [10] classification system, and legal and general aspects within Information Science, Security, and Computing [11].

In depth, aspects of mapping, organisation and multi-lingual data have been discussed, e.g., simple mapping between a classification and an “index” [12], simple conceptual methods for using classification in libraries [13], knowledge organisation [14], multi-lingual lexical linked data cloud [15]. In principle, any multi-disciplinary data resources may be used, e.g., projects like Europeana [16], European Cultural Heritage Online (ECHO) [17] or World Digital Library (WDL) [18]. Although these examples are focussed on providing special information they lack in sufficient content, organisation, and structure.

The main motivation for this research was the lack of multi-disciplinary data-centric approaches, which can be used for long-term creation of knowledge and scalable implementations. The data used here is based on the content and context from the knowledge resources, provided by the LX Foundation Scientific Resources [19]. The LX knowledge resources’ structure and the classification references [20] based on UDC [21] are essential means for the processing workflows and evaluation of the knowledge objects and containers. Both provide strong multi-disciplinary and multi-lingual support. For this part of the research all small unsorted excerpts of the knowledge resources objects only refer to main UDC-based classes, which for this part of the publication are taken from the Multilingual Universal Decimal Classification Summary (UDCC Publication No. 088) [8] released by the UDC Consortium under the Creative Commons Attribution Share Alike 3.0 license [22] (first release 2009, subsequent update 2012). The analysis of different classifications, development of concepts for intermediate classifications, and experiences from case studies from the research conducted in the Knowledge in Motion (KiM) long-term project [23] have contributed to the application of UDC and different classifications and concordance schemes in the context of knowledge resources.

The following term definitions for object, container, and matrix can be helpful in this context. An object is an entity of knowledge data being part of knowledge resources. An object can contain any documentation, references, and other data. Objects can have an arbitrary number of sub-objects. A container is a collection of knowledge objects in a conjoint format. A matrix is a subset of the entirety, the “universe”, of knowledge. A workflow can consist of many sub-workflows each of which can be based on an arbitrary number of knowledge matrices. The output of any sub-workflow or workflow can be seen as an intermediate or final result matrix.

III. DATA-CENTRIC IMPLEMENTATION

In order to concentrate on the challenges of the data itself so-called data-centric, data-defined, document-oriented or

document-centric approaches have been developed. This went along with extending features like Structured Query Language (SQL) and “Not only SQL” (NoSQL) [24], e.g., via MySQL [25] and MongoDB [26] and in consequence [27] also in bridging relational and data- or document-centric approaches [28]. The very minimalistic “map” and “reduce” functions approach of MapReduce [29], which attracted many quick and simple solutions is a nice example for building simple workflow elements. As the knowledge resources’ approach [19] is even much more general [20] it allows for arbitrary measures and also for processing implementing map and reduce functions, which can be based on the creation of objects and concordances.

Regarding a distributed computer system theoretical computer science can state the CAP (Consistency, Availability, Partition tolerance) theorem. In condensed form this means: *Consistency*: All nodes see the same data at the same time, *Availability*: A guarantee that every request receives a response about whether it succeeded or failed, *Partition tolerance*: The system continues to operate despite arbitrary message loss or failure of part of the system. Accordingly, learning from decades of case-studies, regarding the long-term knowledge and information sciences we can state a “CLU” theorem:

- *Consistency*: All knowledge in context used is neither in contradiction to other knowledge in context nor disacording within its content,
- *Long-term sustainability*: Data-centric architectures, the core knowledge resources can be used for an arbitrary number of different implementations,
- *Universal documentation*: Documentation is supported for any knowledge and data, e.g., factual, conceptual, procedural, and metacognitive knowledge.

There is no direct reasonable equivalent for the P and A aspects. Besides the consistency, the items much more important are the long-term sustainability and universal documentation aspects. This includes the requirements for any type of knowledge as well as its multi-lingual documentation and features.

IV. IMPLEMENTATION AND RESOURCES

The implementation for dynamical visualisation and computation is based on the framework for the architecture for documentation and development of advanced scientific computing and multi-disciplinary knowledge [30]. The architecture implemented for an economical long-term strategy is based on different development blocks. Figure 1 shows the three main columns: Application resources, knowledge resources, and originary resources. The central block in the “Collaboration house” framework architecture [20], are the knowledge resources, scientific resources, databases, containers, and documentation (e.g., LX [19], databases, containers, list resources). These can be based on and refer to the originary resources and sources (photos, scientific data, literature). The knowledge resources are used as a universal component for compute and storage workflows. Application resources and components (Active Source, Active Map, local applications) are implementations for analysing, utilising, and processing

data and making the information and knowledge accessible. The related information, all data, and algorithm objects presented are copyright the author of this paper, LX Foundation Scientific Resources [19], all rights reserved. The structure and the classification references based on the LX resources and UDC, especially mentioning the well structured editions [8] and the multi-lingual features [21], are essential means for the processing workflows and evaluation of the knowledge objects and containers. Both provide strong multi-disciplinary and multi-lingual support.

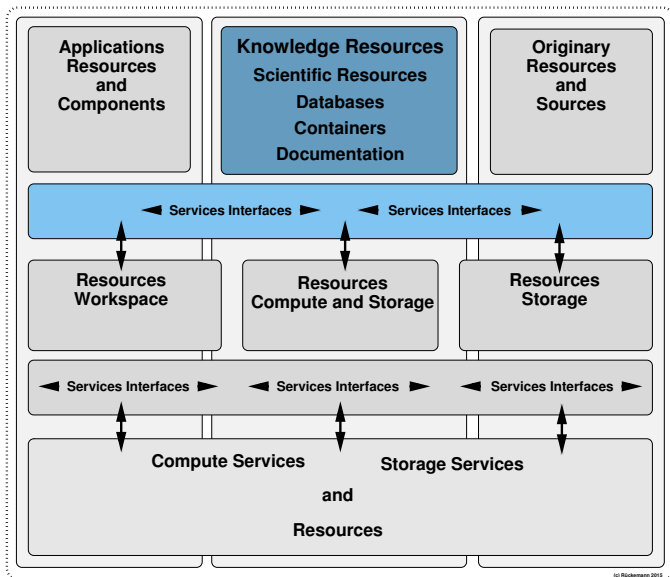


Figure 1. Architecture: Columns of practical dimensions. The knowledge resources are the central component within the long-term architecture.

The three blocks are supported by services’ interfaces. The interfaces interact with the physical resources, in the local workspace, in the compute and storage resources the knowledge resources are situated, and in the storage resources for the ordinary resources.

All of these components do allow for advanced scientific computing and data processing, as well as the access of compute and storage resources via services interfaces. The resources’ needs depend on the application scenarios to be implemented for user groups.

A. Creation of objects

Practical creation of objects has shown to be most efficient when three different categories of creation are considered:

- Manually created objects,
- Hybrid (semi-automatically) created objects, and
- Automatically created objects.

In any case creating objects is supported by universal classification, e.g., references to UDC. Therefore, that can also be applied for the creating concordances with objects. The listing in Figure 2 shows an instance of a simple object excerpt.

```

1 Vesuvius [Volcanology, Geology, Archaeology]:
2 (lat.) Mons Vesuvius.
3 (ital.) Vesuvio.
4 (deutsch.) Vesuv.
5 Volcano, Gulf of Naples, Italy.
6 Complex volcano (compound volcano).
7 Stratovolcano, large cone (Gran Cono).
8 Volcano Type: Somma volcano,
9 VNUM: 0101-02=,
10 Summit Elevation: 1281 m.
11 The volcanic activity in the region is observed by the
12 Oservatorio Vesuviano. The Vesuvius area has been
13 declared a national park on 1995-06-05.
14 The most known antique settlements at the Vesuvius are
15 Pompeji and Herculaneum.
16 Syn.: Vesaevus, Vesevus, Vesbius, Vesvius
17 s. volcano, super volcano, compound volcano
18 s. also Pompeji, Herculaneum, seismology
19 compare La Soufrière, Mt. Scenery, Soufriere
20 ...
    
```

Figure 2. Processed instance of a simple object (excerpt).

The listing in Figure 3 shows an instance of a simple container entry excerpt from a volcanological features container.

```

1 ...
2 CONTAINER_OBJECT_EN_ITEM: Vesuvius
3 CONTAINER_OBJECT_DE_ITEM: Vesuv
4 CONTAINER_OBJECT_EN_PRINT: Vesuvius
5 CONTAINER_OBJECT_DE_PRINT: Vesuv
6 CONTAINER_OBJECT_EN_COUNTRY: Italy
7 CONTAINER_OBJECT_DE_COUNTRY: Italien
8 CONTAINER_OBJECT_EN_CONTINENT: Europe
9 CONTAINER_OBJECT_DE_CONTINENT: Europa
10 CONTAINER_OBJECT_XX_LATITUDE: 40.821N
11 CONTAINER_OBJECT_XX_LONGITUDE: 14.426E
12 CONTAINER_OBJECT_XX_HEIGHT_M: 1281
13 CONTAINER_OBJECT_EN_TYPE: Complexvolcano
14 CONTAINER_OBJECT_DE_TYPE: Komplex-Vulkan
15 CONTAINER_OBJECT_XX_VNUM: 0101-02=
    
```

Figure 3. Processed instance of a simple container entry (excerpt).

The excerpts have been processed with the appropriate `lx_object_volcanology` and `lx_container_volcanology` interfaces, selecting a number of items and for the container also items in English and German including a unique formatting. The resources’ access and processing can be done in any programming language, assuming that the interfaces are implemented. For example, combining scripting, filtering, and parallel programming can provide flexible approaches.

B. Creation of concordances

Many disciplines and large fields of application have developed and used individual adapted frameworks of conceptual knowledge for their purposes. The reasons have been multifold, in that cases either developing a universal approach was too demanding or a distinction for certain reasons might have been considered adequate. In many cases, various classifications required to be “compared” and to be used together.

However, when developing content with conceptual knowledge and classifications sooner or later also the individual classifications get in the focus of development and may require to be “mapped”. In most cases this can be done with the means of concordances, for example, concordances with classification in medicine and health [31] or the creation concordances between two classifications systems [32], in

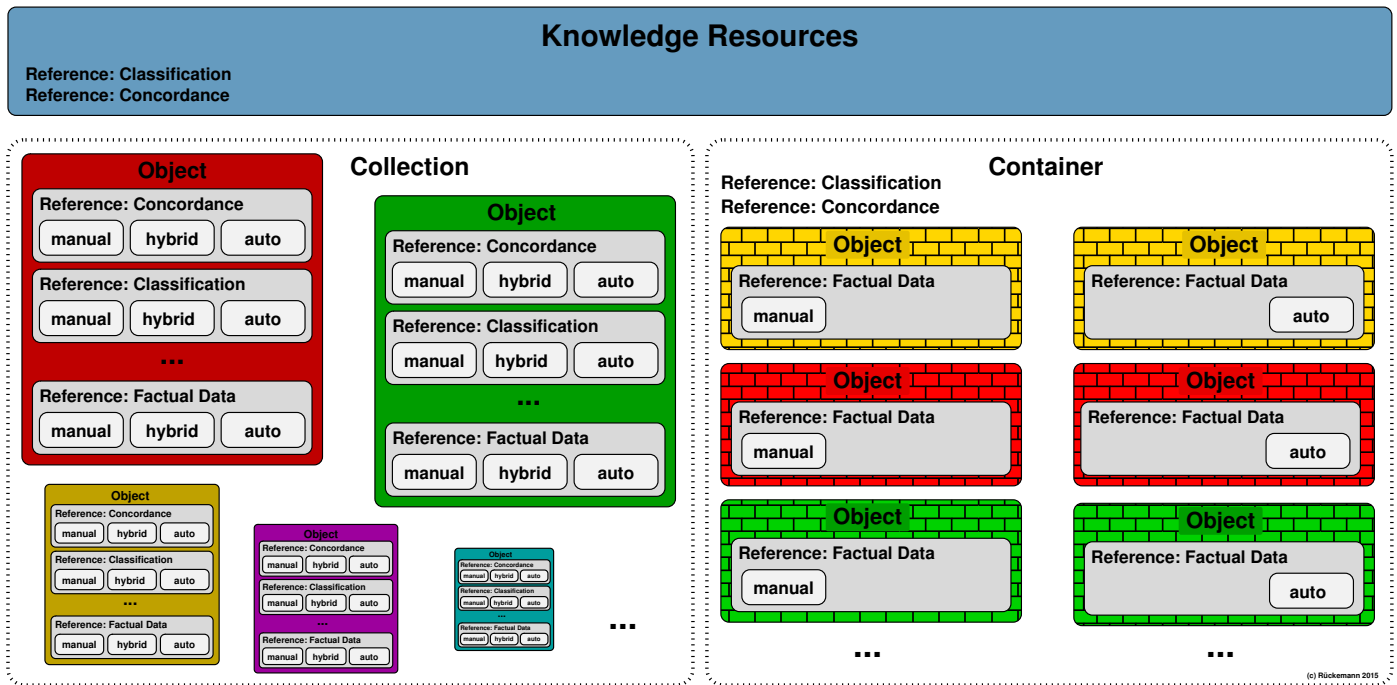


Figure 4. Resources and objects: Selected knowledge resources’ objects containing references for concordances and classifications in collections and containers. In this case, the excerpt shows a distinct handling of manually, hybrid, and automatically created data, especially regarding classifications and concordances.

concordance projects like coli-conc [33] or for benefits in industry classification systems [34]. Therefore, the organisation of the resources and objects is significant for the long-term aspects and the vitality of the data. Taking advantage of the modular architecture of the overall resources (Figure 1) the main objectives are the knowledge resources, services, and interfaces, which are deployed for creating workflows. Figure 4 illustrates an excerpt of selected knowledge resources’ objects. The selected objects are associated to collections and containers and contain references to concordances and classifications. The excerpt in this case shows a distinct handling of manually, hybrid, and automatically created data. The collection objects carry mostly only their individual conceptual knowledge, as there are concordances and classification, for example. The container objects are commonly similar types of objects and structures where the container can carry a respective commonly valid conceptual knowledge for the container (symbolised in the figure by brick-structures for the objects). The respective knowledge resource, on the level integrating collections and resources, can also contain a respective commonly valid conceptual knowledge for the resource.

There are different ways of handling the processes for semi-automatically and automatically created concordances. With the main focus on processing and advanced computing we concentrate of the object and references side of the resources. This concept has shown vital benefits, which enables implementations with comparably high flexibility. Disciplines, services, and resources can be integrated in a very scalable way. Practical creation of concordances has shown to be

most efficient when three different categories of creation are considered:

- Manually created concordances,
- Hybrid (semi-automatically) created concordances, and
- Automatically created concordances.

Manually-created-concordances is a type of concordances, which has resulted from manually inserting references from objects into a concordance instance. Hybrid-created-concordances is a type of concordances, which has resulted from a combination of manual and automated (semi-automatically) processes.

The processes may work on primary concordances or on any level of secondary data in order to support the creation of concordances. Automatically-created-concordances is a type of concordances, which have been generated, e.g., by an automated workflow process. This is mostly done for big data, which are used as quantity data and not due to their quality. In any way, an integration with the knowledge resources’ references and structures is the target.

The workflows can contain several functions comparable to the map and reduce concept. A map function finds the data according to the criteria and creates a map result matrix. A reduce function does the appropriate operation on the map result matrix output. The listing in Figure 5 shows a simple object instance classification and concordances excerpt (Figure 2) from a volcanological object in a collection. The excerpt shows classification concordances in several different classifications as used in different disciplines. Possibly multiple views from different disciplines or author groups

on a certain object are not shown in this reduced view but they can also hold the full spectrum of classifications and concordances.

```

1 ...
2 UCC:UDC2012:551.21
3 UCC:UDC2012:551
4 UCC:UDC2012:902/908
5 UCC:MSC2010:86,86A17,86A60
6 UCC:LCC:QE521-545
7 UCC:LCC:QE1-996.5
8 UCC:LCC:QC801-809
9 UCC:LCC:CC1-960,CB3-482
10 UCC:PACS2010:91.40.-k
11 UCC:PACS2010:91.65.-n,91.
    
```

Figure 5. Classification and concordances excerpt of a simple object instance (knowledge resources collection).

The listing (Figure 6) excerpts classification and concordances of a (volcanological features) container (Figure 3).

```

1 UCC:UDC2012:551.21
2 UCC:UDC2012:551
3 UCC:UDC2012:551.2,551.23,551.24,551.26
4 UCC:UDC2012:902/908
5 UCC:MSC2010:86,86A17,86A60
6 UCC:LCC:QE521-545
7 UCC:LCC:QE1-996.5
8 UCC:LCC:QC801-809
9 UCC:LCC:CC1-960,CB3-482
10 UCC:PACS2010:91.40.-k
11 UCC:PACS2010:91.65.-n,91.
12 UCC:PACS2010:91.40.Ge,91.40.St,91.40.Rs,*91.45.C-,*91.45.
    D-,90
13 ...
    
```

Figure 6. Classification and concordances excerpt of a simple container instance (knowledge resources container).

The differences in classification and concordances are resulting from the different level of detail in the collections and containers as well as in different potential of the various classification schemes to describe certain knowledge as can be seen from the different depth of classification. In integration, together the concordances can create valuable references in depth and width to complementary classification schemes and knowledge classified with different classification.

The term concordance is not only used in the simple traditional meaning. Instead, the organisation is that of a meta-concordances concept. That results from the use of universal meta-classification, which in turn is used to classify and integrate classifications [35]. The samples include simple classifications from UDC, Mathematics Subject Classification (MSC) [36], Library of Congress Classification (LCC) [37], and Physics and Astronomy Classification Scheme (PACS) [38]. For PACS the asterisk (*) indicates entries from the “Acoustics Appendix / Geophysics Appendix”.

The Universal Classified Classification (UCC) entries contain several classifications. The UCC blocks provide concordances across the classification schemes. The object classification is associated with the items associated with the object whereas the container classification is associated with the container, which means it refers to all objects in the containers.

V. KNOWLEDGE PROCESSING AND COMPUTING

The advanced processing of knowledge resources benefits from a significant number of unique attributes in its elements.

These attributes can be references, classification, keywords, textual content, links, and many more. The elements can consist of objects or collections of objects, the containers, integrating factual data with object information and structure.

Workflows for creating arbitrary result matrices (Figure 7) have been based on the organisation and object features (Figure 4) in the knowledge resources.

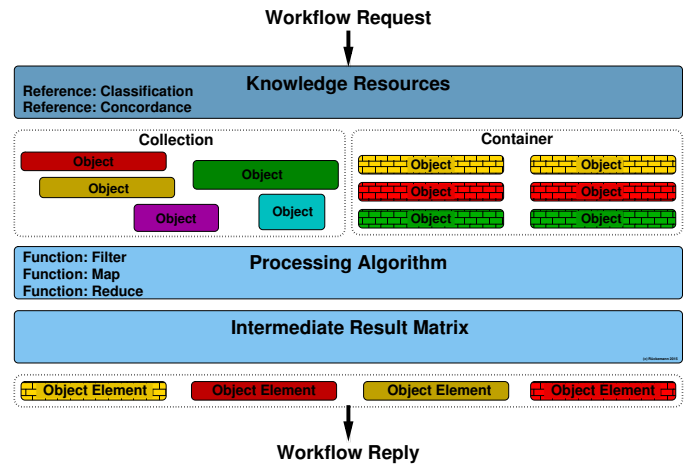


Figure 7. Creation of intermediate result matrices from resources and references (collections and containers) in reply to workflow requests.

The illustration shows that object information is gathered from the objects and references in collections and containers. Configurable algorithms like filters and mapping are then used in order to compute a result matrix. Here, the result matrix is considered “intermediate” because any of such workflows can be used in combination with other workflows, workflow chains or further processing.

For example, there is no “archaeology” in PACS, the concordances refer to resources including “archaeology” via some of the other schemes. MSC also does not contain a classification neither for volcanology or geology nor for associated features. Instead even the geophysics section classifying geological problems refers to computational methods. The above examples (Figures 5 and 6) also illustrate this. The concordances’ blocks allow to bridge between classifications and data resources, which can efficiently increase the available data pool size. Common options are in-depth computation with the container, or in-width with the general object collections. The concordances’ blocks allow to follow in-depth or in-width references within data resources, which efficiently supports to improve the quality of result matrices and the quantity of elements, which also impacts on scalability and efficiency of workflows. Table I shows the shares of items regarding processing and computing with the main steps at knowledge resources, processing algorithms, and intermediate result matrices for the Vesuvius/volcanology case (Figures 2–6).

TABLE I. PROCESSING AND COMPUTING WITHOUT (/w) AND WITH (w/) CLASSIFICATION & CONCORDANCES (VESUVIUS/VOLCANOLOGY CASE).

Items of Processing and Computing	Values	
	/w	w/
Knowledge resources		
Collection	10,000	200
Container	300	5
Processing algorithm		
Mean	750	230
String comparison	90,000	8,000
Associations	344	127
Phonetics	34	22
Weighting	296	84
Intermediate result matrix		
4 result matrix elements	120	20

The number of operations is based on subset of 100,000 collection and container objects from the knowledge resources, which have been accessed for the study. The number of items to be handled by the processing and computing for creating a comparable or higher quality result matrix have been much smaller in the major number of practical workflows when classification and concordances are included in the workflows. Especially, the primary number of requests on the collections and containers can be reduced. Consequently, the number of algorithm calls is reduced. The number of string comparisons and associated algorithms is most prominent here as the majority of objects in the resources contain text. Figure 8 shows an elementary sample workflow batch implementation of a generated caller script used for the processing parallelisation for the computing tasks, e.g., calling from Integrated System components like actmap [39].

```

1 #!/bin/bash
2 #PBS -A ruckema
3 #PBS -N PARA_Discover
4 #PBS -j oe
5 #PBS -l feature=mpp1
6 #PBS -l nodes=16:ppn=6
7 #PBS -l walltime=00:60:00
8 cd $PBS_O_WORKDIR
9 msub para_discover.sh
    
```

Figure 8. Generated workflow parallelisation with PARA_Discover.

Every instance of this sample Portable Batch System (PBS) script uses 16 compute nodes and 6 processors per node in order to execute a `para_discover` call for maximum 60 minutes walltime. A regular run with the above values requires about 5 minutes walltime per instance without and about 25 seconds with classification and concordances. With four times the nodes and cores we can handle about four times the subset data.

However, it is important to choose a right knowledge representation for universal long-term data. The Resource Description Framework (RDF) [40] is a simple example for representing Web data. In many cases, simple directed labeled graphs are not sufficient to represent knowledge. References to directed labeled or other kinds of representations should be possible.

The structure should provide an intuitive and flexible access to the data. There should be features for integrating any kind

of external data, e.g., objects, references, links, from structured to unstructured data with the available data. The elementary means of accessing the data should be independent from a certain implementation or certain purpose. The integration, interfaces, and interchange of data should be provided in most sustainable ways. This means any kind of structure and references and conceptual knowledge representation can be integrated. For example, in case of Web data even RDF can be deployed for Uniform Resource Identifiers (URI) naming relationships between data at the “ends” of a link, which in simple context enables to use graph analytics even on powerful High End Computing resources.

VI. EVALUATION

As shown, objects and containers can carry complementary information and knowledge. The classifications and concordances feature a fuzzy bridging between resources, which allows modular in-depth as well as in-width workflows. In addition to that, workflows can require strongly adaptive code and algorithms. This may result in significant variations of runtime behaviour and resources’ requirements. The workflows can integrate any objects for the processing, e.g., from collections and containers. These objects and their content may result from manual to automated origin. For example, the spectrum of creation includes use of classification, keywords, text analysis, and context analysis for the purpose of integration.

All the elements like classification, concordances, and factual data can result from manual, hybrid, and automatic processes. For example, Big Data resources can be automatically outfitted with classifications and concordances following the container components. The level of details in content, context, and structure is arbitrary and can be scaled defined by the focus of the creator of the respective data. Therefore, associated conditions can be used in workflows for weighting the types of processes and qualities involved.

In practice, during the processing and computing, the numbers of algorithm calls for requests on the collections and containers can be significantly reduced with considering classification and concordances in workflows even when creating result matrices of comparable or higher quality. There will always be non-automated resources, which might be the knowledge intensive ones. The knowledge review can also be supported by distributed authorities as well as by means of automation.

VII. CONCLUSION

The types of objects and concordances shown in this paper have been successfully created and further developed within the knowledge resources. These results have also been integrated into the knowledge resources. The workflows for creating the structures and the features for the advanced processing and computing based on these resources have been successfully implemented with in the last years. From this research, we have learned some major results.

Experiences with the creation and development of objects within the knowledge resources have resulted in the fact that the data-centric approach neither conflicts with the long-term aspects nor with the deployment of advanced processing and

computing features. This way, it should be possible to keep knowledge persistent even under changes of technology and paradigms.

The integration of objects, classification, and concordances has provided new means of documenting and accessing knowledge as well as for the efficient application of computational means. The structure of the long-term multi-disciplinary and multi-lingual knowledge resources' components enables to easily integrate objects from collections and containers. In more depth the conceptual knowledge, e.g., the classification can improve the quality of the result matrices. It enables to integrate more objects via strong means of knowledge instead of statistics or pattern matching algorithms only.

The implementation of the concordances and workflows has shown that the integrability of objects regarding multi-disciplinary and multi-lingual aspects has improved. The introduction of a universal classification and concordances is an excellent means of breaking up the isolated state of knowledge resources' content and associated data. In this context, creating concordances mainly contribute in two ways. On the one hand, concordances enable to consider different views of different and even special disciplines with the knowledge processing. On the other hand, concordances can be used to build bridges between isolated data resources.

The flexibility of the knowledge processing benefits from the advanced organisation of the data, which enables various scalable computational means for implementing directed graphs to fuzzy links, for which High End Computing resources can be deployed. Future work will be focussed on intelligent ways how classifications, concordances, and references, e.g., intermediate classifications, can be created and deployed for the use with long-term knowledge resources, classification, and workflows.

ACKNOWLEDGEMENTS

We are grateful to all national and international partners in the GEXI cooperations for their support and contributions. Special thanks go to the scientific colleagues at the Gottfried Wilhelm Leibniz Bibliothek (GWLB) Hannover, especially to Dr. Friedrich Hülsmann, for prolific discussion, inspiration, and practical case studies. Many thanks go to the scientific colleagues at the Leibniz Universität Hannover, especially to Mrs. Birgit Gersbeck-Schierholz for practical multi-disciplinary case studies and discussions. We are grateful to the Knowledge in Motion (KiM) long-term project and its members, DIMF, for the contributions on the development and application of classifications.

REFERENCES

[1] R. S. Sasser, U.S. Geological Survey library classification system. U.S. G.P.O., 1992, USGS Bulletin: 2010.

[2] E. Dodsworth and L. W. Laliberte, Eds., *Discovering and using historical geographic resources on the Web: A practical guide for Librarians*. Lanham: Rowman and Littlefield, 2014, ISBN: 0-8108-914-1.

[3] "Die Internationale Patentklassifikation (International Patent Classification, IPC)," 2014, Deutsches Patent- und Markenamt (DPMA), Germany, URL: <http://dpma.de/service/klassifikationen/ipc/> [accessed: 2015-02-01].

[4] P. Cerchiello and P. Giudici, "Non parametric statistical models for online text classification," *Advances in Data Analysis and Classification – Theory, Methods, and Applications in Data Science*, vol. 6, no. 4, 2012, pp. 277–288, special issue on "Data analysis and classification in marketing" ISSN: 1862-5347.

[5] I. Dahlberg and M. R. Schader, Eds., *Automatisierung in der Klassifikation, Proceedings, 7. Jahrestagung der Gesellschaft für Klassifikation e.V. (Teil 1), Königswinter/Rhein, Deutschland, 5.–8. April 1983*. Indeks Verlag, Frankfurt a. M., 1983, ISBN: 3-88672-012-X, URL: https://openlibrary.org/books/OL21106918M/Automatisierung_in_der_Klassifikation [accessed: 2015-02-01].

[6] W. Gaul and M. Schader, Eds., *Classification As a Tool of Research. North-Holland, Amsterdam, 1986, Proceedings, Annual Meeting of the Classification Society, (Proceedings der Fachtagung der Gesellschaft für Klassifikation)*, ISBN-13: 978-0444879806, ISBN-10: 0-444-87980-3, Hardcover, XIII, 502 p., May 1, 1986.

[7] J. Templin and L. Bradshaw, "Measuring the Reliability of Diagnostic Classification Model Examinee Estimates," *Journal of Classification*, vol. 30, no. 2, 2013, pp. 251–275, Heiser, W. J. (ed.), ISSN: 0176-4268 (print), ISSN: 1432-1343 (electronic), URL: <http://dx.doi.org/10.1007/s00357-013-9129-4> [accessed: 2015-02-01].

[8] "Multilingual Universal Decimal Classification Summary," 2012, UDC Consortium, 2012, Web resource, v. 1.1. The Hague: UDC Consortium (UDCC Publication No. 088), URL: <http://www.udcc.org/udccsummary/php/index.php> [accessed: 2015-02-01].

[9] P. Cousson, "UDC as a non-disciplinary classification system for a high-school library," in *Proc. UDC Seminar 2009, Classification at a Crossroads: Multiple Directions to Usability*, 1992, pp. 243–252, URL: http://www.academia.edu/1022257/UDC_as_a_non-disciplinary_classification_system_for_a_high-school_library [acc.: 2015-02-01].

[10] A. Adewale, "Universal Decimal Classification (UDC): A Most For All Libraries," *library 2.0, the future of libraries in the digital age*, 2014, URL: <http://www.library20.com/forum/topics/universal-decimal-classification-udc-a-most-for-all-libraries> [accessed: 2015-02-01].

[11] "EULISP Lecture Notes, European Legal Informatics Study Programme, Institute for Legal Informatics (Institut für Rechtsinformatik, IRI), Leibniz Universität Hannover," 2014, URL: <http://www.eulisp.de> [accessed: 2015-02-01].

[12] F. Heel, "Abbildungen zwischen der Dewey-Dezimalklassifikation (DDC), der Regensburger Verbundklassifikation (RVK) und der Schlagwortnormdatei (SWD) für die Recherche in heterogen erschlossenen Datenbeständen – Möglichkeiten und Problembereiche," *Bachelorarbeit im Studiengang Bibliotheks- und Informationsmanagement, Fak. Information und Kommunikation, Hochschule der Medien Stuttgart*, 2007.

[13] T. Riplinger, "Die Bedeutung der Methode Eppelsheimer für Theorie und Praxis der bibliothekarischen und der dokumentarischen Sacherschließung," *BIBLIOTHEK Forschung und Praxis*, vol. 28, no. 2, 2012, pp. 252–262, 01/2004, DOI: 10.1515/BFUP.2004.252.

[14] S. A. Keller and R. Schneider, Eds., *Wissensorganisation und -repräsentation mit digitalen Technologien*. Walter de Gruyter GmbH, 2014, *Bibliotheks- und Informationspraxis*, ISSN: 0179-0986, Band 55, ISBN: 3-11-031270-0.

[15] E. W. De Luca and I. Dahlberg, "Die Multilingual Lexical Linked Data Cloud: Eine mögliche Zugangsoptimierung?" *Information Wissenschaft & Praxis*, vol. 65, no. 4–5, 2014, pp. 279–287, Deutsche Ges. f. Information und Wissen e.V. (DGI), Ed., De Gruyter Saur, ISSN: 1434-4653, (title in English: *The Multilingual Lexical Linked Data Cloud: A possible semantic-based access to the Web?*).

[16] "Europeana," 2015, URL: <http://www.europeana.eu/> [accessed: 2015-02-01].

[17] Max Planck Institute for the History of Science, Max-Planck Institut für Wissenschaftsgeschichte, "European Cultural Heritage Online (ECHO)," 2015, Berlin, URL: <http://echo.mpiwg-berlin.mpg.de/> [accessed: 2015-02-01].

- [18] “WDL, World Digital Library,” 2015, URL: <http://www.wdl.org> [accessed: 2015-02-01].
- [19] “LX-Project,” 2015, URL: <http://www.user.uni-hannover.de/cpr/x/rprojs/en/#LX> (Information) [accessed: 2015-02-01].
- [20] C.-P. Rückemann, “Enabling Dynamical Use of Integrated Systems and Scientific Supercomputing Resources for Archaeological Information Systems,” in Proc. INFOCOMP 2012, Oct. 21–26, 2012, Venice, Italy, 2012, pp. 36–41, ISBN: 978-1-61208-226-4.
- [21] “UDC Online,” 2015, URL: <http://www.udc-hub.com/> [accessed: 2015-02-01].
- [22] “Creative Commons Attribution Share Alike 3.0 license,” 2012, URL: <http://creativecommons.org/licenses/by-sa/3.0/> [accessed: 2014-01-12].
- [23] F. Hülsmann and C.-P. Rückemann, “Summary on Algorithms and Workflows,” KiMrise, Knowledge in Motion Winter Meeting, December 12, 2014, Knowledge in Motion, Hannover, Germany, 2014.
- [24] S. Tiwari, Professional NoSQL. Wrox, John Wiley & Sons, Inc., 2011, ISBN: 978-0-470-94224-6.
- [25] “MySQL,” 2015, URL: <http://www.mysql.com/> [accessed: 2015-02-01].
- [26] “MongoDB,” 2015, URL: <http://www.mongodb.org/> [accessed: 2015-02-01].
- [27] K. Chodorow, MongoDB: The Definitive Guide. O’Reilly Media, 2013, ISBN: 9781449344689.
- [28] J. Roijackers and G. H. L. Fletcher, “On Bridging Relational and Document-Centric Data Stores,” in Proc. 29th British National Conf. on Databases (BNCOD 2013), Oxford, UK, July 8–10, 2013, Big Data, LNCS, vol. 7968, 2013, pp. 135–148, ISBN: 978-3-642-39466-9.
- [29] J. Dean and S. Ghemawat, “MapReduce: Simplified data processing on large clusters,” in Proceedings of the 6th Conference on Operating Systems Design and Implementation (OSDI 2004), 2004, DOI: 10.1.1.163.5292.
- [30] C.-P. Rückemann, “High End Computing Using Advanced Archaeology and Geoscience Objects,” International Journal On Advances in Intelligent Systems, vol. 6, no. 3&4, 2013, ISSN: 1942-2679, LCCN: 2008212456, URL: http://www.iariajournals.org/intelligent_systems/intsys_v6_n34_2013_paged.pdf [accessed: 2015-02-01].
- [31] U. Balakrishnan, “Eine DDC-RVK-Konkordanz - Erste Erkenntnisse aus dem Gebiet Medizin & Gesundheit,” 2012, Projekt Colibri / DDC Teilprojekt coli-conc, URL: <http://nbn-resolving.de/urn:nbn:de:bsz:ch1-qucosa-82838> [accessed: 2015-02-01].
- [32] I. Rauner, “Erstellung einer Konkordanz zwischen BK (Basisklassifikation) und RVK (Regensburger Verbundklassifikation) für das Fachgebiet Germanistik,” Master’s thesis, University of Vienna, Universitätslehrgang Library and Information Studies, 2010.
- [33] U. Balakrishnan, “Das Konkordanzprojekt coli-conc,” 2012, Verbundzentrale des GBV, 13. November 2013, Göttingen, Germany, URL: https://www.gbv.de/cls-download/fag-erschliessung-und-informationsvermittlung/arbeitsdokumente-fag-ei/presentation-zu-konkordanzen/at_download/file [accessed: 2015-02-01].
- [34] “North American Industry Classification System (NAICS), Concordances,” 2014, URL: <https://www.census.gov/eos/www/naics/concordances/concordances.html> [accessed: 2015-02-01].
- [35] F. Hülsmann and C.-P. Rückemann, “Classifying Classifications: Meta-classification as a New Method to Deal with Knowledge,” in Proceedings of The Fifth International Conference on Advanced Communications and Computation (INFOCOMP 2015), June 21–26, 2015, Brussels, Belgium. XPS Press, 2015, ISSN: 2308-3484, ISBN-13: 978-1-61208-416-9, (to appear).
- [36] “Mathematics Subject Classification (MSC2010),” 2010, URL: <http://msc2010.org> [accessed: 2015-02-01].
- [37] Fundamentals of Library of Congress Classification, Developed by the ALCTS/CCS-PCC Task Force on Library of Congress Classification Training, 2007, Robare, L., Arakawa, S., Frank, P., and Trumble, B. (eds.), ISBN: 0-8444-1186-8 (Instructor Manual), ISBN: 0-8444-1191-4 (Trainee Manual), URL: <http://www.loc.gov/catworkshop/courses/fundamentalslcc/pdf/classify-trnee-manual.pdf> [accessed: 2015-02-01].
- [38] “Physics and Astronomy Classification Scheme, PACS 2010 Regular Edition,” 2010, American Institute of Physics (AIP), URL: <http://www.aip.org/pacs> [accessed: 2015-02-01].
- [39] C.-P. Rückemann, “Active Map Software,” 2001, 2005, 2012, URL: <http://wwwmath.uni-muenster.de/cs/u/ruckema> (information, data, abstract) [accessed: 2012-01-01], URL: <http://www.unics.uni-hannover.de/cpr/x/rprojs/en/index.html#actmap> [accessed: 2015-01-02].
- [40] W3 Consortium (W3C), “Resource Description Framework (RDF),” 2015, W3C Semantic Web, URL: <http://www.w3.org/RDF/> [accessed: 2015-01-24].

An Investigation into Game Based Learning Using High Level Programming Languages

Ragab Ihnissi

School of Computing and Engineering
University of Huddersfield
Huddersfield, West Yorkshire, UK
e-mail: Ragab.Ihnissi@hud.ac.uk

Joan Lu

School of Computing and Engineering
University of Huddersfield
Huddersfield, West Yorkshire, UK
e-mail: J.lu@hud.ac.uk

Abstract— Game-related education within mobile learning spheres is a matter of great debate for university students across the globe. It is the case that programming languages often pose a sizeable challenge for university students. This research paper aims to develop a game based learning platform “iPlayCode”, designed to offer a new and exciting method of learning programming language. Xcode 5.0.2 was used to develop the game by using the cocos2d-x development tool and the Adobe Photoshop graphic design tool. In addition, iOS 7.0.3 (11B508) Simulator was used to test the application and the application was deployed in different models of mobile devices such as the iPhone and iPad. The application outcomes are presented by a mobile game that teaches programming languages in an easy, attractive and effective way.

Keywords—User interface; application; m-learning; mobile game based learning.

I. INTRODUCTION

Mobile technology usage has come of age to such a degree in recent years that now it has surpassed the increase of personal computers in our professional and social lives [1]. Improvements in innovative mobile and wireless technologies have also had a positive impact in our educational settings, thereby creating a new method or means for technology enhanced or improved learning called m-learning (mobile learning) [2]. Mobile technologies provide an opportunity for a fundamental change in education and due to the success of the m-learning community, recently we have noticed a rapid growth in mobile learning in all educational sectors [3]. In addition, game-based learning in m-learning environments has been a subject of interest amongst young people all over the world [4]. Nevertheless, the advancement of digital and mobile technologies has so far been restricted or confined to social communication. However, there is considerable interest in incorporating mobile learning into schools and the potential of these devices for educational use cannot be ignored. The current state of the art games that are used as an auxiliary to learning activities, are restricted to just a single programming language, and they often do not cover the whole vocabulary of this language [5]. The key objective of this project is to develop games that can be played by students who want to build upon and surpass their current level of programming language using their mobile devices. Edutainment is designed to educate, as well as to amuse by adding elements of interest to learning activities and the related contents; game-based learning is a part of edutainment [6]. In recent years, there have been many studies on the use of mobile platforms for education and these include:

A. User Interface (UI)

Bowen and Reeves [7] explored ideas for the improvement of UIs, and aimed to demonstrate ways in which more relaxed perspectives of UIs can be established, based on conventional ideas of improvement.

Several studies focus on usability but ignore the aesthetics of user interfaces. In contrast, many user interface designs emphasise usability yet reduce efficiency and effectiveness. The authors believe that the evaluation of a user interface should depend on the usability, efficiency and effectiveness.

B. Theoretical Studies

Kadyte explored the significance of beginning from the viewpoint of the mobile user, with a theoretical model for enhancing mobile platforms for education [8]. Valk et al [9] assessed support for the function of mobile education, in regards to its impact on the enhancement of learning outcomes in emerging economies.

C. Game-based M-learning

Tan and Liu [10] developed a mobile-based interactive learning environment (MOBILE) to help elementary school students with their English learning. Ab Hamid and Fung outlined a framework, built around the use of mobile gaming devices, designed to aid the comprehension of programming languages [6]. The key objective of this study is to develop games that can be played by students who want to build upon and surpass their current level of programming language using their mobile devices. The iPlayCode project is in its second version and the focus is to ensure that the correct types of games are developed and are suitable to be applied with the domain subject, which includes Objective C, Java for Android, Java, C#, C++ and Python programming languages. In addition, the developer also had to look at the user interface aspect because of the limited screen size of mobile devices. The paper is organised as follows: Section II explains the interface and game design. The implementation is discussed in Section III and the results are presented in Section IV. Section V shows the testing and evaluation of the application and Section VI is the discussion. The conclusion and future work is given in Section VII.

II. USER INTERFACE AND GAME DESIGN

This section represents an effective design of the user interface and the game layout to enhance the iPlayCode application for the users.

A. User Interface

The creators of the game need to maximise the amount of screen space dedicated to learning to compensate for the small screen size of mobile devices.

Ware [11] stated that “effective design should start with a visual task analysis; determine the set of visual queries to be supported by a design, and then the use color, form and space to efficiently serve those queries”.

The user interface must, therefore, be easy to navigate if it is to attract and engage users. Results of evaluation and analysis conducted by the researchers indicate that game creators prefer to use bright colours, including purple, blue, cream and light brown (see Figure 3) [12].

Red, black and white are also used to appeal to users and respond to questions. Bright wooden colour is employed for background and to represent the game’s levels. The colours help to emphasise the learning content of the application. To maximise usability, the authors of the game unified the background and the user interface to decrease the size and number of screens opened by the application [13]. The single screen features a tab that enables users to select any of the three levels. Once a level has been selected, another screen is displayed that comprises the question and answer part of the application.

This guarantees ease of use. The user is further stimulated by the inclusion of a tally of their score (see Figures 5, 6 and 7).

B. Game Layout

iPlayCode has three main attributes that make ‘our case games’ inspiring and fun to learn: a requisite level of challenge, use of fantasy and abstractions to make it more interesting, and triggering the curiosity of the player [14].

There are times when fantasy is missing in games, however. Where this is the case, it is preferable to create a multi-level game where factors such as interaction are the main motivation for playing. Interaction is hugely important for games used in lectures as it promotes student participation. iPlayCode also offers sound or music to motivate players to learn programming languages while playing the game. A game session in iPlayCode begins when the player registers his or her name on the start screen. When the game starts, the application displays a main screen that consists of programming language icons; the player then chooses from the programming languages on their mobile phone.

Each programming language has three levels of difficulty: level 1, level 2 and level 3. The level of the challenge in iPlayCode is adjusted by changing the difficulty level of the questions, which are subsequently randomly generated. Every level has a set of sub-functions, each containing ten questions. The questions time-out after a specified number of seconds, and the countdown time is displayed next to the question. To answer the questions, the user clicks either the right answer button or the wrong answer button on the user

interface. An additional button helps the player figure out the correct answer. Each question has a ten-second time limit, and every second represents a point. Therefore, the player needs to decide quickly whether the answer is true or false. There are two screens that display the results. The first screen shows the exposure points obtained by the player in the sub-function level. This screen includes three buttons: the first to repeat the same sub-function; the second to return to another level, and the third to move to the second screen, which displays the total points for the three levels. Gold, silver or bronze medals represent the points obtained by the player. There is also a button to return to the main screen so that the user can select another programming language. Figure 1 shows the iPlayCode design structure.




Figure 1. Game flow for iPlayCode

III. IMPLEMENTATION

A. User Interface Implementation

1) *Reduce buttons*: The results and the smaller screen size of mobile phones mean the user interface would benefit from having a limited number of buttons [13], this makes the interface simple to use and creates more space to display essential data. There are three buttons in the user interaction portion of the screen. The CCMenuItemImage class is used to decrease the number of buttons as it includes a menu image and features an upper layer of text in collaboration with CCLabelTTF,

which can be manipulated using the setString method (See Figures 6 and 7). The following code was used to execute this function.

```
Void Gameplay::nextQuestionCallback () {
    RightLabel->setString(Help); //change text
    if (questionNumber < 10) {
        WrongLabel->setString(Next); //change text
    } else {
        WrongLabel->setString(End); //change text }}
```

2) *Minimise pages:* In accordance with the research findings [13], minimal design is used on the application's screens, which includes a main screen that contains three navigation levels. This removes the necessity to create a separate screen for each level (see Figure 5 (a, b and c)).

Navigation between the three levels contains two front and rear backgrounds. The text menu is at the front and controls the contents of the page using CCMenuItemFont class. The background can also be changed to the front to show any of the three levels. The getTag() method obtains the tag of the label for the chosen level. Users can select one of three levels by clicking on the relevant button.

```
void MenuSelect::menuCallback(CCOBJECT* pSender){
    CCMenuItemImage* item =
    (CCMenuItemImage*)pSender;
    if (item->getTag() == 0) { //level1}
    else if (item->getTag() == 1) { //level2}
    } else if (item->getTag() == 2) { //level3}}
```

B. Game Implementation

To assess the quality of the application and ascertain whether the purpose for which the game was developed has been met, the application was tested for its functionality after development. The game was developed using Xcode 5.0.2 through the cocos2d-x development

tool and the Adobe Photoshop graphic design tool. The iOS 7.0.3 (11B508) Simulator was used to test the application. The application was deployed in different models of mobile devices such as the iPhone and iPad.

Although the iPlayCode application implements on mobiles, there is a user interface design challenge for devices that have a small screen. To fully utilise this application these challenges need to be addressed.

Figure 2 shows the interface scrolling up and down through question levels 2 and 3 for each language.

Initially, due to the length of the questions, it was not possible to see the entire question on the iPlayCode display. However, this has been remedied by using the code below, enabling the user to scroll up or down by clicking and dragging the question on the interface.

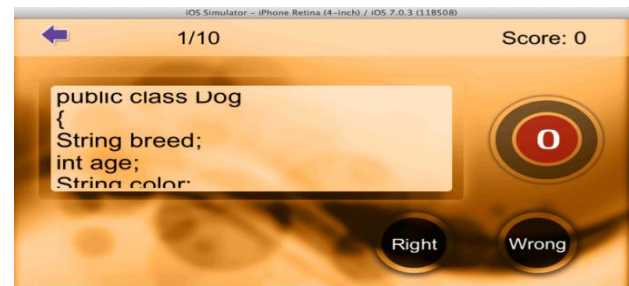


Figure 2. Scrolling down and up of the questions

```
pScrollView = CScrollView::create(CCSIZEMake(500,
230));
pScrollView->setContentSize(pQuestionLabel-
>getContentSize());
pScrollView->setContainer(pQuestionLabel);
pScrollView->setBounceable(true);
```

IV. RESULTS

The default language of the iPlayCode application is English. However, other languages such as Arabic or

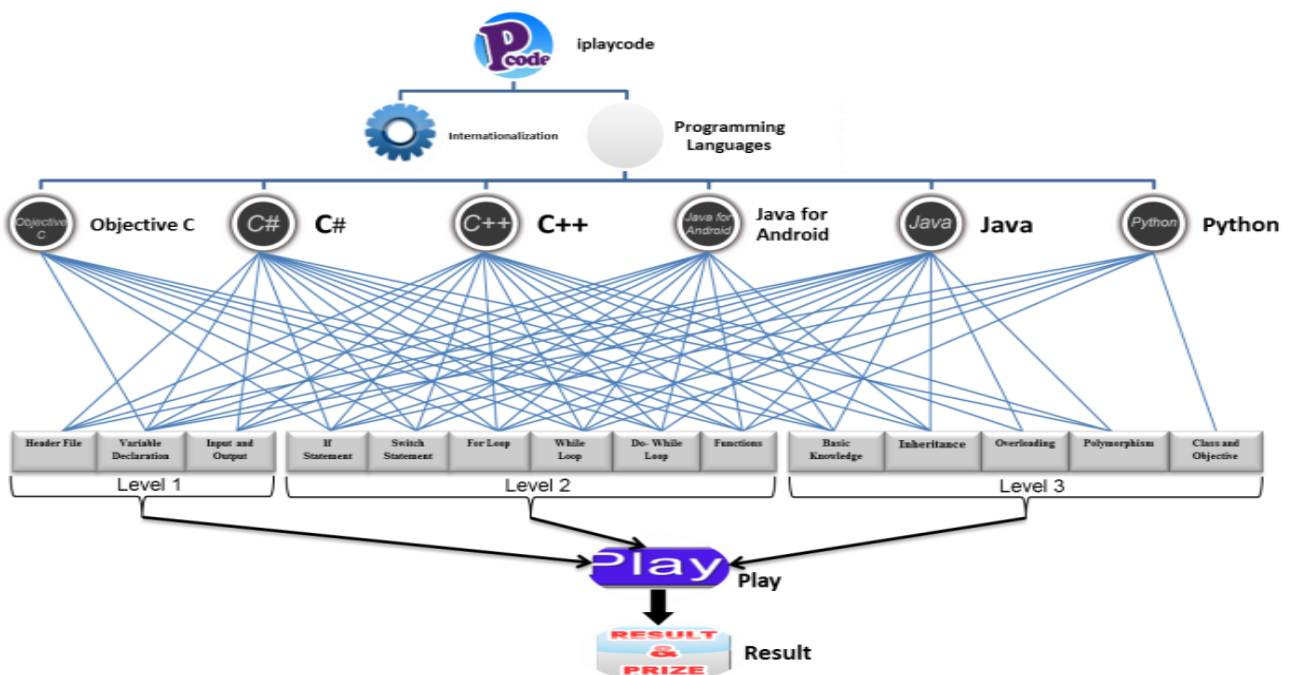


Figure 3. The architecture of iPlayCode.

Polish can be selected using the internationalisation option of the iPlayCode application. The user enters his or her name and a programming language of the user’s choice is selected. There are six programming languages in the iPlayCode application: Objective C, C#, C++, Java for Android, Java and Python (Figure 3). Each programming language has three levels but the sub-function in each level differs. The following screenshots (Figures 4 to 10) show the game interface and display for the iPlayCode application. When the game loads, the start screen is displayed. As shown in Figure 4, the players start the game by pressing the ‘Play’ button after entering their username. The user can then press the ‘Settings’ button to change the language for the interface if his or her preferred language is not English.



Figure 4. The start screen.

The interface of the iPlayCode after pressing the ‘Play’ button is shown in Figure 5. The game starts at the main game environment, and players either select the programming language they want to study by pressing one of the buttons or return to the start screen to change the username or language for the user interface.



Figure 5. The main game environment.

The different levels of the gameplay process are shown in Figure 6. Figure 6(a) is level 1 and it is made up of one to three sub-functions in each programming language. Figure 6(b) is level 2 and it consists of five to six sub-functions. Figure 6(c) is level 3 and it comprises of one to four sub-functions. In addition, each level screen displays the username, a button to return the main screen and the chosen programming language.



(a)



(b)



(c)

Figure 6. The game screens: (a) level 1, (b) level 2 and (c) level 3.

The questions screen includes the score, the number of questions, the back button to display levels, a timer and two buttons to select whether it is either the correct or the wrong answer (see Figure 7).

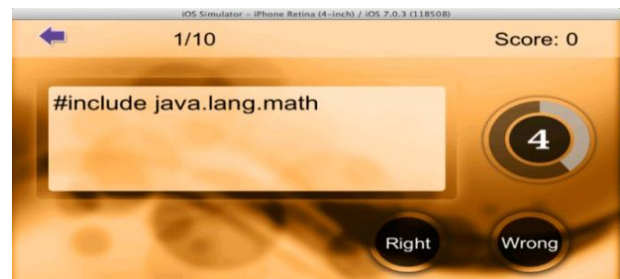


Figure 7. The questions screen

The answers screen (see Figure 8) displays the total score, the number of questions, the back button to display levels, the score for each question and an additional two buttons: one to move to the next question and the other for help to find out the correct answer.

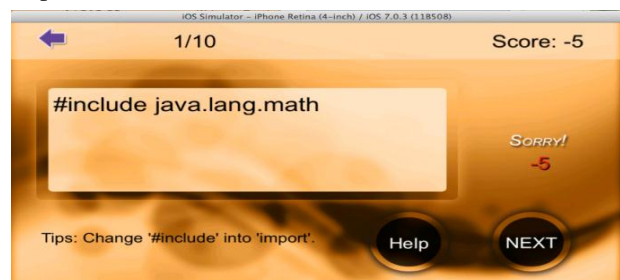


Figure 8. The answers and aid screen

The first screen to display the points for the sub-function played is shown in Figure 9. This screen also includes buttons to return to the sub-function to play again, to choose another level, to go to the final result screen, and a return button to go to the main screen to choose another programming language. In addition, the screen displays the points obtained using a medal graphic.

The type of medal is determined as follows:

- Bronze, if score>0 &&score<40

- Silver if score>39 &&score<70
- Gold if score>69 &&score<101



Figure 9. The first screen to display the points for the sub-function played.

The final results screen for the three levels displays the username (see Figure 10). It also includes the return button and displays the points obtained alongside the corresponding level and medal. The total points and medals are calculated as follows:

$$\text{Total score} = \frac{\text{Total score for all sub-functions in the level}}{\text{The number of sub-functions in the level}} \quad (1)$$

In the final result screen, the type of medal awarded to any level achieved is obtained using:

$$\frac{\text{Total score for all sub-functions in the level}}{\text{The number of sub-functions in the level}} \quad (2)$$

The award criteria are:

Bronze, if score>0 &&score<40; silver, if score>39 &&score<70; and gold, if score>69 &&score<101

For instance, in Figure 10, level 1 has a total score of 108 and the number of sub-functions is three, thereby awarding a score of 36, which falls within the bronze range. Level 2 has a total score of 345 and the number of sub-functions is six, hence awarding a score of 57.5, which falls within the silver range. Level 3 has a total score of 340 and the number of sub-functions is four. This gives a score of 85, which falls within the gold range.



Figure 10. The final result screen of the three levels

V. TESTING AND EVALUATION

Functionality testing was used to trial the game. Game testing is associated with the functionality of the game design and involves playing the game to ascertain whether the game is functioning as specified. The main aim of the functionality test is to unearth any general problems within the design or user interface. The application was tested using the iOS 7.0.3 (11B508) simulator.

Subsequently, the application was deployed in mobile devices including the iPhone and iPad. The results of the tests prove that the game is free of any defects since the application operated smoothly without any errors or

breakdowns. The game mechanics were also tested to ensure that it worked properly on several mobile devices.

To ensure consistency the test has been repeated. The results obtained provide further proof that the game functions without errors and faults. The asset of utmost importance, however, is the iPlayCode, which ensures the integrity of the application. Table (1) shows the results of the test, which confirms that devices like the iPhone and iPad are able to support the iPlayCode.

TABLE I. THE TEST THAT DEVICES LIKE IPHONE AND IPAD ARE ABLE TO SUPPORT THE IPLAYCODE

Functions	Work	Not work	Description
Login function	Yes		Successful
Menus functions	Yes		Successful
Button 'Setting' in starting screen	Yes		Successful
Button 'Play' in starting screen	Yes		Successful
Button 'Return' in all screens	Yes		Successful
Buttons 'Right, Wrong, Help, End' in answering screen	Yes		Successful
Buttons 'MENU, RESULT & PRIZE' in the first result screen	Yes		Successful
Timer in answering screen	Yes		Successful
Labels in all screens	Yes		Successful
Scroll label in answering screen	Yes		Successful
sound	Yes		Successful
Background	Yes		Successful

A detailed breakdown of the different tests is shown in table (1), which includes: functional testing, performance testing and usability testing. These results show that there are no problems with the functionality of the application.

VI. DISCUSSION

iPlayCode exists to resolve the problems caused by the conventional, rigid method of teaching utilised by tertiary institutions. To be effective, a game-based learning platform must emphasise the learning objectives and provide students with a resource that is both entertaining and supportive. This goal is fulfilled through clear and simple user interface development that helps students to learn programming languages. The user element is the core feature of the application, particularly games.

Three factors need to be considered when creating a user interface: beginning the game, basic mechanics and ease of use.

A. Beginning the Game

The start of the game is an important stage for the user. If the early stages are confusing and difficult it impedes the user's enjoyment, causes frustration and may result in the user quitting. iPlayCode begins with an enticing and clear interface that simply asks the user to specify their preferred language and enter a username (see Figure 3). The evaluation of the platform shows that it incorporates most of the essential aspects of what is deemed a 'good game'. The objectives and regulations are straightforward and this creates a product that is simple and enjoyable. It provides students with a learning tool that is relevant, practical and interactive. The interactive experience is enhanced by the iPlayCode application, which gives the students tasks that are relevant to the acquisition of valuable information.

B. Game Mechanics

Entertainment and learning are influenced by the game's mechanics. Two key features are incorporated into iPlayCode: incentives and competition.

When students interact with games, positive competition is important. iPlayCode encourages the completion of tasks by prompting the students to replay each level until they obtain a high score (see Figure 8).

This repetition ensures that the students easily recall the information and solve the problems. The students are encouraged and engaged by the penalties and rewards contained within the application. The game offers gold, silver and bronze medal achievements that are awarded when the students earn a certain number of points for each level (see Figure 9). A penalty of minus five points is given if a student answers incorrectly, does not give an answer within the allotted time, or fails to answer the question correctly before the time is up (see Figure 7).

C. Ease of Use

The enjoyment of the game is closely related to its usability. A game that is difficult to operate frustrates the user, which means they do not take pleasure in using the application. Figures 5, 6 and 7 illustrate that the simple iPlayCode user interface decreases user processes, appeals to students and guarantees fun.

VII. CONCLUSION AND FUTURE WORK

This study presents the development of the game-based learning platform iPlayCode, which offers university students a new and exciting method of learning programming languages. Many existing educational games developed for programming languages are restrictive, and they often do not cover the vocabulary of this process in its entirety. The iPlayCode platform has conquered this disadvantage by incorporating six different programming languages: Objective C, C#, C++, Java for Android, Java and Python. Moreover, the application screens are simple and easy to use. The conclusions drawn from this paper are:

- 1) A good user interface design that is easy to use makes an application attractive and interactive; as a result, users' learning needs are more effectively met.
- 2) Game-based learning creates a viable educational environment that further promotes interactive learning.
- 3) Game-based learning improves the efficiency of learning and makes it engaging and fun.

It is hoped that the aspects of the game discussed in this research paper will be implemented and tested by different users, groups and individuals, and potentially develop to include other programming languages. The long-term aim is that the platform will reach increasingly varied groups of learners. The findings gathered from this study can also be used to pinpoint even more valuable

conclusions related to the worth of mobile games and the ways in which they can be further refined.

REFERENCES

- [1] A. Herrington and J. Herrington, "Authentic mobile learning in higher education," presented at the AARE 2007 International Educational Research Conference, Fremantle, Western Australia, 2007, November 28, pp. 1-10.
- [2] M. Milrad, "How should learning activities using mobile technologies be designed to support innovative educational practices?," Big issues in mobile learning, University of Nottingham, Nottingham, 2006, pp. 27-29.
- [3] N. Winters, "What is mobile learning," Big issues in mobile learning, University of Nottingham, Nottingham, 2006, pp. 4-7.
- [4] C.-C. Chao, "An investigation of learning style differences and attitudes toward digital game-based learning among mobile users," in Wireless, Mobile and Ubiquitous Technology in Education, 2006. WMUTE'06. Fourth IEEE International Workshop on, Athens, Greece, 2006, November 16-17, ISBN: 0-7695-2723, pp. 29-31.
- [5] M. O. M. El-Hussein and J. C. Cronje, "Defining Mobile Learning in the Higher Education Landscape," Educational Technology & Society, vol. 13, pp. 12-21, 2010, October, ISSN: 1436-4522, PP. 12-21.
- [6] S. H. Ab Hamid and L. Y. Fung, "Learn programming by using mobile edutainment game approach," in Digital Game and Intelligent Toy Enhanced Learning, 2007. DIGITEL'07. The First IEEE International Workshop on, Jhongli, Taiwan 2007, March 26-28, ISBN: 0-7695-2801-5, pp. 170-172.
- [7] J. Bowen and S. Reeves, "Refinement for user interface designs," Electronic Notes in Theoretical Computer Science, Proceedings of the 2nd International Workshop on Formal Methods for Interactive Systems (FMIS 2007), Hamilton, New Zealand, vol. 208, 2008, April 14, ISSN: 1571-0661, pp. 5-22.
- [8] V. Kadyte, "Learning can happen anywhere: a mobile system for language learning," Learning with mobile devices, Learning and Skills Development Agency, London, UK, 2004, ISBN: 1-85338-833-5, pp. 73-78.
- [9] J. J.-H. Valk, A. T. Rashid, and L. Elder, "Using mobile phones to improve educational outcomes: An analysis of evidence from Asia," The International Review of Research in Open and Distance Learning, Edmonton, Canada, vol. 11, 2010, March, ISSN: 1492-3831, pp. 117-140.
- [10] T.-H. Tan and T.-Y. Liu, "The mobile-based interactive learning environment (MOBILE) and a case study for assisting elementary school English learning," in Advanced Learning Technologies (ICALT04), 2004. Proceedings. IEEE International Conference on, Washington, USA, 2004, 30 Aug.-1 Sept, ISBN: 0-7695-2181-9, pp. 530-534.
- [11] C. Ware, Visual thinking: For design: Morgan Kaufmann, 2010.
- [12] J. Zhang and J. Lu, "Using Mobile Serious Games for Learning Programming," in the proceedings of The Fourth International Conference on Advanced Communications and Computation (INFOCOMP 2014), Paris, France, 2014, July 20-24, ISBN: 978-1-61208-365-0, pp. 24-29.
- [13] R. Ihnissi and J. Lu, "An investigation into the problems of user oriented interfaces in mobile applications," presented at the The 2014 World Congress in Computer Science, Computer Engineering, and Applied Computing (WORLDCOMP'14), Las Vegas, USA, 2014, July 21-25, pp. 1-7.
- [14] T. W. Malone, "What makes things fun to learn? Heuristics for designing instructional computer games," in Proceedings of the 3rd ACM SIGSMALL symposium and the first SIGPC symposium on Small systems, New York, USA, 1980, ISBN: 0-89791-024-9 pp. 162-169.

Combining Artificial Bee Colony and Genetic Algorithms to Enhance the GPGPU-based ANN Classifier for Identifying Students with Learning Disabilities

Tung-Kuang Wu, Shian-Chang Huang,
Chin-Yu, Hsu, Chih-Han Tai

Dept. of Information Management
National ChangHua University of Education
Changhua City, Taiwan
e-mail: tkwu@im.ncue.edu.tw,
shhuang@cc.ncue.edu.tw, sw9856@gmail.com,
p198711@gmail.com

Ying-Ru Meng

Dept. of Special Education
National HsinChu University of Education
HsinChu City, Taiwan
e-mail: myr321@mail.nhcue.edu.tw

Abstract—Diagnosis of students with learning disabilities (LD) is a difficult procedure that requires extensive man power and takes a long time. Fortunately, through genetic-based (GA) parameters optimization, artificial neural network (ANN) classifier may be a good alternative to the above procedure. However, GA-based ANN model construction is computation-intensive and may take quite a while to process. In this study, we examine another optimization algorithm, the artificial bee colony (ABC) algorithm, which is based on the foraging behavior of honey bee swarm, to search for the appropriate parameters in constructing ANN-based LD classifier. We also integrate ABC algorithm with GA evolution strategy by first applying the former to derive a set of values of the ANN parameters and then use these values as the starting points for the latter GA evolution procedure. In addition, to speed-up the above process, a low-cost general purpose graphics processing unit (GPGPU), specifically, the nVidia graphics card, is adopted for the ANN model training and validation. The experimental results show that ABC can achieve better correct identification rate (CIR) than GA with less computation time. In addition, the strategy of using ABC as a pre-processing step for GA evolution has improved the correct identification rate by as much as 2.5% in two of our three data sets when compared to using GA alone.

Keywords-learning disabilities; neural network; CUDA; ABC

I. INTRODUCTION

The term “learning disabilities” (LD) was first used in 1963 [1]. However, experts in this field have not yet completely reach an agreement on the definition of LDs and its exact meaning [2]. In fact, a person can be of average or above average intelligence, without having any major sensory problems (like visual or hearing impairment), and yet struggles to keep up with people of the same age in learning and regular functioning. Due to such implicit characteristics of learning disabilities, the identification of students with LDs has long been a difficult and time-consuming process. In the United States, the so called “Discrepancy Model” [3], which states that a severe discrepancy between intellectual ability and academic achievement has to exist in one or more of these academic

areas: (1) oral expression, (2) listening comprehension (3) written expression (4) basic reading skills (5) reading comprehension (6) mathematics calculation, used to be one of the commonly adopted criteria to evaluate whether a student is eligible for special education services. However, the newer Diagnostic and Statistical Manual of Mental Disorders, Fifth Edition (DSM 5) by the American Psychiatric Association (APA) [4] has eliminated this requirement and replace it with four other criteria.

In Taiwan, the diagnosis procedure pretty much follows the “Discrepancy Model” despite the shift of criteria made by the APA. The sources of input parameters required in such prolonged process include information from parents, general education teachers, students’ academic performance and a number of standard achievement and IQ tests. To guarantee collection of required information regarding students suspected with LD, usually checklists of various aspects are developed to assist parents and regular education teachers. The Learning Characteristics Checklists (LCC), a Taiwan locally developed LD screening checklist [5], is commonly used in most counties of Taiwan. Among the standard tests, the Wechsler Intelligence Scale for Children, Third or Fourth Edition (WISC III or IV) plays the most important role in this LD diagnosis model. WISC-III consists of 13 sub tests [6]. The scores of the sub-tests are then used to derive 3 IQs, which include Full scale IQ (FIQ), Verbal IQ (VIQ), Performance IQ (PIQ), and 4 indexes, which include Verbal Comprehension Index (VCI), Perceptual Organization Index (POI), Freedom from Distractibility Index (FDI), Processing Speed Index (PSI). There are also a number of locally developed standard achievement tests (AT), which typical consist of reading, math, and fields that are related to students’ academic achievement.

Diagnosis of students with LDs then involves mainly interpreting the standard test scores and comparing them to the norms that are derived from statistical method. As an example, in case the difference between VIQ and PIQ is greater than 15, representing significant discrepancy between a student’s cultural knowledge, verbal ability, and his/her ability in recognizing familiar items, interpreting action as depicted by pictures, is a strong indicator in differentiating between students with or without LD [6]. A number of

similar indicators together with the students' academic records and descriptive data (if there is any) are then used as the basis for the final decision. Confirmed possible LD students are then evaluated for one year before admitting to special education. However, it is important to note that a previous study reveals that the certainty in predicting whether a student is having a LD using each one of the currently available predictors is in fact less than 50% [7].

The above identification procedure involves extensive manpower and resources. Furthermore, a lack of nationally regulated standard for the LD diagnosis procedure and criteria result in possible variations on the outcomes of diagnosis. In some cases, the difference can be quite significant [8].

With the advance in artificial intelligence (AI) and its successful applications to various classification problems, it is interesting to investigate how these AI-based techniques perform in identifying students with LDs. In our previous study, we have shown that ANN classifier does well in positively identifying students with LDs [8]. In subsequent studies, we combined various feature selection techniques and genetic-based parameters optimization with the ANN classifier, which further improve the overall identification accuracy [9]. However, despite the ANN-based classifier performs well in LD diagnosis problem, the procedure is computation-intensive and may take quite a while to process. Accordingly, multi-threaded programming, grid-based and cloud-based parallel computing have been used to speedup the ANN model training and validation [10][11][12].

In this paper, we are still focusing on the ANN classification model and work on porting the computation intensive ANN classifier to the general purpose graphics processing units (GPGPU). In addition, the ABC algorithm is evaluated and compared to the GA approach in terms of performance in optimizing the above mentioned ANN classifier.

The rest of the paper is organized as follows. Section 2 briefly describes the history of applying AI techniques to special education and gives a short introduction to ABC algorithm that is used in our implementation. Sections 3 and 4 present our experiment settings, design and corresponding results. Finally, Section 5 gives a brief conclusion of the paper and lists issues that deserve further investigation.

II. RELATED WORK

Artificial intelligence techniques have long been applied to special education. However, most of the studies occurred more than one or two decades ago and mainly focused on using the expert systems to assist special education in various ways [8]. There were also numerous classification techniques other than neural networks that were developed and widely used in various applications [13]. Among all the classification techniques, ANN has received lots of attentions due to its demonstrated performance and has gained wide acceptance [14].

An ANN is a mathematical representation that is inspired by the way the brain processes information. Many types of ANN models have been suggested in literature, with the most popular one for classification being the multilayer

perceptron (MLP) with back propagation. The goal of this type of network is to create a model that correctly maps the input to the output using historical data so that the model can then be used to predict the outcome when the desired output is unknown. MLP with back propagation is typically composed of an input layer, one or more hidden layers and an output layer, each consisting of several neurons. Each neuron processes its inputs and generates one output value that is transmitted to the neurons in the subsequent layer. Figure 1 provides an example of an MLP with one hidden layer and one output neuron.




Figure 1. MLP with one hidden layer.

The output of i -th hidden neuron is computed by processing the weighted inputs and its bias term b_i as follows:

$$h_i = f^h \left(b_i + \sum_{j=1}^n w_{ij} x_j \right) \quad (1)$$

where w_{ij} denotes the weight connecting input x_j to hidden unit h_i . Similarly, the result of the output layer is computed as follows:

$$y = f^{output} \left(b + \sum_{j=1}^n w_j x_j \right) \quad (2)$$

with n being the number of hidden neurons and w_j represents the weight connecting hidden unit j to the output neuron. A threshold function is then applied to map the network output y to a classification label. The transfer functions f^h and f^{output} allow the network to model non-linear relationships in the data. Also note that the number of hidden layer nodes does not need to be the same as the number of input nodes.

The training of a neural network is the process of presenting the network with sample data and modifying the weights to approximate the desired function. In particular, an epoch indicates one iteration through the process of providing the network with a sample input and updating the network's weights. Let N_i , N_h and N_o respectively represent input feature size, number of hidden and output nodes, the total order of complexity is then $O(N_i \times N_h \times N_o + N_h \times N_o)$ for one

epoch [14]. Since a typical ANN training process usually takes 500 epochs, the computation complexity for training of an ANN model is roughly equal to $500 \times N \times O(N_i \times N_h \times N_o + N_h \times N_o)$, where N represents the size of input samples for training.

In the field of special education, ANN has been used in a number of applications [8]. To improve the ANN classification accuracy, genetic-based algorithms have been used in the training and construction of ANN model [15]. A number of other approaches, such as particles swarm (PSO), ant colony (ACO) and asynchronous parallel pattern search (APPS) have also been proposed and applied to various optimization problems. In addition to the above optimization methods, a newer ABC algorithm simulating the foraging behavior of honey bees has also been proposed and shown to be performing better than GA and the other two optimization methods [16]. In this model, the colony of artificial bees consists of three different types of bees, which are employed bees, onlookers and scouts. The first half of the colony consists of employed bees and the second half includes onlookers. Each employed bee is in charge of one food source. In other words, the number of employed or on-looker bees is equal to the number of food sources. The employed bee may become a scout when its food source has been exhausted. Without regard to all the details, the procedure of ABC can be simply presented in Table I.

TABLE I. ABC PROCEDURE.

1. Generate initial n food sources $S_i, i=1, 2 \dots, n$.
2. Each employed bee $i, i=1, 2 \dots, n$, computes and memories the fitness of each food source S_i .
3. $cycle = 1$
4. repeat
4.1. For each on-looker bee $i, i=1, 2 \dots, n$, selects a new food source with roulette wheel method (with probability calculated using fitness values, $P_i = \frac{CIR(i)}{\sum_{k=1}^n CIR(k)}$) and generates its new position (v_{ij}) according to the following equation: $v_{ij} = x_{ij} + \phi_{ij} * (x_{ij} - x_{kj})$, $i, k \in \{1, 2, \dots, n\}, j \in \{1, 2, \dots, D\}$, x_{ij} is the current position of food source evaluated by employed bee i for parameter j to be optimized, k is a randomly selected number that is not equal to i , ϕ_{ij} is also a random number in the range $[-1,1]$, and D is the number of parameters to be optimized.
4.2. For each employed bees i , computes the new fitness value. If the fitness value of the food source is not improved for a continuous <i>limit</i> tries, abandons the food source. The employed bee, i , becomes a scout and randomly generates a new food source.
4.3. Memorizes the best food source.
4.4. $cycle = cycle + 1$
5. Until $cycle = MCN$

According to Table I, there are three major parameters with ABC procedure: (1) maximum cycle number (*MCN*), which is similar to the number of generations in GA, (2) the number of *food source* (also, the number of employed as well as on-looker bees), which resembles the population size in GA, (3) the maximal continuous exploitation attempts to a food source without improvement, *limit*, before an employed bee becomes a scout. This parameter may prevent ABC from trapping into some local maximum, which is somewhat similar to mutation mechanism in GA.

However, optimization procedures such as those mentioned above (like GA or ABC) usually require numerous applications of the ANN training and validation processes (depending on the number of chromosomes / food sources and evolution generations / number of cycles), and thus usually takes quite a long time to process. Accordingly, researches have been applying parallel processing, which may provide affordable computational power, to speedup the time-consuming process. For network connected cluster or grid environment, message passing interface (MPI) is usually used to coordinate computing nodes for completing a common task. On the other hand, to take full advantage of the currently available multi-core processor technology, OpenMP may be used explicitly to direct multi-threaded, shared memory parallelism. In addition to nodes or CPU-level parallelism, the fast advancing GPU technology now finds its way to all kinds of applications that require computation power [17]. Although originally designed for 3D graphics application, the Compute Unified Device Architecture (CUDA) by nVIDIA has made the low-cost and general purpose use of GPU possible [17].

In this study, we will work on porting the ABC algorithm to the enhancement of the ANN classifier for LD identification, and use GPU to speed-up the ANN training and validation processes.

III. HARDWARE PLATFORM & IMPLEMENTATION ISSUES

The data sets used in this study are summarized in Table II, which together with the corresponding pre-processing (such as normalization and feature selection) are exactly the same as those used in [10][11][12].

TABLE II. DATA SETS AND THEIR FEATURE SIZE USED IN THIS STUDY

	sample size	number of features
data set 1	652	7
data set 2	125	7
data set 3	159	10

A workstation running Ubuntu 12.04 with hardware specifications listed in Table III is set up for the above objectives. The communication between CPU and GPU is accomplished through a PCI express bus with theoretical maximum 16GB/sec bandwidth (Although the GPU supports higher PCIe 3.0 standard, the main board we use supports only up to PCIe 2.0).

TABLE III. HARDWARE DETAILS USED IN THIS STUDY

	Processing Unit	No. of cores	Memory
PC	Intel Xeon Processor E3-1230 v2 @ 3.30GHz	4 physical cores	16 GB
GPU	nVIDIA GeForce GTX 660 @ 1.03GHz	960 physical cores (5 stream multi-processors, each with 192 stream processors)	2 GB

The implementation of the ANN classifier (for LD identification) is divided into two parts: (1) the PC host is responsible for the optimization procedures (either GA or ABC), and (2) the GPU takes care of the most time-

consuming neural network model training and validation procedure, as shown in Figure 2. Switching in between the PC host and GPU, there are two I/O operations (data transfer through PCI express slot) in each iteration of GA or ABC in Figure 2. For GA, each generation may contain only two I/Os, while for ABC the number may reach six within one cycle in the worst cases. This is due to the fact that in each cycle of ABC, there may consist of fitness function computations as a result of employed, on-looker and scout bees. Considering the bandwidth between CPU and GPU, the I/O operation may potentially become the bottleneck of overall computation. However, the advantage of such a design is that it is easy to incorporate different kind of optimization algorithms in the future.

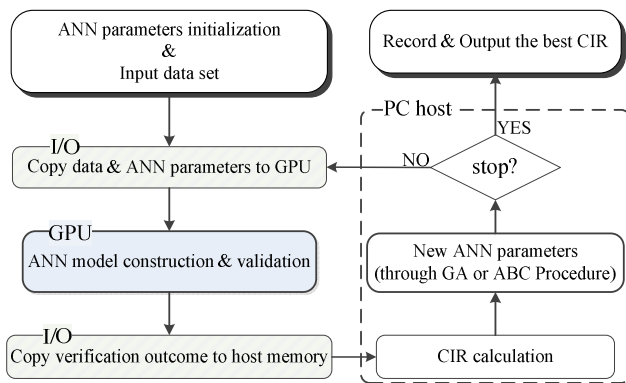


Figure 2. Computation workload distribution between the PC host and GPU.

Three parameters of an ANN classifier with one hidden layer, include number of hidden nodes (ranges in between 1 and 26), learning rate and momentum (both range in between 0.0 and 1.0), together with random number seed (which might affect the initial weights and bias of neural network) are encoded into the food sources (for ABC) or chromosomes (for GA). The training and validation of ANN model are sub-divided into six (for training) and four (for validation) independent procedures as shown in Table IV. Each of the ten procedures contains various numbers of computations that may potentially be parallelized by transforming them to the so called kernel functions.

TABLE IV. MAPPING OF ANN PROCEDURE TO GPGPU.

1. ANN Training Stage	
For (i=1 → EPOCH) {	
	Initialize ANN input vector
	Compute results of Input layer neurons
	Compute results of Hidden layer neurons
	Compute result of Output layer neuron & its error
	Adjust Output layer neurons' weights
	Adjust Hidden layer neurons' weights }
2. ANN Validation Stage	
	Initialize ANN input vector
	Compute results of Input layer neurons
	Compute results of Hidden layer neurons
	Compute result of Output layer

As an example, for a GA-optimized (with size of population equals P) N_i -fold cross-validation ANN

experiment with N_i input nodes and N_h hidden nodes, the total number of multiplications added up to $P \times N_i \times N_h \times N_h$ for the second procedure (Compute results of Input layer neurons) in the training stage of Table IV.

```

1. Determine threadID
2. Determine input vector & hidden layer weight with threadID
3. Compute the product of input vector & hidden layer weight
    
```

Figure 3. Pseudo kernel function code for computing results of Input layer neurons in the ANN training stage.

Each of the multiplications of the above example corresponds to the so called thread in CUDA's term and implemented as a kernel function as shown in Figure 3. These threads are then distributed to the 960 available cores of the GTX 660.

IV. EXPERIMENT DESIGNS AND RESULTS

The long-term goal of our series studies has been constructing an accurate ANN classifier for the identification of students with learning disabilities. The purposes of this study are to achieve the goal by evaluating ABC algorithm and exploring ways to further improve its performance. In addition, we would also like to speedup the above process with the adoption of GPGPU. The outcomes of this study will be compared to results by using the other optimization methods (such as GA).

We have designed and conducted four experiments for the purposes addressed above. In all the four experiments, CIR of ANN five-fold cross validation is used to evaluate the fitness and final outcomes of optimization (GA or ABC).

In the first experiment, we compare the performance of ABC and GA algorithms in optimizing our LD ANN classifier without the use of GPU. *MCN* and *Limit* (as explained in Table I) are set to 50 (the same as GA's number of generation) and 20 (will be examined in more details in experiment 3), while varying the number of food source. For genetic algorithm, real-value encoding is adopted with the crossover rate, mutation rate, population size and number of generation set to 0.8, 0.1, 100 and 50, respectively. Note, both ABC and GA codes are not parallelized (i.e., with OpenMP), which mean they utilize only one of the four available CPU cores. The results are shown in Table V.

TABLE V. PERFORMANCE COMPARISON BETWEEN SEQUENTIAL GA AND ABC IMPLEMENTATIONS. (ALL NUMBERS ARE AVERAGES OF 10 CONSECUTIVE EXECUTIONS OF CODE AND ALL TIME IN SECONDS)

data set \ method	1		2		3	
	CIR	execution time	CIR	execution time	CIR	execution time
GA	87.5%	9296	84.9%	3317	86.9%	6363
ABC(10*)	87.2%	1307	85.3%	713	86.5%	1026
ABC(20*)	87.6%	3390	85.6%	1261	87.0%	2126
ABC(30*)	87.6%	5399	85.8%	1751	87.2%	3188

* number of food sources.

As we can see, ABC can achieve approximately the same (or slightly better) CIR as GA algorithm (using 100 chromosomes) does with number of food sources equals 20, while take only about one third of time.

The second experiment is pretty much similar to the first one except the ANN code is now ported and executed in GPU with the number of food sources varying from 20 to 100. The results are shown in Table VI.

TABLE VI. PERFORMANCE COMPARISON BETWEEN GPU-ASSISTED GA AND ABC IMPLEMENTATIONS. (ALL NUMBERS ARE AVERAGES OF 10 CONSECUTIVE EXECUTIONS OF CODE AND ALL TIME IN SECONDS)

data set \ method	1		2		3	
	CIR	execution time	CIR	execution time	CIR	execution time
GA	87.9%	1304	85.2%	368	87.3%	641
ABC(20 [*])	87.9%	1342	85.2%	335	88.5%	480
ABC(40 [*])	88.1%	1590	86.2%	451	89.0%	722
ABC(60 [*])	88.1%	1889	86.3%	537	89.0%	894
ABC(80 [*])	88.2%	2339	86.7%	645	88.9%	1100
ABC(100 [*])	88.2%	2598	87.0%	729	89.1%	1267

* number of food sources.

The first thing to note in Table VI is that the GPU version ABC no longer has the edge in execution time as the sequential one. The major reason would be the increased number of I/Os (in each cycle of ABC) as we mentioned in Section III. We also notice that more food sources correspond to better CIR, which is not much surprise. However, the execution time does not seem to increase proportionally with the number of food sources. Apparently, more food sources mean the time spent in GPU computation is longer, which reduces the percentage of time wasted in I/O. Accordingly, using more food sources to improve the overall CIR of LD classifier should be a viable option in this case. Finally, like the sequential version, ABC with number of food sources equals 20 performs comparatively to GA in terms of CIR.

In the third experiment, we try to evaluate the effect of adjusting the *limit* parameter of ABC. As explained in Table I, *limit* is the threshold for an employed bee to escape away from some potential local maximum. We try to vary *limit* from 20 to 40 on data set 2 and record the count of final best CIR that is contributed by the scout bee (referred to as success count hereafter). For example, in Table VII the success count equals 3 when *limit* is set to 20, which means three final best results (out of ten consecutive executions of ABC optimized ANN classifier) are achieved as a result of new food sources discovered by scout bees. However, the success count dropped as *limit* increases, which indicates the effect of scout bee becomes less significant (or completely no contribution when success count drops to zero). Accordingly, it appears that setting *limit* to 20 may be a good balance in our study. In addition, with the original ABC algorithm, the scout bee operates by abandoning all the previous work and then randomly generates a new food source for a fresh start. Instead of just throwing all the efforts so far, we record the best food source (among all exploitation before reaching

limit) by the employed bee and apply the mutation strategy (as that used by the GA algorithm) on this food source to derive a new one for the scout. The results of the modified mutated strategy (using data set 2) are also presented in Table VII. As we can see, the modified mutated strategy appears to guarantee the contribution of scout bee. Also, in general the averaged CIR does show some improvement as a result of this slight modification.

TABLE VII. EFFECT OF *LIMIT* PARAMETER ON DATA SET 2 (ALL NUMBERS ARE AVERAGES OF 10 CONSECUTIVE EXECUTIONS OF CODE)

strategy \ limit	limit					
		20	25	30	35	40
random scout	CIR	87.0%	87.0%	86.9%	87.0%	87.0%
	success count*	3	1	0	0	0
mutated scout	CIR	87.3	86.9	87.3	87.2	87.0
	success count*	3	2	3	3	2

* success count indicates the number of times that the best results are derived by the scout (out of ten consecutive runs).

In the last experiment, we use ABC or GA as the preprocessing step of each other. More specifically, we perform the procedures depicted in Figure 2 twice by first adopting ABC (or GA) and preserving the optimized ANN parameters and random number seeds in the form of food sources (or chromosomes). The food sources (or chromosomes) are then used as the initial values for the following GA (or ABC) procedure. The two variations are referred to as ABC2GA (first ABC followed by GA) and GA2ABC (first GA followed by ABC), respectively. Both numbers of cycle / generation and food source / population are set to 50 and 100. In the case of ABC, *limit* is set to 20 with the mutated strategy (as used in previous experiment) adopted. The outcomes are listed in Table VIII. Also shown in Table VIII are results in previous studies [10][11][12], referred to as GA-grid1, GA-grid2 and API2APPS.

TABLE VIII. PERFORMANCE COMPARISON AMONG VARIOUS METHODS (ALL NUMBERS ARE AVERAGES OF 10 CONSECUTIVE EXECUTIONS OF CODE AND ALL TIME IN SECONDS)

data set \ method	1		2		3	
	CIR	execution time	CIR	execution time	CIR	execution time
GA	87.9%	1,304	85.2%	368	87.3%	641
ABC	<u>88.2%</u>	<u>2,598</u>	<u>87.3%</u>	<u>729</u>	<u>88.9%</u>	<u>1,267</u>
GA2ABC	88.3%	3,870	87.4%	1,092	89.3%	1,892
ABC2GA	88.4%	3,881	87.7%	1,091	89.8%	1,893
GA-Grid1	87.8%	4,931	87.5%	1,623	87.5%	2,521
GA-Grid2	-	-	87.2%	7,769	88.3%	13,234
API2APPS	88.2%	3,207	87.3%	1,285	87.6%	2,178

It appears that ABC2GA performs slightly better than its counterpart (GA2ABC). Yet, most important of all,

both perform better than methods used in all other previous implementations with the same data sets in terms of average CIR. Another matter deserved mentioning is the cost incurred by the leasing of Amazon EC2 virtual hosts for the experiments in [11] (referred to by GA-grid2) was nearly 2500 US dollars (including environment setup, code testing and final production). However, in this study, we have demonstrated that with the right methodologies (like ABC), a 200 US dollars low cost graphics card can achieve much better results both in CIR and execution time.

V. CONCLUSIONS AND FUTURE WORK

In this study, we modify the original ABC algorithm, i.e., the characteristic of scout bee, and applied this modified ABC to enhance the ANN classifier for identifying students with learning disabilities. A low cost GPU is also used to speed-up the ANN training and validation processes. As our experiments show (Table VIII), the resulted solution (ABC-optimized and GPU-assisted ANN classifier) itself not only outperforms its GA counterpart (for as many as 2.1% gain), but also in average achieve better CIR than our previous studies. Furthermore, when this modified ABC is used as the pre-processing step (for finding initial starting points) for GA evolution, the CIRs can be further improved by nearly 1% (data set 3). This is the best that we have ever got on these three data sets in terms of average CIR.

However, there are still a few things that we may be working on. The first of all would be improving the performance in terms of execution efficiency. It appears that our current implementations take too much time in the I/O operations (data transfer between CPU and GPU memory). We will try to port the optimization algorithms (GA or ABC) to GPU to reduce unnecessary I/Os in the future. On the other hand, we may also distribute the optimization code into a number of virtual nodes, which according to our past experience should further improve the CIR, and use (multiple-)GPU as the computation engine.

ACKNOWLEDGMENT

This work was supported in part by the National Science Council of Taiwan, R.O.C. under Grant NSC 103-2511-S-018-007.

REFERENCES

- [1] S. A. Kirk, "Behavioral diagnosis and remediation of learning disabilities," Proc. of the Conference on the Exploration into the Problems of the Perceptually Handicapped Child, 1963, pp.1-7.
- [2] J. M. Fletcher, W. A. Coulter, D. J. Reschly, and S. Vaughn, "Alternative approach to the definition and identification of learning disabilities: some questions and answers," *Annals of Dyslexia*, vol. 54, no. 2, 2004, pp. 304-331.
- [3] J. Schrag, "Discrepancy approaches for identifying learning disabilities," <http://www.specialed.us/discoveridea/topdocs/nasdse/discltd.pdf>, retrieved: April, 2015.
- [4] American Psychiatric Association, "Diagnostic and Statistical manual of mental disorder-V (5th ed)," Washington, DC: Book Promotion & Service LTD, 2013.
- [5] Y.-R. Meng and L.-R. Chen, "On discussing the differences about the learning characteristics of LD," *Bulletin of Special Education*, vol. 233, 2002, pp. 75-93. (in Chinese)
- [6] C. L. Nicholson and C. L. Alcorn, "Interpretation of the WISC-III and its subtests," Paper presented at the 25th Annual Meeting of the National Association of School Psychologists, Washington, DC, 1993, pp. 1-16.
- [7] T.-S. Huang, "A Study on the characteristics of WISC-III for students with learning disabilities," Master thesis, Graduate Institute of Special Education, National HsinChu University of Education, Hsinchu, Taiwan. (in Chinese)
- [8] T.-K. Wu, S.-C. Huang, and Y.-R. Meng, "Evaluation of ANN and SVM classifiers as predictors to the diagnosis of students with learning disabilities," *Expert Systems with Applications*, vol. 34, no. 3, April 2008, pp. 1846-1856.
- [9] T.-K. Wu, S.-C. Huang, and Y.-R. Meng, "Effects of feature selection on the identification of students with learning disabilities using ANN," *Lecture Notes in Computer Science*, Springer Berlin/Heidelberg, vol. 4221, 2006, pp. 565-574.
- [10] T.-K. Wu, S.-C. Huang, Y.-R. Meng, Y.-L. Lin, and H. Chang, "On the parallelization and optimization of the genetic-based ANN classifier for the diagnosis of students with learning disabilities," *Proc. 2010 IEEE Conference on Systems, Man and Cybernetics*, 2010, pp. 4263-4269.
- [11] T.-K. Wu, S.-C. Huang, Y.-R. Meng, and T.-H. Wu, "Experiences on constructing neural network based learning disabilities identification model with the Amazon elastic compute cloud," Paper presented at the 2012 International Conference on Internet Studies (NETs 2012), August 17-19, 2012, Bangkok, Thailand, pp. 1-10.
- [12] T.-K. Wu, S.-C. Huang, T.-J. Chang and Y.-R. Meng (2014), Improving learning disabilities students classification accuracy by integrating api and apps algorithms on the hadoop cloud environment, work presented at the 2014 International Conference on Internet Studies (NETs 2014), August 16-17, 2014, Singapore, pp. 1.
- [13] B. Baesens, T. Van Gestel, S. Viaene, M. Stepanova, J. Suykens, and J. Vanthienen, "Benchmarking state-of-the-art classification algorithms for credit scoring," *Journal of the Operational Research Society*, vol. 54, 2003, pp. 627-635.
- [14] C. M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, Oxford, UK, 1995.
- [15] E. Cantú-Paz and C. Kamath, "An empirical comparison of combinations of evolutionary algorithms and neural networks for classification problems," *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, vol. 35, no. 5, 2005, pp. 915-927.
- [16] D. Karaboga, "An idea based on honey bee swarm for numerical optimization," Technical Report-TR06, Erciyes University, Computer Engineering Department, Kayseri/Türkiye, Oct., 2005.
- [17] I. Buck, "GPGPU: General-purpose computation on graphics hardware-GPU computation strategies & tricks," *ACM SIGGRAPH course notes* 8, 2004.

Analysis of GLCM Parameters for Textures Classification on UMD Database Images

Alsadegh Saleh Saied Mohamed

Computing and Engineering, Huddersfield University
Huddersfield, UK

Email: alsadegh.mohamed@hud.ac.uk

Joan Lu

Computing and Engineering, Huddersfield University
Huddersfield, UK

Email: j.lu@hud.ac.uk

Abstract— Texture analysis is one of the most important techniques that have been used in image processing for many purposes, including image classification. The texture determines the region of a given gray level image, and reflects its relevant information. Several methods of analysis have been invented and developed to deal with texture in recent years, and each one has its own method of extracting features from the texture. These methods can be divided into two main approaches: statistical methods and processing methods. Gray Level Co-occurrence Matrix (GLCM) is the most popular statistical method used to get features from the texture. In addition to GLCM, a number of equations of Haralick characteristics will be used to calculate values used as discriminate features among different images in this study. There are many parameters of GLCM that should be taken into consideration to increase the discrimination between images belonging to different classes. In this study, we aim to evaluate GLCM parameters. For three decades now, GLCM is popular method used for texture analysis. Neural network which is one of supervised methods will also be used as a classifier. And finally, the database for this study will be images prepared from UMD (University of Maryland database).

Keywords- *GLCM Parameters; Haralick Feature Extraction; Texture Classification using window size.*

I. INTRODUCTION

Analysis of the texture in images aims at finding characteristics of textures and representing them in distinctive forms to enable further processing. Extracting features from the texture by algorithm is used for various types of images from different disciplines including: medical images, analysis of aerial and satellite images, and in Remote Sensing Images [1]-[5].

There are four main types of methods that have been used for extracting features from the texture, which are: texture classification, texture segmentation, texture synthesis and shape from texture [6]. Classification of images into one of the classes, which have been prepared in advance, is the main purpose of texture classification.

Most methods used to analyse the texture may divide them into: statistical approaches and filtering based approaches. Statistical approaches such as Co-occurrence matrices and Local binary patterns [2][3][7][8] extract local features from the image depending on the spatial distribution of gray values in the particular image. Statistical methods

can be categorized into: first-order (just one pixel), second-order (couple of pixels) or higher-order (more than three pixels) statistics. Due to the complicated calculations and time involved when dealing with three or more pixels, higher-order is not commonly implemented. Filtering based approaches extract global features from the texture and examples of these are wavelet and Gabor filtering [9]-[12]. Filtering deals with the pattern of texture in the spatial frequency domain of the given image, which the energy distribution in the frequency domain identifies as a texture, and focuses on periodic patterns resulting in peaks in the spatial frequency domain.

In spatial texture analysis, the texture extraction methods analyse the spatial distribution of pixels in gray scale texture. Gray level co-occurrence matrix is a statistical method used to achieve second order statistical texture features. In addition, the texture features may be extracted from the GLCM by Haralick features, using several parameters which GLCM depends upon in its design. These features are: displacement value, orientation value, gray level range ‘quantization level’, and window size [13]-[19]. All of these features influence the accuracy of texture classification.

This section has been used to introduce the paper. The remainder of this paper is structured as follows: Section II examines works related to the study. Section III describes the general concepts of GLCM methods and Haralick features. Section IV proposes the methodology of texture classification. In section V, we give the experimental results and at the end in Section VI is the conclusion.

II. RELATED WORKS

Many techniques that have been proposed for texture classification depend upon GLCM for analysing the image as the main stage in classification.

It was established that using second order statistics to extract texture features such as Angular Second Moment, Homogeneity, Contrast, Angular Second Moment, Energy and Entropy are suitable for classification of color and high resolution images of cities and farmland which are important sources of information for the geographical sector [8]. Another method was established by using GLCM after applying window size 7×7 on the original image of Guizhou karst mountainous region taken by remote sensing

by the use of the synthetic aperture radar (SAR) process. The different characteristic values of GLCM were analysed from different directions of GLCM. This experiment concluded that by using SAR image, a texture can reflect important information of the different land entities [3].

We propose a new method for defining the direction in GLCM, which selects the main direction of the image by measuring the different directions. Characteristically, the value of the main direction in texture is calculated from an average of the three directions, where a set of characteristics which includes more information about the rotation invariance of the image are extracted [18]. On the contrary, using distinct displacements to classify any type of texture does not give good results, and the value of displacement between two pixels depends on Texel size. As a result, the new method tries to compute Texel size of texture [17].

The performance of GLCM was investigated on large database from breast lesions on ultrasound images for classification. The performance depends on a changing number of parameters such as quantization, orientations and distances [13]. Another investigation on GLCM involved testing a number of the parameters of GLCM on mapping sea ice patterns with synthetic aperture radar (SAR) for assessing which one of them had the most effect on mapping sea ice texture [14].

III. GLCM ANALYSIS METHOD AND FEATURE FUNCTION

This section explains the GLCM method and the functional features of Haralick that are used to extract features from GLCM.

A. GLCM

Texture feature can be calculated by GLCM, which is one of the most popular statistical methods used for the analysis of texture. It creates a new matrix dependent on gray level values of the original image matrix. The number of rows and columns in the origin image is equal to the gray tones of new matrix. The GLCM method gives information about the type of texture in the image from the relationship between pairs of pixels. The values inside the new matrix take two parameters into consideration. These are: distance and angle.

The gray level intensity value of two pixels with a particular spatial relationship computes the distance of GLCM. Angles determine the direction of the relationship between two pixels of the same gray-level which can be horizontal; vertical or diagonal.

GLCM determines differences between surface textures through the collection of elements around a diagonal in the matrix for example, rough and smooth surfaces will be different and can easily be classified using GLCM [20].

The dimensions of GLCM are calculated by the gray level of image. More levels give more accuracy in extracting the information from the texture, as the results increase the computation cost [14]. One of the most complicated issues with GLCM is texture size which is usually only estimated.

B. Feature Function of Haralick

Haralick texture features are a function for calculating values from GLCM [21]. They are used to discriminate between different textures in the classification of classes images to determine where they belong. Some of the Haralick texture features are more important than others and this is determined by surface texture where parameters such as second moment, contrast, entropy and correlation are mostly used.

(1) Second Moment

It evaluates the uniformity of an image and focuses on the partial characteristics of the image.

$$f = \sum_{i=0}^{Ng-1} \sum_{j=0}^{Ng-1} pd, \theta(i, j)^2 \quad (1)$$

where Ng is gray tone, i,j coordinate of function P(i,j).

(2) Contrast

It indicates the range of dissimilarity between pairs of pixels over the whole image so it reveals the clarity of the image by extracting the edge information of the objects.

$$f = \sum_{n=0}^{Ng-1} n^2 \left\{ \sum_{i=0}^{Ng-1} \sum_{j=0}^{Ng-1} pd, \theta(i, j)^2 \right\}, \text{where } n = |i - j| \quad (2)$$

(3) Entropy

It is a measure of disorderliness of intensity distribution in the image. If there were no textures in the image, the entropy value would be near to zero, and on the other hand, a bigger entropy value indicates a more complex texture.

$$f = \sum_{i=0}^{Ng-1} \sum_{j=0}^{Ng-1} pd, \theta(i, j) \log(pd, \theta(i, j)) \quad (3)$$

(4) Correlation

It reflects a definite gray value along a certain direction of extended length. The correlation value will be larger if extended or if made longer.

$$f = \sum_{i=0}^{Ng-1} \sum_{j=0}^{Ng-1} pd, \theta(i, j) \frac{(i - \mu_x)(j - \mu_y)}{\sigma_x \sigma_y} \quad (4)$$

μ_x, μ_y and σ_x, σ_y are the means and standard deviations of p_x and p_y .

IV. METHODOLOGY OF TEXTURE CLASSIFICATION

Texture classification is relevant to computer vision. The stages to be followed in classification will be as laid out in Figure 1. After applying multiple windows of different sizes

the stages of classification will be as laid out in Figure 2. Firstly, there is the preparation of the images or databases that are needed for classification which are originally 160*120 in size by dividing each original image into multiple windows of size (80*60 – 40*30 – 20*15). Secondly, we will implement GLCM using the different parameters, mentioned earlier on completed sizes of images and on specific window sizes. Thirdly, we will extract the features by Haralick, and use the average of features function when using multiple windows size for identifying sets of features that describe the visual texture of an image. Finally we will classify textures by their features using machine learning approaches such as Artificial Neural Network (ANN) [19].

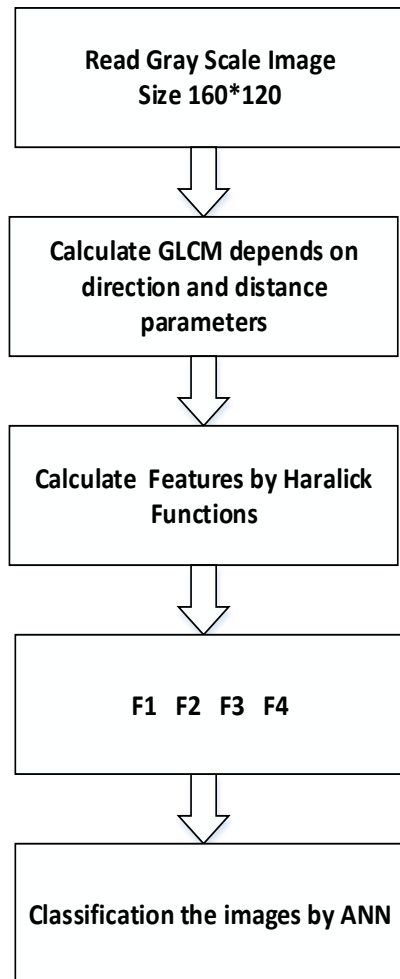


Figure 1. Classification stages on origin image size

Neural Network consists of multi-layer preceptor (MLP) algorithm. It is used in ANN to update the weights through back-propagation and training of ANN, where the neural Network is divided into two stages:

1- Training stage: using the features vector which is extracted by Haralick functions, the MLP feed forward artificial neural network

2- Testing stage: the MLP feed forward artificial neural network using other samples of data from the image, which are extracted as vector of important feature. The sampling in training is typically more than in training stage in ANN.

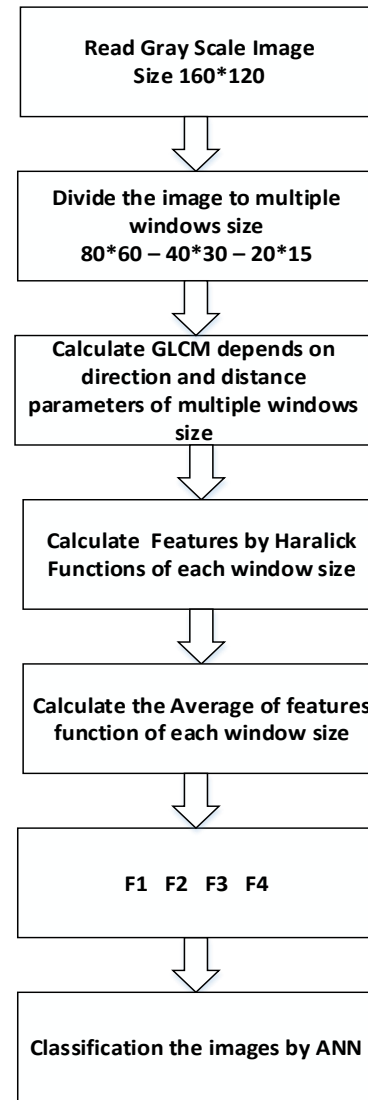


Figure 2. Classification stages on multiple window size

V. EXPERIMENTAL RESULTS

The statistical GLCM are taken as texture features of a texture. Four of the most important statistical properties are calculated for describing image content: contrast, energy, entropy, correlation, and local stationary. The main objective of the experimental investigations was to compare the discriminatory power of GLCM as a method used for analyzing the texture as the main stage for classification, as well as the influence of GLCM parameters. It also divides the texture into multi-windows through extracting the features from the image. The classification accuracy rates on the database groups were compared using MATLAB 2014.

TABLE I. CLASSIFICATION RESULTS ON THE IMAGES BY CHANGING IN DIRECTION PARAMETER WITH DIFFERENT WINDOW SIZE

Texture	angle=0 [0 1]				angle=45 [-1 1]				angle=90 [-1 0]				angle=135 [-1 -1]			
	160*120	80*60	40*30	20*15	160*120	80*60	40*30	20*15	160*120	80*60	40*30	20*15	160*120	80*60	40*30	20*15
1-A	84	86	85	85	80	80	82	87	85	84	84	83	85	84	89	87
1-B	89	88	86	87	89	91	82	84	94	93	89	87	89	90	89	88
1-C	70	75	79	76	79	79	76	80	74	72	73	71	85	84	77	81
1-D	91	92	93	91	90	91	95	88	96	97	95	97	92	92	91	94
2-A	84	81	77	78	69	80	80	80	79	77	81	77	80	78	75	76
2-B	82	90	91	83	86	86	88	85	83	83	85	85	88	84	84	89
2-C	84	84	82	83	84	82	85	88	86	85	83	83	87	87	87	87
2-D	97	98	97	97	97	98	97	97	99	99	99	98	97	97	97	98
3-A	73	77	77	75	75	75	78	80	79	72	78	76	75	78	79	80
3-B	99	99	99	98	99	99	99	99	99	99	100	99	99	99	99	99
3-C	93	94	94	92	94	94	93	93	94	94	93	95	96	94	94	94
3-D	98	98	97	97	95	94	95	98	98	98	98	98	98	98	99	99
4-A	97	90	81	84	82	85	82	85	82	82	83	82	85	84	87	87
4-B	80	83	83	83	76	72	75	77	76	82	79	78	80	82	77	78
4-C	84	70	80	79	86	87	88	91	86	79	86	78	90	81	91	90
4-D	98	97	98	97	97	98	98	97	97	98	99	98	98	98	98	98
Average	87.6	87.6	87.4	86.5	86.1	86.9	87.0	88.06	87.9	87.1	87.8	86.5	89	88.1	88.3	89.0

A. Data Preparation

The data used in the paper are from a UMD dataset [22]. The UMD (University of Maryland, College Park) consists of 25 high resolution texture classes, each one with 40 samples with resolution 1280*900 pixels (see Figure 3). Here, we divided the samples into four groups. Each group consists of another four groups: A, B, C and D.




Figure 3. Samples of texture of UMD database

B. Discussion and Analyses

We exploited the UMD database, by making a number of groups where each group consisted of two classes and each class contained 1000 texture of sizes 160 * 120 and each class was divided into 700 images for training and 300 images for testing. The results of the texture classification were as follows:

The changing in direction parameter of GLCM gives different results if applied on whole images or using multiple sizes of window. The results in Table I explain the influence of directions changing in GLCM on the classification accuracy in the number of image groups with resolution 160*120, after dividing these images to the number of regions by different window size.

In general, the GLCM has given somewhat good results in some groups especially in groups: 1-B, 1-D, 2-D, 3-B, 3-C, 3-D and 4-D that used different angels and different windows sizes. Conversely, it gave somewhat low results in classifying textures in groups: 1-C, 2-A, 3-A and 4-B. The result in rest of the groups: 1-A, 1-C, 2-C, 4-A and 4-C kept fluctuating.

There is an increase in classification accuracy through dividing the original images into sub-images, by sliding windows such as in most groups when the angle equal 45 in which include groups: 1-A, 2-A, 3-A and 4-C. When the angle equals 135, the increase was noted in groups 3-A and 3-D, and the increase was in group 2-B when the angle was 90. On the contrary, there is a decrease in other groups mostly when the angle equals 0 such as groups 2B, 3-C and 2-D, and when the angle equals 135 in groups 2-A and 4-B. In other groups, there was no apparent change such as in groups 3-B with different angles, and in groups 2-C, 3-B and 4-D when the angle equals 135. The conclusion was that each image needed its own particular window size to obtain optimal information.

The direction parameter had a big effect on GLCM. The average classification accuracy as shown in Figure 4 explains the following: when the parameter angle equals to zero we

obtain the best results between all texture classification groups and across all window sizes. Other angle values give different results depending on the texture types and window sizes used to divide the image. When the angle is equal to 135, it gives better results with the origin image as well as when using window size 80*60, whereas the accuracy percentage decreased with window size 20*15. In contrast, the classification accuracy increased significantly when the angle is equal to 45 and window size 20*15. There are fluctuations in the accuracy when the angle is equal to 90 through changing the window size from 160*120 to 20*15. In all results, the effect of angle parameter is clear on the classification accuracy between groups of textures.

A change in the distance value in GLCM leads to a change in the result of GLCM. This can be seen in Table II, which gives results about classification accuracy of images depending on different values of the distance parameter of GLCM. This was done by applying GLCM on the origin image size 160*120 and later, by dividing them to 4 regions with window size 40*30 we found that:




Figure 4. Average Classification Accuracy of Different Angle Values

TABLE II. CLASSIFICATION RESULTS ON THE IMAGES BY CHANGING IN DISTANCE PARAMETERS AND DIFFERENT WINDOW SIZE

Texture	D=2	D=4	D=6	D=8	D=12	D=16	D=20	D=24	D=2	D=4	D=6	D=8	D=12	D=16	D=20	D=24
	160*120	160*120	160*120	160*120	160*120	160*120	160*120	160*120	40*30	40*30	40*30	40*30	40*30	40*30	40*30	40*30
1-A	84	82	83	79	83	81	83	80	88	84	84	84	77	83	80	85
1-B	94	92	92	90	91	89	89	87	93	94	91	90	90	87	82	87
1-C	76	69	74	76	66	68	68	80	71	64	75	61	75	74	70	82
1-D	89	91	92	90	84	86	90	91	90	86	85	88	87	88	88	87
2-A	90	90	91	91	89	87	91	90	90	91	91	87	89	89	86	86
2-B	90	85	81	83	77	81	82	77	82	87	81	80	80	81	79	73
2-C	73	87	78	82	79	79	75	76	88	86	76	82	80	71	76	78
2-D	96	95	94	93	94	94	91	94	96	96	93	94	94	95	94	96
3-A	77	82	77	84	82	82	80	80	83	83	83	87	82	85	81	82
3-B	99	99	99	99	99	98	97	95	100	99	99	99	99	98	97	96
3-C	93	96	95	94	96	95	93	93	93	96	96	95	97	95	95	93
3-D	94	94	93	93	91	88	89	83	95	96	94	92	91	88	87	89
4-A	86	85	85	85	85	83	83	85	84	85	83	87	87	85	83	86
4-B	87	88	85	82	79	78	76	81	88	87	84	84	82	78	79	78
4-C	70	63	66	68	66	71	64	72	70	65	69	67	74	67	66	67
4-D	97	96	95	95	96	95	94	94	97	95	96	95	95	94	95	94
Avr	89	88.8	87.9	88.1	86.2	85.7	85.1	85.5	89.6	88.6	87.9	87.4	87.1	85.9	84.7	85.5

Accuracy of classifications in most groups has been affected by increase in distance in GLCM. For example there is a significant decrease in accuracy of origin image in groups 1-B, 2-C, 2-B, 3-D, and 4-D, whereas in window size 40*30 the decrease in the accuracy was in groups: 2-A, 2-B, 2-C, 3-B, 3-D and 4-C. In groups 1-A, 1-C, 1-D, 3-C and 4-C it is clear how distance caused the fluctuation of accuracy. It was noted that in the groups that have high accuracy such as 3-B, 1-D and 2-D, there was a slight decrease in accuracy.

The distance has an impact on the features which are extracted from the texture, so its influence is clear on classification of the accuracy of images.

The impact of the displacement parameter of GLCM is clear in the average results of the accuracy on the images used in classification. In general, as show in Figure 5 there was a decrease in accuracy after increase in displacement value on classification of the original image, and after dividing it to sub-images by window size.




Figure 5. Average Classification Accuracy with displacement changing

VI. CONCLUSION

Extracting features from texture is considered an effective stage in image processing and is used in many tasks. In this study, it is applied for image classification.

GLCM is one of texture feature extraction methods used in many applications. Features sets which are calculated by GLCM depend on many parameters. In this paper, we tried to evaluate clearly the impact some of these parameters have on classification of well-known database, and how the influence of these parameters differs from one to another. Distance and direction are the most important parameters of GLCM. We noticed from the experimental results, that direction has more influence than the distance, whose effects are unclear on the classification of accuracy of images. The impact of these parameters varies between the image groups. This can be observed in the fact that the impact is different based on the type of images in the group and the size of window used to divide the image. As a result, it is essential to select appropriate parameter values carefully to increase classification accuracy of each type of textures for maximizing the discrimination between images that belong to different classes. In future work, we look forward to introducing studies about the influence of other parameters such as quantization and types of features extraction functions on GLCM and finding out which of them has the most influence on the accuracy of image classification.

REFERENCES

[1] L. Nanni, A. Lumini, and S. Brahmam, "Local binary patterns variants as texture descriptors for medical image analysis," *Artificial Intelligence in Medicine*, vol. 49, 2010, pp. 117-125.

[2] C. Song, F. Yang, and P. Li, "Rotation Invariant Texture Measured by Local Binary Pattern for Remote Sensing Image Classification," *Second International Workshop on Education Technology and Computer Science*, vol. 37, 2010, pp. 3-6.

[3] J. Longhao, Z. Zhongfa, and L. Bo, "Study of SAR Image Texture Feature Extraction Based on GLCM in Guizhou Karst Mountainous Region," *Remote Sensing, Environment and Transportation Engineering*, vol.1, 2012, pp. 1-3.

[4] R. P. Ramos, M. Z. Nascimento, and D. C.r Pereira, "Texture extraction: An evaluation of ridgelet, wavelet and co-occurrence based methods applied to mammograms," *Expert Systems with Applications*, vol. 39, 2012, pp. 11036-11047.

[5] D. Fu and L. Xu, "Satellite cloud texture feature extraction based on Gabor wavelet' *International Congress on Image and Signal Processing*," 2011, pp. 248-251.

[6] M. Tuceryan, "Texture Analysis," *Department of Computer and Information Science*, 1998, pp. 207-248.

[7] Y. Zhao, W. Jia, R. Hu, and H. Min, "Completed robust local binary pattern for texture classification," *Neurocomputing*, vol. 106, 2013, pp. 68-76.

[8] M. Umaselvi, S. Kumar, and M. Athithya, "Color Based Urban and Agricultural Land Classification by GLCM Texture Features," *Third International Conference on Sustainable Energy and Intelligent System*, vol. 1, 2012, pp. 1-4.

[9] S. Arivazhagan and L. Ganesan, "Texture classification using wavelet transform," *Pattern Recognition Letters*, vol. 24, 2003, pp. 1513-1521.

[10] A. Sengur, I. Turkoglu, and M. C. Ince, "Wavelet packet neural networks for texture classification," *Expert Systems with Applications*, vol. 32, 2007, pp. 527-533.

[11] M. Idrissa and M. Acheroy, "Texture classification using Gabor filters," *Pattern Recognition Letters*, vol. 23, 2002, pp. 1095-1102.

[12] F. Riaz, A. Hassan, S. Rehman, and U. Qamar, "Texture Classification Using Rotation- and Scale-Invariant Gabor Texture Features," *Signal Processing Letters*, vol. 20, 2013, pp. 607-610.

[13] W. Gómez, W. C. A. Pereira, and A. F. C. Infantosi, "Analysis of Co-Occurrence Texture Statistics as a Function of Gray-Level Quantization for Classifying Breast Ultrasound," *Transaction on Medical Image*, vol. 31, no. 10, 2012, pp. 1889-1899.

[14] L. Soh and C. Tsatsoulis, "Texture Analysis of SAR Sea Ice Imagery Using Gray Level Co-Occurrence Matrices," *Transactions on Geoscience and Remote Sensing*, vol. 37, no. 2, March 1999, pp. 780-795.

[15] M. De Martino, F. Causa, and S. B. Serpico, "Classification of Optical High Resolution Images in Urban Environment Using Spectral and Textural Information," *Geoscience and Remote Sensing Symposium*, vol. 1, 2003, pp. 467-469.

[16] B. Wang, H. Wang, and H. Qi, "Wood Recognition Based on Grey-Level Co-Occurrence Matrix," *International Conference on Computer Application and System Modeling*, 2010, pp. 269-272.

[17] H. Nikoo, H. Talebi, and A. Mirzaei, "A Supervised Method For Determining Displacement of Gray Level Co-occurrence Matrix," *Machine Vision and Image Processing*, vol. 1, 2011, pp. 16-17.

[18] Y. Hu, C. Zhao, and H. Wang, "Directional Analysis of Texture Images Using Gray Level Co-occurrence Matrix," *Computational Intelligence and Industrial Application*, 2008, pp. 277-281.

[19] M. Arebey, M.A. Hannan, R.A. Begum, and H. Basri, "Solid waste bin level detection using gray level co-occurrence matrix feature extraction approach," *Journal of Environmental Management*, vol. 104, 2012, pp. 9-18.

[20] A. Al-Janobi, "Performance evaluation of cross-diagonal texture matrix method of texture analysis," *Pattern Recognition*, vol. 34, 2001, pp.171-180.

[21] R. M. Haralick and K. Shanmugam, "Textural Features for Image Classification," *Transactions on Systems*, vol. SMC-3, 1973, pp. 610-321.

[22] S. Hossain and S. Serikawa, "Texture databases – A comprehensive survey," *Pattern Recognition Letters*, vol. 34, 2013, pp. 2007-2022.

Copyright in Multiscale Cancer Modelling

I. V. Lishchuk, M. S. Stauch

Institut für Rechtsinformatik
 Leibniz Universität Hannover
 Hannover, Germany
 e-mail: lishchuk@iri.uni-hannover.de

Abstract — *In silico* hyper-modelling is a complex process which requires interdisciplinary effort. Hyper-modelling is a challenge for scientists, and determining applicable legal protection is a challenge for lawyers. Insofar as a computer model is defined as a computer program, software copyright comes into play. At this stage several questions arise: What elements in computer modelling are copyrightable? Is the modelling work protected? What about copyright in a hyper-model design? In this paper, we seek to suggest answers to these questions. In particular, we investigate the scope of copyright, the requirements for protection, and what elements may be subject to copyright in the context of cancer modelling.

Keywords- *in silico* oncology; cancer modelling; computer models; software copyright; copyright in compilations.

I. INTRODUCTION

In silico oncology aims to improve cancer knowledge and treatment by creating reliable computer predictions. Simulation of cancer progression in space and time requires the use of multiscale cancer modelling. “A model is considered to be “multiscale” if it spans two or more different spatial scales and/or includes processes that occur at two or more temporal scales.” [1]. Multiscale is realized in silico by constructing elementary models – the ones which correspond to elementary biological processes, and relation models – the ones which reflect relations across them, into the hyper-models [2].

In the medical research domain the term hyper-model first appeared in 2008 in relation to Virtual Physiological Human [3]; and in 2011 in the context of computer science, where it was defined as “a concrete instance of an integrative model, built as the orchestration of multiple computer models that might run on different computers at different locations using different simulation stacks.” [4]. The first implementation, based on web services, was tested on biochemical models [5].

The research on multiscale cancer modeling is ongoing. In the project CHIC [6], single-scale models (from molecular to compartment models) created by different research groups are linked into integrated multiscale hyper-models. Linking and inter-play between the models is shown in Figure 1[7].

The technical research on cancer modelling is motivated by the perspective of using multiscale cancer models as a clinical tool [7]. Meanwhile, the legal research in the project seeks to investigate solutions for protecting cancer models and hyper-models in the European legal framework. The goals of legal research is to determine the type of protection

applicable to cancer models, to identify protectable elements and specify what requirements such elements must fulfil to be protected, and to filter out elements which are not protectable.

The results which we have reached so far are presented in this paper. Section II looks at the substance of modelling, and identifies protectable elements in terms of copyright law. The applicability of copyright to cancer models, modelling work and hyper-models is then assessed in Section III, IV and V, respectively. The limits of copyright and elements which are not subject to copyright are presented in Section VI, and the key conclusions of the paper summarized in Section VII.

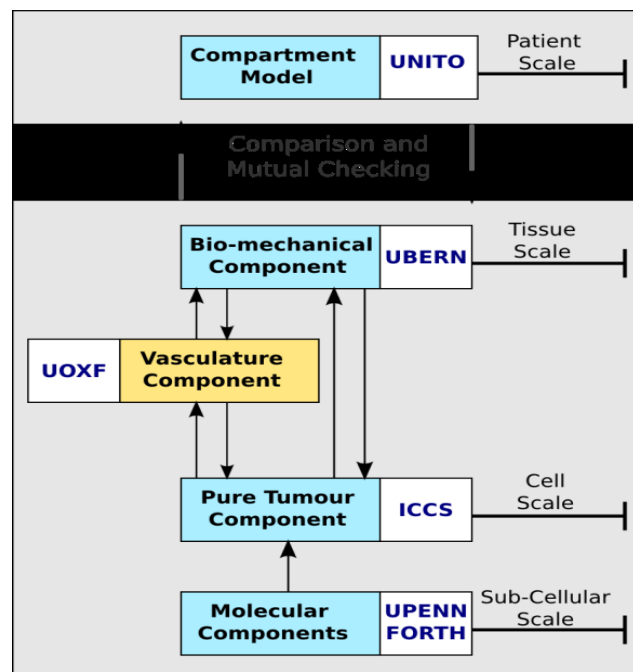


Figure 1. Schematic of the planned modelling framework for the CHIC project with the angiogenesis/vascular component highlighted.

II. CANCER MODELS IN TERMS OF COPYRIGHT

We may begin by distinguishing two types of models: scientific models and computer models. Scientific models are defined as: “finalized cognitive constructs of finite complexity that idealize an infinitely complex portion of reality through idealizations that contribute to the achievement of knowledge on that portion of reality that is objective, shareable, reliable and verifiable.” [8]. These

scientific models are implemented in silico via computer models. In the context of cancer modelling, a computer model is defined as: “*a computer program that implements a scientific model, so that when executed according to a given set of control instructions (control inputs) computes certain quantities (data outputs) on the basis of a set of initial quantities (data inputs), and a set of execution logs (control outputs).*” [9]. For implementation in silico, a scientific model is transformed into its executable form, i.e., encoded in a computer program.

When we consider that cancer models are implemented by computer models, which are written and executed by computer programs, protection of cancer models by copyright comes into question first. Computer programs are recognized as an object of copyright at the level of international law [10].

In addition, a cancer model itself emerges as the outcome of foregoing modelling work. The modelling materials which are recorded and lead to the reconstruction of a computer model for simulation may also constitute object of protection along with the model code [11].

Moreover, hyper-models, organized as choreographies, may also be an object of protection by copyright, provided the selection and arrangement of models in a hyper-model results from free intellectual creation of the modeler [12].

III. COPYRIGHT IN COMPUTER MODELS

To be protected by software copyright, a computer model must stand the criteria for protection. The requirements for copyright protection we consider below.

A. Protection of Software under Copyright

Copyright is a traditional type of protection, which both the European and the international law grant to software. Article 4 WIPO Copyright Treaty [13] and Article 10 TRIPS Agreement [10] protect computer programs as literary works within the meaning of Berne Convention (1886). The same principle is followed by European copyright law. Article 1 of Directive 2009/24/EC on the legal protection of computer programs (the “Software Directive”) recognizes computer programs as object of copyright protection in the EU [14].

B. Prerequisites for Copyright Protection

To be protected by the Software Directive, a cancer model must be “*original in the sense that it is the author's own intellectual creation.*” No other criteria for protecting programs by copyright apply [15]. As interpreted by the CJEU, a work is original when the author in its own work expresses his creativity in an original manner [16].

Hence, to qualify as a copyright work, a model code must be developed (not copied) by the modeller in an original creative way. For this, the code must be written by a human being capable of intellectual activity. Code generated automatically will lack the level of creativity necessary for the copyright. Second, a model code must count as its author's own intellectual creation. Third, the code must be expressed in a way which constitutes expression of a computer program for the purposes of software copyright.

How these criteria apply to the cancer models is described below.

C. Intellectual Creation in Computer Models

As stated, original creation is essential for copyright. However, if we consider that models are written in the era of powerful computing with the use of standard programming techniques and software development toolkits, like libraries, what counts as original creation in the model code might be questionable.

Model codes, as computer programs, are written in a programming language. The models, embodied in source code, are mostly written for interpreted languages, such as Python, Perl or MATLAB; models provided in object code are usually binary C or C++ compiled executables [9]. According to the CJEU SAS Institute decision, isolated symbols, figures, keywords, mathematical concepts, etc., which constitute the material of any programming language, do not constitute the intellectual creation of the author of a computer program and are not subject to copyright per se [17]. However, it is not these symbols that need to be protected. It is the model code which deserves protection most. The model code constitutes a script of symbols arranged into a specific sequence according to the syntax of the programming language so that a functional executable is achieved. And insofar as “*the choice, sequence and combination of those words...*” in the model code reflects original creation of the programmer such intellectual result is potentially subject to copyright protection [17].

At the same time, such protection does not mean that any copying of the code would automatically amount to copyright infringement. According to the CJEU approach, to prove copyright infringement the copying must, first, reflect intellectual expression of the author of the protected code and, second, be substantial [16]. “*Substantiality*” of the copying may be measured by different factors. Here, different jurisdictions operate varied criteria. Under the Australian case law, “*any part of a program is substantial if without it the program would either not work at all or not work as desired.*” [18]. In the US case law, substantiality is measured by the proportion of copying to the entire code. Copying 9 lines of 3,179 lines of code qualifies as “*... an innocent and inconsequential instance of copying in the context of a massive number of lines of code*” and does not count as copyright infringement [19]. Under the UK law substantiality is judged by the degree of skill and labor which the author deployed in the design and coding of the code alleged to be copied [18].

Therefore, before claiming copyright infringement substantiality of copying and intellectual creation reflected in the piece of code being copied first need to be evaluated. Another aspect which matters for the purposes of software copyright is the form of program expression.

D. Program Expression for the Purposes of Copyright

According to Article 10 TRIPS Agreement and case law of the CJEU, the program source code and object code constitute program expression for the purposes of software copyright. And, in line with the criteria established by the

CJEU, only such program expression which “permits reproduction in different computer languages, such as the source code and the object code” is subject to protection under Software Directive [20].

Hence, irrespective of whether the code is released as an executable, i.e. in the object code, or in its source code, the scope of copyright protection would be the same. By contrast, neither model codes generated by technical means, nor the models residing in the human mind or described in publications would count as program expression for the purposes of software copyright.

IV. COPYRIGHT IN PREPARATORY DESIGN MATERIAL

Copyright applicable to a model code may also extend to the preparatory design material, once such material stands the criteria for copyright.

A. Modelling Work as Preparatory Design Material

In addition to the computer model, the underlying modelling work also has the potential to be covered by software copyright. For this, the modelling work would need to qualify as preparatory design material (which is included into the object of protection by Article 1 of the Software Directive), and lead to the development of a computer model. The nature of modelling work must enable a computer model to be recreated from it at a later stage [11].

In assessing whether the modelling materials qualify as preparatory design work it is helpful to consider the modelling process by reference to the overall process of software development.

Thus, modelling of cancer models comprises various stages. These include identification of elementary processes for simulation (e.g., cell cycling, the angiogenesis process, declination of a cell to apoptosis after a particular treatment, etc.), the definition of modelling techniques - discrete, continuum, or hybrid [1], and development of computer codes for simulating those biological processes in silico.

These stages of modelling can be compared with the stages of software development. First, the problem to be solved by a computer is analyzed, then methods of solving the problem are adopted and stages of running the program are identified. Subsequently, detailed instructions for a computer to perform operations necessary for the execution of those stages are developed [21].

With this comparison in mind, insofar as the modelling materials lead to the recreation of a computer model they will have a good chance of qualifying as preparatory design material, and be subject to copyright in the model code.

In contrast to computer programs, there are no specific requirements to the form of expression of preparatory design work. However, the work must be expressed. The development documentation set down in writing, such as: data flow plans, designs of commands and information cycles, exhibits of scientific or technical art, expressed in any form, including mathematical, technical or graphic symbols, should suffice [22].

V. COPYRIGHT IN HYPER-MODELS

As noted in Section II, a further potential candidate for copyright protection in cancer modelling is the creative design by which simple models are integrated into a hyper-model. A hyper-model is defined as “a model that emerges from the composition and orchestration of multiple hypomodels, each one of which is capable of simulating a specific entity or phenomenon. The hyper-model can simulate an entity or phenomenon that may be more complex than the ones simulated by each separate simpler model.” [9].

In order to be protected, the overall design must embody intellectual creation. This is reflected in both Article 5 WIPO Copyright Treaty and Article 10 Paragraph 2 TRIPS Agreement, according to which (only): “Compilations of data or other material, in any form, which by reason of the selection or arrangement of their contents constitute intellectual creations” are subject to the law of copyright.

Although hyper-modelling is facilitated and semi-automated by technical infrastructure [23], nonetheless, an inter-play between the laws of biology and computer engineering make intellectual input in hyper-modelling indispensable. Accordingly, protection would be quite likely.

Copyright in the constituent models and copyright in the integrated hyper-model can subsist together. It allows that one and the same model may be re-assembled into different hyper-models. Any new hyper-model compiled in an original and creative way may be protected in its own right.

However, applicability of copyright to cancer models depends on multiple factors, and copyright also has its limitations. The limitations of copyright relevant to modelling are considered below.

VI. THE LIMITS OF COPYRIGHT PROTECTION

Copyright also has its limits and not all program elements are covered by copyright. The limits of copyright in relation to computer models we consider below.

A. Elements not Subject to Software Copyright

Above we established that copyright protects model codes. Nonetheless, not all aspects of the modelling will enjoy such protection. Most notably, the ideas and principles which lie behind the computer code of a cancer model are not subject to copyright. Consequently, whereas the model code as the end-product of modelling is copyrightable, the modeling background comprising techniques, algorithms, principles of computer science and bio-engineering on which a computer model rests will remain outside the boundaries of copyright. The non-copyrightability of such general modelling techniques means they can be used without restriction in future research for the benefit of the modelling community.

Non-copyrightability of ideas and principles is recognized by the law. Article 2 WIPO Copyright Treaty, Article 9 TRIPS Agreement provide: “Copyright protection shall extend to expressions and not to ideas, procedures, methods of operation or mathematical concepts as such.” Software Directive also explicitly says that ideas and

principles which underlie any element of a computer program, including its interfaces, are not protected by copyright [24]. Following this logic, also algorithms and programming languages which comprise those ideas are not copyrightable. Similarly, US copyright law also places “*ideas, plans, methods, systems, or devices, as distinguished from the particular manner in which they are expressed*” beyond the scope of copyright [25]. A code which is too abstract or straightforward and resembles an idea, rather than a functional executable, may lack the creativity to be protected by copyright in the US [26]. Whereas the computer code which implements some method or function may well be copyrightable, the method or functionality being implemented is not. In a case where the right holder tried to assert copyright in program functionality, the CJEU has adverted to the following argument to reject this: “... *to accept that the functionality of a computer program can be protected by copyright would amount to making it possible to monopolise ideas, to the detriment of technological progress and industrial development.*” [27].

On the other hand, some protection of innovative ideas, either processes or methods, may be obtainable under patent law. “*Unlike copyright, a valid patent does not protect the expression of an idea but the underlying substance of it.*” [28]. Further discussion of the scope of protection under patent law (which contains numerous complexities and uncertainties of its own) is outside the scope of the present paper.

B. Idea-Expression Dichotomy

Where, then, is the line between an idea and expression of an idea, or method of operation and implementation of the method by the code? What program elements are copyrightable expressions and what elements should remain free for the benefit of the public and for interoperability? The courts have sought to give answers to these questions by articulating various tests for segregating ideas and technical elements from copyrightable expression of a work.

Thus, the UK courts have regard to the expression - idea dichotomy and identify circumstances in which ideas “as such” are not subject to copyright [29]. For instance, if a work describes an idea, which has no connection with the literary or artistic nature of the work, such an idea is not covered by copyright in the work itself. For instance, the author of a paper describing a cancer model has a right to prevent copying of his paper, but he may not forbid others from implementing a model which he described. Secondly, commonplace ideas and principles which are used in a work, but do not constitute any substantial part of its nature are not-copyrightable [24]. In the context of computer modelling, such ideas may be common programming or modelling techniques. However, if we consider that such notions as “*judgment, skill and labour*” deployed by the author in a work matter for UK copyright, it may be argued that even if ideas as such are not copyrightable, the efforts a modeller deployed in expressing his ideas may be protected: “...*[UK law] cannot prevent the copying of a mere idea but can protect the copying of a detailed ‘idea’.* It is a question of degree where a good guide is the notion of over-

borrowing of the skill, labor and judgment which went into the copyright work.” [30].

A somewhat different test for distinguishing expression of an idea implemented by computer software from an idea per se has been adopted in the US and is known as the “*abstract-filtration-comparison*” test [31]. According to the test, the program is broken down into its structural components, the non-copyrightable elements are filtered out, and the copyrightable expression remains. For this purpose three elements precluded from copyright have been identified. First, copyright does not subsist in structures dictated by efficiency. To determine this, the court must inquire “*whether the use of this particular set of modules [is] necessary efficiently to implement that part of the program’s process being implemented.*” [31]. If the answer is positive, then the implementation of an idea by the programmer has merged with the underlying idea and is non-copyrightable. Second, copyright does not extend to elements dictated by external factors, such as standards, compatibility, program specifications, etc. Third, copyright in a program does not cover structures already found in the public domain, i.e., program elements which have been freely accessible.

The test may also be useful for filtering non-copyrightable elements in computer models. In this connection, an interesting question, at which we are looking at below, is how it may apply to an application programming interface (shortly known as “API”) – an element which constitutes part of a program, but is dictated by the requirements of interoperability.

C. Copyrightability of Program Interfaces

Interfaces are program elements by means of which a program communicates and interacts with the users and other software and hardware as it is intended to function [32]. For instance, a word processor lets a printer print a .doc file through its API.

Copyrightability of APIs has been tested under the US law. The case in question is C 10-03561 WHA, Oracle America, Inc. v. Google Inc. [19]. The case dealt with copyrightability and copyright infringement in interfaces of the programming language Java. The central question related to the “*extent to which if at all, certain replicated elements of the structure, sequence and organization of the Java application programming interface are protected by copyright*” [33].

Java is a powerful object oriented programming language, developed by Sun Microsystems, first released in 1996, and acquired by Oracle in 2010. Java has a number of pre-written programs, called “methods”, which invoke different functions, such as retrieving the cosine of an angle. These methods are grouped into “classes” and organised into “packages”. Software developers get access to those classes through the Java APIs [34]. In 2008 Java APIs had 166 “packages”, split into more than six hundred “classes”, all divided into six thousand “methods”.

Google built its Android platform on Java language and used 37 Java APIs which were the core for the smartphones. Google wrote its own implementations of the methods and classes which it needed. The only one substantial element

which Google copied from Java into Android was the names and headers of 37 API packages in question. Such copying of the headers amounted to replication of the structure, sequence and organisation of Java APIs. Oracle claimed copyright infringement, and Google defended with fair use arguing that Java is an open solution (as Oracle did not dispute) and there was no literal copying of the Java code.

Unlike code itself, the structure, sequence and organization of a program constitutes an element of non-literal expression of a program. As discussed in Section V, the protectability of such elements under the law of copyright will depend on whether the structure, sequence and organization of a program in question qualifies as an expression of an idea or an idea itself [35].

The District Court trying the case qualified the headers and method names in Java APIs as non-copyrightable. According to the interpretation criteria of the US Copyright Office, “*Even if a name, title, or short phrase is novel or distinctive or lends itself to a play on words, it cannot be protected by copyright.*” [25]. This lends support to non-protectability of isolated code items by copyright, as recognised by the CJEU [17].

As regards copying of the declarations and duplicating the command structure of Java APIs, the judge found that the command structure of Java APIs amounts to a method of operation – a material not subject to copyright in the US.

In Java programming, the specific declarations in the Java APIs designate a method. A method can be implemented in different ways, but is invoked by that specific declaration only. In Java, each symbol in a command structure is more than a simple name - each symbol carries a task to invoke a pre-assigned function:

java.package.Class.method()

Considering that for using Java class methods software developers need to replicate the Java declarations, the judge qualified the command structure of Java APIs as a method of operation – a functional element essential for interoperability not subject to the US Copyright Act. This position was based on the merger doctrine and non-copyrightability of structures dictated by efficiency: “*... When there is only one way to express an idea or function, then everyone is free to do so and no one can monopolize that expression.*” [36].

The decision of the District Court to keep Java APIs free of copyright was welcomed by the software industry as essential for interoperability. Freedom to implement or decompile program APIs would allow the developers to write compatible software and let people switch platforms and services without the need for a license [37].

However, on appeal, the Federal Circuit Court reversed that ruling. The court found the declaring code and the structure, sequence and organisation of packages in Java APIs to be protected by copyright, since Java programmers were not restrained by the factor of efficiency to arrange the declarations in a different way [38]. The implications of this appellate decision may be rather dangerous for the software industry, because copyright in APIs would deny the potential to extend program functionality by compatible apps without a license [37]. The case is now under

consideration by the US Supreme Court and the final ruling is expected soon.

As regards the copyrightability of APIs from the perspective of EU law, it is possible, following the approach of the CJEU in its BSA decision [20], to characterize API as an interactive interface which provides connectivity to a program, but which does not enable reproduction of a program in another computer language. Accordingly, a key element, which would count as a program expression in the meaning of Article 1 Software Directive, would be missing, and the API fall outside of protection. However, as long as the CJEU has not decided on copyrightability of APIs, it is a hypothetical approach only.

Copyrightability of program APIs is important for computer models as well. Such models, in common with any other computer programs, communicate with each other through APIs. In the context of CHIC, in order for the models in the project to be interoperable and communicate with each other as designed, the model interfaces are standardized. However, for expanding the cancer models beyond CHIC and for making the CHIC models interoperable with external models, the release of the models’ APIs will be required. The conditions for such release may well depend on their copyrightability and the intentions of the modelers (commercialization or release open source).

VII. CONCLUSIONS

From the above observations follows that copyright may reside in different elements of cancer models and hyper-models. First, the code of a computer cancer model may well qualify as a computer program and be protected as such. Second, copyright in the model code may extend to the modelling work, provided it is recorded in writing and leads to reproduction or subsequent recreation of the computer model. Third, hyper-models which by reason of the selection or arrangement of the models in them constitute intellectual creations may qualify as compilations and be protected as such.

At the same time, not all elements of the models and modelling efforts are protectable. Thus, copyright does not extend to ideas, principles, methods of operation, algorithms, mathematical formulae, etc. Such pieces of the modelling toolkit may not be monopolized. Further, the copyrightability of program and model interfaces remains questionable both in the EU and the US. However, in view of the open source commitment followed by the modelling community, copyrightability of interfaces plays a secondary role.

The research on cancer modelling is ongoing. Further legal research will focus on the semantic linking of in silico cancer models and applicability of copyright to hyper-models.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement No 600841.

REFERENCES

- [1] T. Deisboeck, Z. Wang, P. Macklin, V. Cristini, "Multiscale Cancer Modelling", Annual Review of Biomedical Engineering, vol. 13, pp. 127-155, August 2011, doi: 10.1146/annurev-bioeng-071910-124729.
- [2] "Computational Horizons In Cancer (CHIC): Developing Meta- and Hyper-Multiscale Models and Repositories for In Silico Oncology", CHIC, Project, Objectives, <<http://chic-vph.eu/project/objectives/>>2015.05.07.
- [3] M. Viceconti, "The skeleton of the Virtual Physiological Human", presented at: "The Living Human Project: building the musculoskeletal physiome" mini-symposium, MBEC2008 Antwerpen, Belgium, 2008.
- [4] P. A. Fishwick, "Hyperm modelling: an intergrated approach to dynamic system modelling", Journal of simulation, vol. 6, pp. 2-8, February 2012.
- [5] R. Guha, "Flexible Web Service Infrastructure,for the Development and Deployment of Predictive Models", Journal of Chemical Information and Modelling, vol. 48 (2), pp. 456-464, 2008.
- [6] "Computational Horizons In Cancer (CHIC): Developing Meta- and Hyper-Multiscale Models and Repositories for In Silico Oncology", CHIC, Project Summary, <<http://chic-vph.eu/project/>> 2015.05.07.
- [7] J. A. Grogan, P. K. Maini, J. Pitt-Francis and H. M. Byrne, "Simulating Tumour Vasculature at Multiple Scales", Proc. 2014 6th Int. Adv. Res. Workshop on In Silico Oncology and Cancer Investigation (IARWISOCI), Athens, Greece, Nov.3-4, 2014 (www.6thiarwisoci.iccs.ntua.gr), pp. 23-26, (open-access version), ISBN: 978-618-80348-1-5.
- [8] M. Viceconti, "A tentative taxonomy for predictive models in relation to their falsifiability", Philos. Transact A Math Phys Eng Sci, vol. 369, pp. 4149-61, 1954.
- [9] CHIC, Deliverable No. 7.1, "Hyperm modelling specifications", CHIC, Downloads, Deliverables, Public Deliverables, http://chic-vph.eu/uploads/media/D7-1_Hyperm modelling_specifications.pdf [accessed: 2015-05-07].
- [10] Agreement on Trade-Related Aspects of Intellectual Property Rights, the TRIPS Agreement, Annex 1C of the Marrakesh Agreement Establishing the World Trade Organization, Marrakesh, Morocco, 15 April 1994.
- [11] Article 1 Paragraph 1, Recital 7, Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs, OJEU, L 111/16, 5 May 2009.
- [12] Article 5, WIPO Copyright Treaty, Geneva, 20 December 1996; Article 10 Paragraph 2 TRIPS Agreement.
- [13] WIPO Copyright Treaty, Geneva, 20 December 1996.
- [14] Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs, OJEU, L 111/16, 5 May 2009.
- [15] Article 1 Paragraph 3, Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs, OJEU, L 111/16, 5 May 2009.
- [16] CJEU, Judgment of 16 July 2009, Case C 5/08, Infopaq International A/S v Danske Dagblades Forening, Recital 45.
- [17] CJEU, Judgment of 2 May 2012, Case C 406/10, SAS Institute Inc v World Programming Ltd.
- [18] S. Stokes, "Digital Copyright, Law and Practice", 4. Edition, Hart Publishing, Oxford and Portland, Oregon, 2014, p. 121.
- [19] U.S. District Court for the Northern District of California, Ruling of 31 May 2012, Case C 10-03561 WHA, Oracle America, Inc., v. Google Inc., p. 14, lines 1- 4.
- [20] CJEU, Judgment of 22 December 2010, Case C 393/09, Bezpečnostní softwarová asociace – Svaz softwarové ochrany v Ministerstvo kultury, Recital 35.
- [21] Report by the International Bureau, "Measures to enhance International Cooperation in the field of Legal Protection of Computer Software", WIPO expert group on the legal protection of computer software, First Session, Geneva, 27-30 November 1979, LPCS/I/2, 30 September 1979.
- [22] Federal Court of Justice of Germany, Judgment of 09 May 1985, Case I ZR 52/83, BGHZ 94, pp. 276 – 292.
- [23] D. Tartarini, K. Duan, N. Gruel, D. Testi, D. Walker, and M. Viceconti, "The VPH Hyperm modelling Framework for Cancer Multiscale Models in the Clinical Practice", Proc. 2014 6th Int. Adv. Res. Workshop on In Silico Oncology and Cancer Investigation (IARWISOCI), Athens, Greece, Nov.3-4, 2014 (www.6thiarwisoci.iccs.ntua.gr), pp. 61-64, (open-access version), ISBN: 978-618-80348-1-5.
- [24] Article 1 Paragraph 2, Recital 11, Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs, OJEU, L 111/16, 5 May 2009.
- [25] U.S. Copyright Office, Circular 34; "Copyright Protection 'Not Available for Names, Titles or Short Phrases", rev. January 2012.
- [26] R. T. Nimmer, "Legal Issues in Open Source and Free Software Distribution", adapted from Chapter 11 in Raymond T. Nimmer, The Law of Computer Technology, 1997, 2005 Supp.
- [27] CJEU, Judgment of 02 May 2012, Case C 406/10, SAS Institute Inc v World Programming Ltd., Recital 40.
- [28] A. M. St. Laurent, "Understanding open source & free software licensing", O'Reilly, 1. Edition, 2004, p. 2.
- [29] S. Stokes, "Digital Copyright, Law and Practice", 4. Edition, Hart Publishing, Oxford and Portland, Oregon, 2014, p.117.
- [30] Lord Bingham of Cornill, Opinions of the Lords of Appeal for Judgment in the case Designer Guild Ltd. v. Russel Williams (Textiles) Ltd. (Trading as Washington DC), 23 Nov. 2000, 1 WLR 2416,2422D-2423E.
- [31] U.S. Court of Appeals for the Second Circuit, Computer Associates International, Inc. v. Altai, Inc., 982 F.2d 693, 2d Cir.1992.
- [32] Recital 10, Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs, Official Journal of the European Union, L 111/16, 5 May 2009.
- [33] U.S. District Court for the Northern District of California, Ruling of 31 May 2012, Case C 10-03561 WHA, p. 1, lines 17-22.
- [34] Oracle, Java Platform, Standard Edition 6, API Specification, Overview, <<http://docs.oracle.com/javase/6/docs/api/>> 2015.05.08.
- [35] U.S. District Court for the Northern District of California, Ruling of 31 May 2012, Case C 10-03561 WHA, p. 27, lines 4-8.
- [36] U.S. District Court for the Northern District of California, Ruling of 31 May 2012, Case C 10-03561 WHA, p. 3, lines 20-21.
- [37] C. McSherry, "Dangerous Decision in Oracle v. Google: Federal Circuit Reverses Sensible Lower Court Ruling on APIs", Electronic Frontier Foundation, Deeplinks Blog, <<https://www.eff.org/deeplinks/2014/05/dangerous-ruling-oracle-v-google-federal-circuit-reverses-sensible-lower-court>> 2015.02.18.
- [38] U.S. Court of Appeals for the Federal Circuit, Ruling of 09 May 2014, Appeals from the United States District Court for the Northern District of California in No. 10-CV-3561.

Assessing Creativity: A Test for Drawing Production using Digital Art Tools

The concept, application and assessment of digital art teaching as a means of enhancing creative proficiency

Aber Aboalgasm
 School of Computing & Engineering
 University of Huddersfield
 Huddersfield, UK
 e-mail: u0874270@hud.ac.uk

Rupert Ward
 School of Computing & Engineering
 University of Huddersfield
 Huddersfield, UK
 e-mail: r.r.ward@hud.ac.uk

Abstract—This paper describes the Test for Creative Thinking - Drawing Production (TCT-DP), including its design, concept and mode of assessment, and the practical consequences of its application in a specific context. The test was used to evaluate the performance of groups of students as part of a case study exploring the use of digital art tools for drawing in a junior school. The students used specific digital art software via both computers and tablets, and also drew manually using a variety of devices. TCT-DP evaluates drawing production by means of a set of 14 criteria. At the same time, this study used the Technology Acceptance Model (TAM) theory to assess the ease of use and usefulness of the digital tools. The test was trialled with students aged 9-10 years in different ability groups. There were no significant differences in performance between male and female participants. Details of various related studies, together with data concerning the reliability and validity of the TCT-DP test, are also provided. The study finds that motivation is an important factor in improving young people’s artistic ability.

Keywords - Digital art tools; artistic ability; assessment of digital artwork; assessment of digital tools

I. INTRODUCTION

When considering the adoption of a technological approach to art, it can be argued that creative individuals should be able to develop their intellect through digital drawing activities, along with the development of their imagination. It can also be said that the visual image could become a means of diagnosing intellectual development. The research discussed here found that digital technology certainly helped to improve children’s artistic ability and enhanced their creative activity.

Some previous research has suggested that the production of art through computational technology may lessen creativity. For example, Lanier discovered that technology can lead participants to accept the lowest common denominator [1]. Another researcher, Pinsky, found that work attempted through computer-based routes was “too sanitized” and did not have the “human touch” [2]. However, although some have opposed the use of computer technology in learning for young children, Cordes and Miller argue that the outcomes for youngsters’ development as a result of the use of technology in educational settings have been widely documented and illustrate positive

feedback [3]. For example, children who use PCs have been found to show several benefits in their intellectual development and constructive knowledge, as well as in their problem solving and language abilities, in comparison with those who do not apply technology in their learning. Thus, it seems clear that the use of technology in educational environments has shown highly beneficial results.

The effects have also proved to be positive in the area of art and creativity. Matthews and Seow did a study of 12 children, aged from 2 to 11 years, using electronic paint on tablet computers [4]. The observers videotaped children drawing with both tablet computers and traditional media (pencil, markers, paint and paper) in normal surroundings. Similarities were discovered in the children’s work using both sorts of media tools, but it was also discovered that stylus-interfaced technology offered a unique tool for drawing when contrasted with the findings of earlier work conducted by Matthews et al. [5], which utilized mouse-driven electronic paint only. Tzafestas also provides an example of working with both digital and traditional tools, and using special software in the laboratory of Athens Technical University [6]. This involved an attempt to integrate traditional methods with digital drawing and painting by developing a tool called an ‘Ant Brush’. It was found that digital drawing tools could add additional factors in terms of colour and line to the user’s design. In fact, the drawing tool possessed a limited degree of autonomy. This factor proved to be helpful in improving and motivating the students’ work, and in giving the children more confidence. Additionally, Arrowhead et al. [7] found that using computers improved the motivation of elementary children in the writing process [7], whilst another researcher, Katsiaticas, demonstrates that digital tools can motivate and build artistic creativity, as well as being easy to use, particularly in circumstances when flexibility is needed with regard to source materials and techniques [8]. Haugand further adds that the provision of images and sounds to support pupils’ natural engagement in the creative process is directly concurrent to motivation [9].

This study intended to investigate whether digital tools could help children to express their ideas through drawing. It aimed to explore whether there is optimal methods that can be used to both strengthen technical skill and develop artistic imagination, and to consider how artistic creativity can be

assessed. The study involved the observation and interviewing of students in a primary school, in order to evaluate how students interacted with both digital and traditional tools.

The remaining sections of the paper are structured as follows. Section II presents the methodology used to test students in primary school in their use of digital art tools. Section III describes how the study used different approaches to gain an accurate assessment by using both the psychomotor domain taxonomy theory and Technology Acceptance Model (TAM) theory. Section IV presents the new model created to summarize the results derived through both the theories mentioned, and Section V provides examples of the outcomes achieved using the 14 criteria of TCT-DP. Section VI suggests how these 14 criteria may be used for assessing artwork, while Section VII discusses how the students became more skilful in their use of the digital tools. Section VIII evaluates the primary results and finally, Section IX provides a conclusion.

II. METHODOLOGY

The study population comprised a group of about 25 students within the age range of 9-10 years old in a primary school. The case study was conducted in the UK and tested the use of digital art tools in an ICT room. It was carried out using a drawing project finalized in collaboration with the school, which took 12 weeks. Although the project was limited in terms of size, time and setting, it did allow an in-depth study. The entire research actually involves two case studies, and while the first study has been done, the second study is still in progress. Comparison between the two cases will enable the research to present accurate results.

In the first study, the students were tested in their use of both traditional and digital tools, and the observer attempted to be as non-subjective as possible. The researcher observed the pupils' activity, took notes and recorded how the pupils dealt with the technology, as well as noting which tools were easy to use and useful. The students used two forms of art software, which were Sketchpad (see Figure 7) and Art Rage (see Figure 6). The first was utilized via PC and the second via tablet.

The idea chosen by the observer was to create a link between an ancient civilization and modern life by using a single project to combine ideas. This would allow the students to think intensely and use their creative imagination effectively. The students first used various traditional tools such as pencil, eraser, sharpener, colours and drawing paper. After this task, they worked on the same idea using digital tools via PC and tablet, but this time applying only drawing and painting tools.

A teacher assisted the students with their work. In noting how the children dealt with the technology, and which tools were easy to use, the observer considered a number of questions. For example, are the pupils using the tools correctly to complete the artwork? Do the tools appear to encourage creative self-expression? What is the motivation behind the pupils' use of some tools more than others? Is it because they are easy to use or because of their usefulness? Do the various tools stimulate and motivate children to

produce creative pieces of artwork? The children's motivation was assessed via questionnaire before the test, along with an interview after the test was completed. The usefulness of the tools can be assessed through the artwork produced, using clear assessment boundaries, as well as through continuity of activity and completion of work.

The teacher's method and skills can motivate the students to improve and build on their artistic ability. For example, the teacher can ask the student to use different tools in order to explore the function of each. The results gained from observation, interviews and questionnaire suggest that the majority of pupils (approximately 90%) felt that using digital tools increased their artistic skills. They also felt that having pictures in the art room enhanced their creativity and that their computer skills had improved. In addition, a significant majority, particularly among the boys, felt that the tools provided a lot of motivation to improve their art work. The primary results of observation show that about 60% struggled with the 'Save Picture', 'Select' and 'Lines' functions, the latter being the hardest tool of all. About 80% of student struggled with this, so pupils preferred to ignore it rather than learn how to use it. Varying numbers of pupils found the tools difficult to access and remember. All students found the computer useful, with very many finding its tools adequate, but some found the tablet the most useful. 80% of students found that they could not control some tools such as pencil or black felt. The reason for this may be that all the tools are controlled by mouse, and the mouse cannot always be as flexible and adjustable as the human hand for controlling tools.

The teacher first asked the pupils to work separately, and to draw a simple sketch by pencil and paper. Each of them did this using their own imagination. After they had finished their work, the teacher asked two or three students to combine their ideas into one artwork using digital tools. Upon completion, the observer noticed how the project had changed from a simple idea to a complex, more creative one. Specific tools helped the pupils to be creative, such as tools for Designing, Stamp and Text, Magic and the Colour tools.

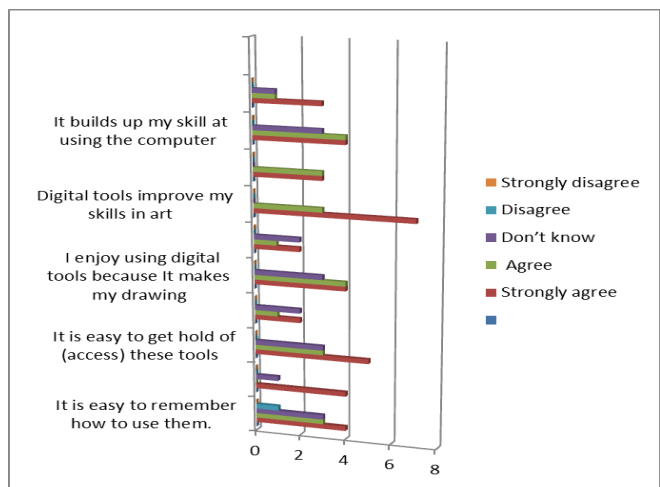


Figure 1. Summary of information from questionnaire regarding pupils' use of digital tools.

Figure 1 shows the results of the questionnaires and how students interacted with the digital tools. The majority of both male and female participants said that they liked to use technology, but a higher number of males than females said they preferred to use digital tools. PC software was the most favoured digital technology, which was preferred by 90% of the students and was more popular than the tablet and iPad art tools. The pupils said this was because they use computers more than other devices in school and also at home. Although half of the group that owned iPads/tablets found them enjoyable for artwork, 90% found that traditional tools were more effective than these digital tools.

The research can be regarded as valid and reliable because the researcher recorded and took pictures of the students working in the ICT room, as illustrated in Figure 2. Furthermore, other researchers in this field have used similar methods with young students. For example, Couse et al. [10] examined the viability of tablet computers as a medium for digital artwork with children; in particular, they explored how easily the children adapted to using them for drawing over a six week study in a classroom setting. Both quantitative and qualitative data was obtained. Ortegren (2012) also conducted a study on the subject of art in conjugation with digital media, and investigated how pupils perceived certain aspects of teaching and work methods. Ortegren’s case study focused on pupils of 7-9 years old [11], and also showed the benefits of digital image media along with multivoiced teaching. Tanir, et al. [12] conducted a further study over 14 weeks with an experimental group. Implemented within the undergraduate programme of a primary school teaching department, the study suggested that learning through visual arts was an effective method for primary school teaching.



Figure 2. Students working in ICT room using digital art tools for drawing.

III. THE ASSESSMENT OF TOOLS AND THE PSYCHOMOTOR DOMAIN

It would seem that an equal preference for both qualitative and quantitative methods can be observed in previous research, with case studies, questionnaires and observations being the most important methods for research with young children. In addition, in a study such as that described here, the researcher should be dedicated to helping children use

different approaches and equipment, and to encouraging them to express their opinions, ideas and preferences. It is important to “listen to what the children are saying; to be non-judgmental and let children ask their questions” [13]. Children have different ways of expressing themselves and therefore, different methods are needed to capture this expression. It has been found that the mixed method approach “gives importance to the children’s actions and the contexts in which they occur” [14].

As well as considering the pupils’ perceptions, this research attempted to test two theories in the art classroom, one of which was the psychomotor domain taxonomy model [15]. The psychomotor domain relates to skills that require the use of the muscles of the body which may be measured according to Bloom’s Taxonomy [15]. These include sports skills, writing skills and drawing skills. The theory describes several stages involved in learning these skills. As the psychomotor domain taxonomy can be used to evaluate the development of pupil’s skills, it can be appropriately applied to young children who are learning art and gaining creativity. The theory was used with students studying art and visual media in an ICT room, in order to determine whether the use of digital tools for drawing can help students to learn in specific stages as they improve their artistic ability and develop technological skill.

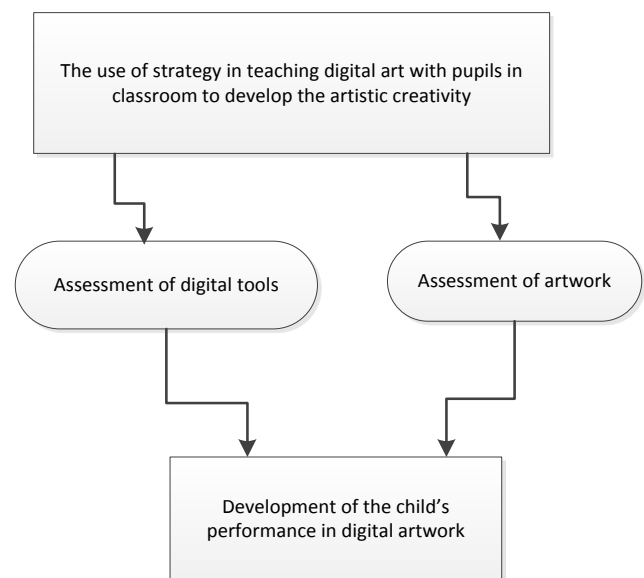


Figure 3. The process of measuring improvement in pupils’ achievement.

The present research concentrates on the stages involved in acquiring drawing skills, and how children improve over these stages. This domain was tested in relation to pupils’ drawing in class using digital tools. The students were tested on their response to a new topic they had not tried before, which had been chosen by the observer. When the principle of psychomotor domain taxonomy was applied, the test showed how students quickly learned new skills appropriate to their age. The aim was to investigate how pupils would improve their performance in drawing using digital tools, and how they would respond visually to a project involving a new idea. More details about the psychomotor domain

taxonomy model in relation to children studying art and visual media can be found in [16].

This model was then linked to a modified version of the Technology Acceptance Model (TAM) [17], which is used to assess the ease of use and usefulness of digital tools. The TAM model is an information systems theory that shows how users come to accept and use a particular technology. The child’s motivation to create good artwork enhances the desire to use digital drawing tools perfectly and to improve their performance. This gives the result that when a pupil understands how to use the digital tools, and has learned which are easy and which are difficult to use, the child will grasp the usefulness of particular digital tools. Perceived ease of use (PEOU) refers to the degree to which a person believes that using a particular system would involve making an effort [17]. Perceived usefulness and perceived ease of use can be considered cognitive factors.

The present study suggests that both theories, modified TAM model [18] and the psychomotor domain taxonomy, [19] can work in parallel to assess how students use digital art tools. Both factors can contribute to assessing students’ artwork, in order to determine the extent to which digital tools help students both artistically and technically. Figure 3 illustrates the process of measuring improvement in pupils’ achievement.

IV. NEW MODEL FOR THE ASSESSMENT OF CREATIVITY IN ART WORK

In order to summarize the results of the research and to test the theories mentioned, it was necessary to modify the TAM model for use with children in the field of digital assessment. It is very difficult for young people to assess ease of use and usefulness based purely on their own knowledge, because they still

have a lack of experience. Therefore, the researcher created an assessment model based on observation, interviews and a questionnaire. Figure 4 shows the features of a child’s creative levels in association with surrounding factors such as culture, background, financial situation, opportunities and psychological factors. If all these subjective norms in a child’s personality are considered, this can help the teacher to motivate young students to explore the ease of use and usefulness of tools, because the more the child learns, the more motivated he/she will become by new ideas. This can provide greater incentive to learn how to use new tools effectively. This in turn will allow them to achieve technical skill, to improve artistic ability and enable them to have sufficient understanding to evaluate the tools.

However, this method of evaluation is not sufficient to assess the benefits of the tools in terms of their effect on the students’ creativity. Therefore it was necessary to find an effective method of assessing and evaluating artwork, and one such method is “The Test for Creative Thinking - Drawing Production” (TCT-DP). This concept uses a set of 14 criteria to assess artistic creativity [19]. It was designed to evaluate artwork produced by traditional methods, but this study intended to use these same criteria to test visual art work produced using digital tools.




Figure 4. New prototype based on Modified TAM Model

V. EXAMPLES OF CRITERIA

Before attempting to measure the creative use of digital tools, it is important to consider what is meant by creativity and why it is important in education. There are many definitions. Traditionally, only very talented people were called ‘creative’, such as da Vinci, Einstein, Darwin, Shakespeare and others. More recently, it has been widely accepted that all individuals can be creative, and creativity does not occur only in the traditional arts and sciences. Creativity is regarded as an essential skill for any individual and for society. The National Advisory Committee on Creative and Cultural Education (NACCE), established in 1999, offers a definition of creativity which is useful for education, describing it as “imaginative activity fashioned so as to produce outcomes”. This definition is useful because it includes five characteristics of creativity: using imagination; a shaping process; having and achieving purpose; being original; and judging value. All these things, successfully taught, can result in work that has value to others, as well as helping the individual who produces the work to develop a type of mental activity that may be described as “possibility of thinking” [20].

In order to assess artistic creativity, however, certain criteria must be used. With regard to assessment of traditional art work, the conceptual deliberations of experts led to the development of a set of 14 key criteria which, as a whole, constitute the TCT-DP construct. In this study, these 14 criteria were used to assess the digital artwork, along with the tools already mentioned.

A. Assessment of Drawing Production (TCT-DP) as used in this study

Davis summarizes in his review of TCT-DP that it is “...a new, carefully developed, and possibly useful creativity test. Some researchers are impressed with its

potential for identifying creatively gifted children [21]. It is difficult to assess whether it will do more than existing American creativity tests. Efficiency of administration and scoring is a definite plus". In addition, Cropley's (1996) review maintains that "The TCT-DP is a major addition to the battery of creativity tests. It offers an approach to creativity tests that goes beyond the divergent-convergent thinking distinction [22]. It also goes some way towards incorporating non-cognitive aspects into measurement of creativity". Davis adds that "The procedure itself is interesting for the people being tested as well as for those scoring the test. The manual reviewed here is highly readable, and is also thorough, in providing not only practical instructions but also convincing theoretical and technical material justifying use of the test by both researchers and practitioners" (p.227). The test author totally agrees with Davis' final statement, that "As with limitations that plague all creativity tests, the TCT-DP should be used in conjunction with other information (e.g., another creativity test, or teachers' or parents' ratings) in order to minimize false negatives - missing creative children whose variety of creativity is different than that measured by a single test" (p.91).

The present research used the design, concept and evaluation scheme of TCT-DP, as described by previous experiences and results of its application [20]. The adoption of this approach was designed to reflect a more holistic concept of visual media, as it was used to evaluate proficiency in the use of digital art tools. More specifically, it was used to assess the ability of children of primary age to use such tools to produce creative artwork.

VI. NEW APPLICATION OF TCT-DP

The current study suggested that the application of these 14 criteria could be used to help assess artwork created using digital tools (see Figure 5), in order to investigate whether the tools could contribute to and strengthen children's artwork. This application could be used in the art classroom to evaluate enhancement of students' artistic ability and imaginative development. Two sides of the application would be required, one side for teaching and a second side for students. Each student would have their own storage area where they could save all their work.

All the applications would be linked to the internet, to enable the students to send their art work to the teacher via email, and to allow the teacher to give the students feedback and scores, again by email. Students would not be able to see each other's work or feedback. The teacher would have a list of criteria showing clearly what tools students had used in their work and which they had not, because this record of tools used would enable the teacher to assess any creative enhancement in students' work correctly. In addition, the students would know from the feedback where they had made mistakes.

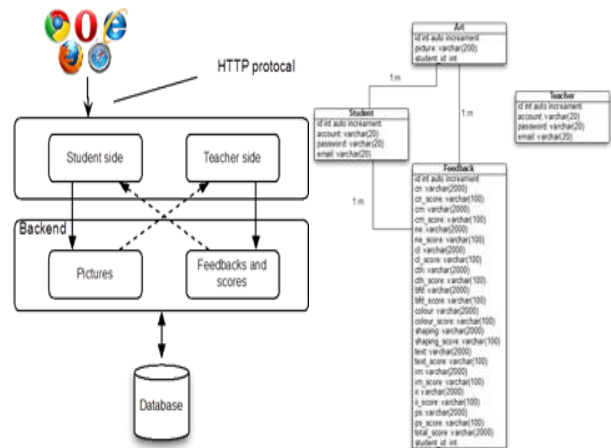


Figure 5. Application tested to ascertain its effectiveness in assessing pupils' artwork with digital tools.

This application was tested to determine whether students' artwork could be assessed, whilst at the same time evaluating the creative enhancement of their work using digital tools and the suitability of the tools for their age range.

VII. DISCUSSION

The main aim of the case study was to investigate whether digital art tools were able to motivate the children to improve their artistic ability, express their imagination and develop their technical skill. It appears that digital drawing tools can help to improve artistic creativity in children by providing ways of trying out new ideas, and new ways of thinking and problem solving. The research found that the concepts and digital tools complemented each other. The more the students used and worked with the tools, the more they were able to progress their ideas. As they became skilful at utilizing the tools, they were motivated to produce good work and learned how to be creative with them. The teacher's role is very important in stimulating students to use the tools correctly, encouraging them to apply the appropriate tools for each step, and motivating them by assessing their artwork. By the end of the project, the students had learned how to evaluate the tools correctly, after correcting mistakes many times and sometimes using inappropriate tools before finding ones that worked more effectively. The observer noted how the children's approach changed during the sessions. For example, when the students started working, their thinking was simple and realistic. Later they changed from naturalism to more abstract ideas because their thinking had improved since they learned how to use the digital tools. The students also learned from each other, especially when they were encouraged by their parents, family, the environment and the school.

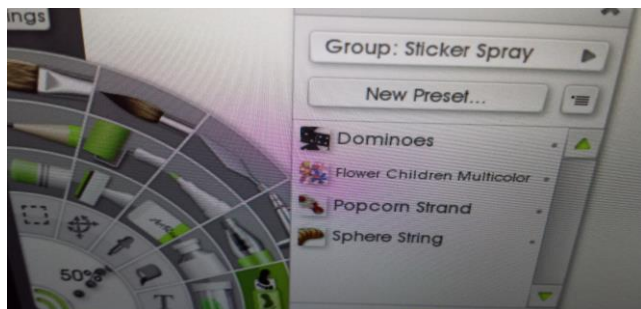


Figure 6. Students drawing with Art Rage software using a tablet.




Figure 7. Students drawing with Sketchpad software using a computer.

The research found that motivation is an essential factor in the progress of the student, as is interaction between them, for example, to give the student new ideas they had not used before. Sometimes they repeated the same work at home, creating competition among the students. All these factors helped to motivate them to interact with the tools. It can therefore be said that the research produced successful results.

VIII. EVALUATION

The researcher found that it is important to explore a range of strategies and methods of teaching art, not just because art is a practical and technical subject, but also because it is different from any other school subjects. Although the outcome was successful, it must be accepted that not everyone is interested in working with digital tools, or in creating art by digital means only. Another observation was that many of the students struggled to understand the distinction between ease of use and usefulness. They could grasp the meaning of ease of use, but not the concept of usefulness.

However, it was interesting to note that, in general, ease of use was not the main motivation when the children used digital tools. They found more ease and enjoyment, generally, when using traditional methods. The most popular digital methods were sometimes the easiest to use, but according to their own statements, it was not the ease of use, but the successful effects the tools provided, which motivated the children to use them. It seems, therefore that enjoyment and artistic satisfaction counts more, to the pupils, than ease of use. Thus, usefulness is more valuable than easiness.

IX. CONCLUSION

Based on the information from observation, interview and questionnaire, it can be said that the children worked well and felt that the digital tools had improved their creative ability. They were quite strongly motivated to use them by a wish to create good artwork, and also by the drive to improve their technological skills. There was also some understanding of the value of these tools in learning other subjects.

It seems that, in general, ease of use was not the main motivation when the children used digital tools. They found more ease and enjoyment, generally, when using digital and traditional methods. The most popular digital methods were sometimes the easiest to use, but according to their own statements, it was not the ease of use, but the successful effects, the tools provided, which motivate the children to use them. It seems, therefore that enjoyment and artistic satisfaction counts more, to the pupils, than ease of use; the usefulness is more valuable than easiness. Also, the motivation is important for young students to enhance their artistic ability and improve their creative activity. The modification TAM model showed itself to be very useful in assessing the value of traditional and digital tools.

Psychomotor domain taxonomy was also useful according to the children's progress in learning to use digital software. It helped in noting and assessing the students' progress generally. The Test for Creative Thinking - Drawing Production (TCT-DP) is still under testing. We hope this test will help us gain more results in assessing digital artwork. In our opinion, the digital tools did improve the children's artistic expression and creativity.

This study is still in progress and will be completed soon. These are only the primary results from the sample test. The final outcome will enable a more detailed consideration of the effectiveness of the TCT-DP application in assessing the pupils' artwork. This will in turn enable a fuller evaluation of the extent to which digital tools can enhance students' creative proficiency.

REFERENCES

- [1] J. Lanier, "You are not a Gadget: A Manifesto". Journal of Communication, New York, ISSN0021-9916, 2010.
- [2] M. I. Pinsky, "The gospel according to Disney: Faith, trust, and pixie dust", Westminster John Knox Pr., 2004.
- [3] C. Cordes and E. Miller, "Fool's gold: a critical look at computers in childhood". College Park, MD: Alliance for Childhood. Children and computers in pre-school157, British Educational Communications and Technology Agency, 2005.
- [4] J. Matthews and P. Seow, "Electronic paint: Understanding children's representation through their interactions with digital paint". International Journal of Art & Design Education, vol. 26, 2007, pp. 251-263.
- [5] J. Matthews and J. Jessel, "Very young children use electronic paint: A study of the beginnings of drawing with traditional media and computer paintbox". Visual Arts Research, 19(1), 1993, pp. 47-62.
- [6] E. S. Tzafestas, "Integrating drawing tools with behavioral modeling in digital painting". Proceedings of the 2000 ACM workshops on Multimedia, ACM. 2000, pp 39-42.

- [7] D. Arrowood and T. Overall, "Using technology to motivate children to write: Changing attitudes in children and preservice teachers". Society for Information Technology & Teacher Education International Conference, 2004, pp. 4985-4987.
- [8] D. Katsiaficas, "Digital drawing exploring the possibilities of digital technology as an essential tool and component in contemporary drawing". Department of Art, University of Minnesota, Minneapolis, USA. *Fine Art in Egypt* 20, 2008, pp. 1-8.
- [9] S. W. Haugland, "What Role Should Technology Play in Young Children's Learning? Part 1", *Young Children*, vol. 54, 1999, pp. 26-31.
- [10] L. J. Couse and D. W. Chen "A tablet computer for young children? Exploring its viability for early childhood education". *Journal of Research on Technology in Education*, vol. 43, no. 1, 2010, p. 75.
- [11] Ö. Hans, "The scope of digital image media in art education". *Computers and Education Journal*, ISSN: 0360-1315 vol. 59, Issue 2, 2012, p. 793.
- [12] T. A. Kaynar "Teaching visual arts in primary school teaching departments with postmodern art education approach". *Procedia, Social and Behavioral Sciences Journal*, ISSN:1877-0428, Vol. 51, DOI:10.1016/j.sbspro.2012, pp. 1044-1049.
- [13] A. Clark, S. McQuail, P. Moss and T. Coram, "Exploring the field of listening to and consulting with young children". Research Report (London, DfES), 2003, p. 445.
- [14] A. Greig and J. Taylor, "Doing Research with Children: A practical guide". Sage, 2012.
- [15] M. Wighting, J. Baker and L. Grooms, "Development of an instrument to measure perceived cognitive, affective, and psychomotor learning in traditional and virtual classroom higher education settings". *The Internet and Higher Education* 12.1. 2009, pp. 7-13.
- [16] A. S. Aboalgasm and R. Ward, "Evaluating the Use of Digital Art Tools for Drawing to Enhance Artistic Ability and Improve Digital Skill among Junior School Students". *International Journal of Social Education, Economics and Management Engineering*, Vol.8, No.10, 2014, pp. 3287-3288.
- [17] F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology". *MIS Quarterly*, vol. 13, no. 3, 1989, p. 319.
- [18] A. S. S. Aboalgasm and R. Ward, "Applying a Modified Technology Acceptance Model to the Use and Assessment of Digital Art Tools". *The International Journal of the Image*, Volume 4, Issue 4, August 2014, pp.71-83.
- [19] K. K. Urban, "Assessing Creativity: The Test for Creative Thinking-Drawing Production (TCT-DP)". *International Education Journal*, vol. 6, 2005, pp. 272-280.
- [20] B. Jeffrey and A. Craft "Teaching creatively and teaching for creativity: distinctions and relationships". The Open University, UK, published online: 07 Oct 2010, p. 77-87.
- [21] G.A. Davis, "Review: Test for Creative Thinking - Drawing Production". *Gifted and Talented International*, 10 (2), 1995, pp. 90-91.
- [22] A. Cropley, "Review of Test zum Schoepferischen Denken - Zeichnerisch (TSD-Z) [Test of Creative Thinking - Drawing Production (TCT-DP)]". *High Ability Studies*, 7, 1996, pp. 224-22.

Nonparametric Estimation of Demand Structures in Airline Revenue Management

Johannes Ferdinand Jörg, Catherine Cleophas

Lehr- und Forschungsbereich Advanced Analytics (ADA)
RWTH Aachen University, Germany

Email: johannes.ferdinand.joerg@ada.rwth-aachen.de

Email: catherine.cleophas@ada.rwth-aachen.de

Abstract—Airline revenue management employs forecasting and optimization techniques to offer the right price at the right time to the right customer. With the ability to store large amounts of data comes the challenge to incorporate the information contained in those data sets. This contribution considers the estimation of demand segments present in a specific market using nonparametrical methods on panel data. We employ finite mixtures to model booking events in different time frames and to obtain an estimator for the number of demand segments. Via an airline revenue management simulation tool, we perform an experimental study of the derived estimator. We discuss the results with respect to the underlying demand structure of the simulation and identified demand segments. The findings are discussed with respect to theoretical and practical use. Finally, we discuss the real world applicability and possible further research intentions.

Keywords—nonparametric; demand segments; estimation; revenue management

I. INTRODUCTION

A central theme of revenue management is analyzing historical sales data to draw conclusions on the underlying demand structure. A multitude of sources influence booking data in airline revenue management. In addition, availability control and product restrictions censor sales. This leads to a discrepancy of historical sales data and actual demand. Studying this issue helps to understand the market and to improve revenue management optimization parameters.

To apply revenue management techniques, we have to be able to segment the market. The goal is to optimize the availability of products over time such that the customers' willingness to pay is exploited. Talluri and van Ryzin [1] give a detailed overview of revenue management techniques. Naturally, this motivates identifying demand segments and their behavior to forecast the number of customers arriving in a certain time frame. Airline revenue management usually segments demand by characteristics including customers' willingness to pay, utility costs for different product properties (e.g., weekend stay, minimum stay, economy or business class, number of transfers) and date of the request or cancellation. One example would be to simply segment the market into business travellers with a high willingness to pay and late booking requests and private travellers with lower willingness to pay and earlier booking requests. In practice, this relatively simple distinction does not suffice to obtain a satisfying forecast or optimization. Also, most of the related works assume a fixed number of demand segments as input into estimation or optimization procedures, e.g., [2], [3]. Thus, we are motivated to find ways to obtain a suitable amount of demand segments.

Revenue management forecasting usually relies on historical data to extrapolate demand. It is important to accurately forecast demand, as for example a study by Pölt [4] suggests that a 20% increase in forecasting accuracy may lead to an 1% increase in revenue. As the amount of data stored by companies steadily increases, methods to analyze these data sets are needed. Most current, practice-oriented approaches explain demand structures with parametric statistics or heuristics. For an overview of forecasting methods see, e.g., the taxonomy of Azadeh [5]. Parametric estimation needs specified underlying distributions. We focus on nonparametric estimation of demand structures here, which uses large data sets to remove the assumption of a specific underlying distribution.

One example of the application of nonparametric statistics to revenue management is presented by Van Ryzin and Vulcano [6]. They propose an expectation-maximization approach for analysing market structures. They identify demand segments with their preferences over the set of alternatives and use an iterative algorithm to create new sets of demand segments, such that the probability for the observed booking distribution is maximized.

To model customer decisions, we employ a discrete choice model: Customers buy one of a set of products at each specified time period. In the case of quantity-based airline revenue management, we identify products with discrete booking classes on a specific flight itinerary and time periods correspond to flight departures. Our proposed method to estimate the number of demand segments works as follows: In a first step, we formulate our model for panel data of two time periods using a finite mixture model to represent the probability of a booking event. An introduction to finite mixture models can be found in McLachlan and Peel [7]. In the next step, we decompose the model such that the estimation of the number of demand segments becomes a rank estimation problem. This idea is based on an approach of Kasahara and Shimotsu [8], in which the authors derive sufficient conditions for the nonparametric identifiability for various finite mixture models in the framework of discrete choices.

This short paper is structured as follows: In Section 2, we define a mathematical model of the booking process and derive a lower bound for the number of demand segments in a market. Section 3 discusses a study, which employs an airline revenue management simulation tool in order to test the methods in a theoretical environment. Finally, Section 4 reviews our findings, provides an outlook of the intended extensions to the model, and addresses practical applicability.

II. METHODS

The notion of finite mixture models is often applied in economics, chemistry and health care [9]. In principle, this modeling technique can represent an observation set as a composition of several subpopulations. In the context of airline revenue management, this can explain seemingly homogeneous booking data by presuming the presence of a number of demand segments.

As input for our estimation method, we use panel data. Panel data is the generic term for data which is collected at different times for the same population and the same indicators. Here, we will observe customers for a set amount of flight departures and store their booking decisions.

Suppose we have panel data of individuals over a number of T time periods. In each time period $t \in T$, an individual chooses to buy one alternative x_i out of the set of available alternatives X . Let M be the number of mixture components contributing to the observed data. The probability of buying a product is denoted by $p_m^*[x_1]$ for the first time period and by $p_{m,t}[x_t]$ for $t > 1$. Each mixture component m contributes a factor $\omega_m \in [0, 1]$ to the observation. The baseline model of our investigation then looks as follows:

$$P[\{x_t\}_{t=1}^T] = \sum_{m=1}^M \omega_m p_m^*[x_1] \prod_{t=2}^T p_{m,t}[x_t], \quad (1)$$

where $\sum_{m=1}^M \omega_m = 1$ ensures that the probability (1) consists of the contribution of demand segments. This means that the probability for an event $\{x_t\}_{t=1}^T$ is completely explained by the M mixture components.

Now, we can identify the nature of demand segments: Each mixture component m represents a demand segment with its own probability to book a specific class x_t at different times t . Thus, our objective to estimate the number of demand segments corresponds to estimating parameter M of our model. We assume that panel data of customers for $T = 2$ time periods is available.

Since the temporal spacing of time periods is not yet defined and we consider an airline revenue management context, let the time periods correspond to flight departures. On all flights, the same set of booking classes X is offered such that a booking event for two time periods is a tuple (x_1, x_2) . Therefore, the model (1) simplifies to

$$P[(x_1, x_2)] = \sum_{m=1}^M \omega_m p_m^*[x_1] p_{m,2}[x_2]. \quad (2)$$

Let N be the number of individual observations in the panel data and denote by $x^i = (x_1, x_2)$ the observation tuple of the i^{th} individual. Given this data, we can estimate the probabilities $P[(x_1, x_2)]$. In a first approach, we use plug-in estimates for this probability, i.e.,

$$p_{i,j} = \hat{P}[(x_1 = i, x_2 = j)] = \frac{\sum_{i=1}^N \mathbf{1}_{\{(x_1=i, x_2=j)\}}(x^i)}{N}. \quad (3)$$

With these estimates, we can calculate an observable $|X| \times |X|$ matrix $P = (p_{i,j})$

$$P = \begin{pmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,|X|} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,|X|} \\ \vdots & \vdots & \ddots & \vdots \\ p_{|X|,1} & p_{|X|,2} & \cdots & p_{|X|,|X|} \end{pmatrix} \quad (4)$$

Note here, that in order to identify M , it is required to have $|X| \geq M$.

Defining $V = \text{diag}(\omega_1, \dots, \omega_M)$, this admits the decomposition $P = P_1 V P_2$, where P_1 and P_2 are $|X| \times M$ and $M \times |X|$ matrices consisting of the entries p_m^* and $p_{m,2}$, respectively. We can now apply a simple rank argument: Since $\text{rank}(V) = M$ and $\text{rank}(P_i) \leq M$, we have that

$$\text{rank}(P) \leq \min \{ \text{rank}(P_1), \text{rank}(V), \text{rank}(P_2) \} \leq M \quad (5)$$

Therefore, the rank of matrix P yields a lower bound for the number of demand segments M , i.e.,

$$M \geq \text{rank}(P). \quad (6)$$

III. RESULTS

Our results are based on an airline revenue management simulation. It models the complete revenue management process from forecasting demand to optimizing available booking classes, to taking reservations of arriving customer requests. It also includes an extensive stochastic demand model to create artificial demand. Through simulation experiments, we obtain data sets, which provide input for the estimation process. Customers are generated according to several parameters (e.g., willingness to pay, preferred flight departure time, utility costs for specific product properties). The algorithm draws these values with a given error term from a normal distribution, to create realistic variation. Demand segments in the context of our simulation tool are then a specific set of these underlying parameters. The set of available alternatives X is defined by a set of 12 booking classes and the no-purchase alternative. We distinguish booking classes by price, compartment and additional product properties, such as being refundable or rebookable or requiring a minimum stay. These product properties impose, depending on the demand segment, either a penalty on the utility of the customer or an acceptance probability, such that the customer has to decide if he accepts this restriction or not without having further implications on the utility. As expected, the more limitations are imposed on a booking class, the cheaper its price. It is also possible for booking classes to only differ in price. Three of the 12 booking classes are located in the business compartment, the remaining 9 are located in the economy compartment.

Customers are identified by a unique customer identification number. Thus, we can track their purchases over several flight departures. We also assume that the simulation parameters do not change over the course of both time periods, i.e., the willingness to pay, product restriction costs, and request and cancellation date are constant for both flights for which they requested tickets. In order to obtain different bookings given constant demand, we varied the forecasts for both flights such that the revenue management system does not offer the same product availabilities in both cases. This was achieved by introducing a forecast error as follows: For each booking class on one flight itinerary, we draw a realization of a random variable with uniform distribution on $[0.8, 1.2]$. The forecast

for this booking class is then multiplied by the value of the realization. This ensures that customers do not always book the same class and that the panel data obtained is sufficiently diverse.

To compute the rank of matrix P , we used a QR decomposition with varying tolerance for eigenvalues which are considered to be zero. This will be denoted by QRx , where x is the number of decimal places which are considered non-zero and the values of QRx are the lower bound for the number of demand segments computed. We also modified the number of customers we tracked over several simulation runs, to assess the amount of data needed for convergence.

We performed the estimation procedure with three different scenarios. The first data set is based on a simulation with only two demand segments. These customers who book mainly prefer the economy compartment, one demand segment that has higher willingness to pay and the other has a lower willingness to pay. The second data set describes a scenario with six demand segments. Each of the demand segments has its own set of parameters and therefore its own subset of the available booking classes, which customers consider buying. While this may be the most realistic scenario considered here, we plan to perform the estimation for additional scenarios. The third data set is similar to the first, but one demand segment is mainly business oriented while the other one is mainly economy oriented, such that there is an inherent separation within the data set.

TABLE I. SIMULATION WITH TWO DEMAND SEGMENTS

Number of customers	QR7	QR4	QR3	QR2	QR1
100	7	7	7	7	7
1000	7	7	7	7	7
10000	7	7	7	7	5
50000	8	8	7	6	5
100000	8	8	7	6	4

At first glance, the results exhibit some peculiarities. Table I shows that the lower bound of demand segments actually increases when the number of observations increases, i.e., from seven to eight demand segments when observing 10000 and 50000 customers. We would expect the number of demand segments to decrease as demand observations increase. The reason that this is not always the case lies in the simulation design. Since customers are stochastically generated, outliers buy products that were not bought in simulation runs before. The algorithm recognises these customers as a new demand segment.

TABLE II. SIMULATION WITH SIX DEMAND SEGMENTS

Number of customers	QR7	QR4	QR3	QR2	QR1
100	12	12	12	12	12
1000	13	13	13	13	13
10000	13	13	13	13	12
25000	10	10	10	10	10
50000	8	8	8	8	8
100000	7	7	7	7	7

Table II shows the results for the simulation with six demand segments. A similar effect to the one in Table I takes place for 100 and 1000 observations: Again, the lower bound for the number of demand segments increases due to the increase in generated customers and thereby customers who

booked booking classes that were not observed before. Here, we also note that the sensitivity to the non-zero eigenvalues is lowered compared to Table I. The only occurrence of a difference is at 10000 customers observed from 13 to 12 demand segments.

TABLE III. SIMULATION WITH TWO DEMAND SEGMENTS WITH DISTINCT BOOKING BEHAVIOUR

Number of customers	QR7	QR4	QR3	QR2	QR1
100	10	10	10	10	10
1000	10	10	10	10	8
10000	8	8	8	8	6
50000	7	7	7	7	4
100000	5	5	5	5	4

Table III shows the results of the estimation for the third scenario setup. It includes two demand segments, which largely differ in their booking behaviour. This means that customers from one demand segment usually book cheaper economy classes while the others book the more expensive economy and business classes. A low amount of observations leads to a high amount of estimated demand segments, since there is a broad spectrum of booking classes booked. As the number of observations increases, we obtain an even lower amount of demand segments than in the first scenario (compare Table I). Since it should be easier to separate two demand segments, which are clearly distinct, than two which overlap in booking behaviour, this meets our expectations.

These preliminary results indicate that the approach does better for a higher amount of demand segments. The identification of demand segments in this model still remains open, as the presented methods do not explain the relationship between the simulation-generated demand segments and those identified by the estimation procedure presented here.

IV. CONCLUSION

We have developed a procedure to estimate a lower bound of the number of demand segments from panel data. This approach does not rely on the specification of an underlying demand structure and thus can be used without the knowledge of specific distribution functions. First preliminary results, which are based on a simulation study of an airline revenue management tool, show promising results.

There are several gaps to be filled in order to use the presented methods in real world applications. Some simplifications limit the generality of the model, e.g., the assumption that the probabilities are not dependent on previous time periods. This is clearly not the case in real world scenarios. In the introductory example, we can easily deduce that for business travellers, the probability to book a business class again should be higher than buying a lower economy class. We expect that the probabilities are not too erratic. Still, assuming that we obtained a number of demand segments in a market, we may now use that knowledge to derive the probability distribution of each segment. This requires panel data for at least three time frames as is described in [8]. With these probability distributions and the information of the booking data, we may calculate a forecast for future flights.

Another point of interest is the distribution of requests for the demand segments. As the estimation procedure uses final booking data, i.e., booked classes after departure of the flight,

the information about when the customer booked his flight is lost. However, the request date is a crucial part of dynamic revenue management optimization techniques. We may argue that if we can obtain panel data in the first place, we should be able to know when the booking was made. This would present us with the possibility to make inference of when the demand segments likely request over the booking horizon.

It is also difficult to assess the quality of the preliminary results, as we do not know yet which initially created segments are identified. The ultimate test should be to compare the forecast, which resulted from this estimation procedure, with forecasts of state-of-the-art methods currently used in practice. One may also substitute different parts of already established forecasting methods with nonparametric estimation procedures in order to alleviate the need for a specific underlying distribution which has to be specified beforehand. For example, we may be able to represent price elasticities with the probabilities of the demand segments.

In current practice, creating observation sets in the form of panel data is difficult, since we usually only have access to the information of how many bookings were observed in each offered class after the flight departed. Neither the time of the request nor the set of available classes at that time is stored. The need to track each single customer for a set amount of time frames is imperative for the presented estimation method. In airline revenue management, we may consider the use of bonus cards or improved tracking of online purchases of flights to create such data sets. Other areas, where revenue management is commonly practised, e.g., hotels and car rentals, may also use the fact that customers usually have to leave personal information when renting rooms or cars. The access to this kind of information should be available and give more detailed data sets.

The next steps to extend this model include to incorporate more observable characteristics into the data set in order to lower the amount of observations needed to obtain reasonable results. We may additionally consider the availability of booking classes at the request time as an explaining characteristic. Also, other panel data sets have to be constructed from other simulations to assess the behaviour of the lower bound. We plan to create scenarios with up to 12 demand segments present. This would amount to a scenario where each demand segment has exactly one booking class which he is willing to buy. In revenue management, these kinds of customers are called product customers as opposed to priceable customers, e.g. see [10], since their choice set only consists of the specified booking class and the no-purchase alternative. Another approach would be to allow for incomplete observations of customers, i.e., assess the viability of this method for customer data sets which do not track every single customer for two or more time periods. Incomplete or inaccurate observations may also be created when the technique of collecting this panel data is not able to exactly assign recurring customers. We plan to use different rank estimation methods for our matrix rather than an explicit calculation, e.g., the rank statistic of Robin and Smith [11] or Kleibergen and Paap [12].

REFERENCES

[1] K. T. Talluri and G. J. Van Ryzin, *The theory and practice of revenue management*. Springer Science & Business Media, 2004, vol. 68, ISBN: 978-1-4020-7701-2.

[2] P. Hall and X.-H. Zhou, "Nonparametric estimation of component distributions in a multivariate mixture," *Annals of Statistics*, 2003, pp. 201–224, ISSN: 0090-5364.

[3] E. S. Allman, C. Matias, and J. A. Rhodes, "Identifiability of parameters in latent structure models with many observed variables," *The Annals of Statistics*, 2009, pp. 3099–3132, ISSN: 0090-5364.

[4] S. Pölt, "Forecasting is difficult—especially if it refers to the future," in *AGIFORS-Reservations and Yield Management Study Group Meeting Proceedings*, 1998, pp. 61–91.

[5] S. Sharif Azadeh, "Demand forecasting in revenue management systems," Ph.D. dissertation, École Polytechnique de Montréal, 2013.

[6] G. van Ryzin and G. Vulcano, "A market discovery algorithm to estimate a general class of nonparametric choice models," *Management Science*, vol. 61, no. 2, 2015, pp. 281–300, ISSN: 1526-5501.

[7] G. McLachlan and D. Peel, *Finite mixture models*. John Wiley & Sons, 2000, ISBN: 978-0-471-00626-8.

[8] H. Kasahara and K. Shimotsu, "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, vol. 77, no. 1, 2009, pp. 135–175, ISSN: 1468-0262.

[9] P. Schlattmann, *Medical applications of finite mixture models*. Springer Science & Business Media, 2009, ISBN: 978-3-540-68650-7.

[10] T. Fiig, K. Isler, C. Hopperstad, and P. Belobaba, "Optimization of mixed fare structures: Theory and applications," *Journal of Revenue & Pricing Management*, vol. 9, no. 1, 2010, pp. 152–170, ISSN: 1476-6930.

[11] J.-M. Robin and R. J. Smith, "Tests of rank," *Econometric Theory*, vol. 16, no. 02, 2000, pp. 151–175, ISSN: 0266-4666.

[12] F. Kleibergen and R. Paap, "Generalized reduced rank tests using the singular value decomposition," *Journal of econometrics*, vol. 133, no. 1, 2006, pp. 97–126, ISSN: 0304-4076.

Environmental Performance Evaluation of Different Glazing-Sunshade Systems Using Simulation Tools

Adeeba A. Raheem
Department of Civil Engineering
University of Texas – El Paso
El Paso, TX 79968
e-mail: aaaraheem@utep.edu

Svetlana Olbina
Rinker School of Construction Management
University of Florida
Gainesville, FL
e-mail: solbina@ufl.edu

Raja R. A. Issa
Rinker School of Construction Management
University of Florida
Gainesville, FL
e-mail: raymond-issa@ufl.edu

Ian Flood
Rinker School of Construction Management
University of Florida
Gainesville, FL
e-mail: flood@ufl.edu

Abstract— With the increasing capabilities of computer modeling and simulation technology, the analysis of options for maximizing gains in energy efficiency for buildings can be realized more efficiently and cost effectively. Conducting building performance simulations allows for the analysis of the environmental impacts of buildings at the early design stages. Commercial buildings consume nearly one fifth (18 quads) of all the energy used in the United States, costing more than \$200 billion each year. The building envelope plays a key role in determining how much energy is required for the operation of a building. Individual thermal and solar properties of glazing and shading systems only provide information based on static evaluations, but it is very important to assess the efficiency of these systems as a whole assembly under site-specific conditions. This paper presents a case study that was conducted using computer simulation tools to evaluate the environmental impacts of using different types of glazing-sunshade systems on the overall performance of an office building. The case study results show how early stage building performance studies using computer simulation tools help practitioners in achieving the goals of reduced energy consumption and increased indoor comfort in an economical manner.

Keywords- *Simulation; Energy consumption; Indoor comfort; Commercial buildings.*

I. INTRODUCTION

The use of computer modeling and simulation tools during the early design phase can be very helpful to obtain more reliable building performance predictions [8][19]. The models can be a simple building information model (Level Of Development (LOD) 200) or detailed one (LOD 300) depending on the owner's requirements or other physical conditions [3]. A major advantage of using these tools is the comparison of the environmental performance of different design alternatives to improve the overall building design efficiency [13].

According to the U.S. Department of Energy [21], commercial buildings consumed about 18.26 quads of primary energy in 2010, which represents 46% of building

energy consumption, 19% of U.S. energy consumption and 18% of U.S. carbon dioxide emissions. Space heating, lighting and cooling were the major end uses representing 26.6%, 13.6% and 10.1% respectively of the total energy consumed by the commercial buildings during the same year [21]. The energy performance of building envelope components is critical in determining how much energy is required for heating, cooling and lighting. Energy efficiency measures such as properly designed envelopes in commercial buildings provide an opportunity to reduce energy consumption and costs, and to reduce greenhouse gas (GHG) emissions at the same time.

Florida has become the nation's fourth largest energy consuming state in the commercial building sector utilizing about a Quadrillion BTU's in commercial consumption and having a gross expenditure of over \$10 billion per year in this sector [24]. Of the total energy that Florida produces per year, more than 90% comes from non-renewable sources like coal and gas contributing 4.8 million metric tons of energy related carbon-dioxide emissions from the commercial building sector to the total emissions per year. In an effort to decrease the carbon footprint of this high energy consumption in commercial buildings, more stringent rules for building envelope design have been adopted in the Florida Building Code's Energy Conservation Code section [11]. Much of the emphasis is given to the window to wall ratio, U-factor and solar heat gain coefficient (SHGC) of window glass and frame type while describing the energy efficient window strategies in Section 502 (Building envelope requirements) of FBC 2010. Although there is a potential for the use of advanced window systems such as switchable electrochromic or gasochromic windows in reducing the overall energy loads, widespread use is unlikely to occur in the near future due to high initial costs and lack of technical expertise [16][14]. Hence other related options such as automated shading systems could be deployed while satisfying the thermal and daylighting requirements of the occupants.

II. LITERATURE REVIEW

Wherever Individual thermal and optical properties of glazing/shading systems only provide information based on static evaluations. However, it is very important to assess the efficiency of these systems as a whole assembly under site-specific conditions. Cazes [6] found that a positive energy balance (which depends on the season, building type and operation of the building) can be achieved using advanced static glazing combined with well-insulated window systems and architectural shading optimized for seasonal impacts. Past research has shown that the proper use of shading devices may reduce the cooling loads by 15-20% (depending on the amount and location of the windows) [4] [5] [9]. The occupants in commercial buildings often complain about too much solar heat and glare, both of which can be reduced by the use of advanced glazing systems that are tuned to reject as much heat as possible while transmitting high levels of visible light and preventing glare. Combining these advanced solar control glazing (static SHGC) and exterior architectural shading offers an improved solution [14]. Although there is a great potential for advanced window systems such as switchable electrochromic or gasochromic windows in reducing the overall energy loads, still more research and economies of scale are needed so that these systems can become cost-effective (market viable) for mainstream markets (Table I) [14].

There are several variables that influence the thermal comfort of building occupants such as personal variables, environmental variables and physiological variables [1] [10] [15]. The personal and physiological variables are controlled/owned by the occupants. The effects of external environmental variables on the indoor comfort of the occupants can be controlled through a proper fenestration design. It is critical to design south façades properly since during the day they receive a large amount of energy from the sun through the glazing and usually most of the sunlight gets concentrated in certain areas of the space if the facade is not properly designed [17][18]. This may result in glare on work surfaces causing discomfort for the occupants [12] [25].

TABLE I. BEST AVAILABLE TECHNOLOGIES (BAT) FOR WINDOWS AND CLASSIFICATION BASED ON MARKET READINESS (ADAPTED FROM [13])

Key Technical Attribute	BAT (Market Viable)	BAT(Pre-market Viable)	Future Technology
Low U-value	Triple-glazed, dual low-e-coating, advanced frames	Quadruple-glazed, exotic inert gases, aerogel-filled frames	Vacuum-insulated glass, market- viable, multiple-glazed cavity system
Variable SHGC	Automated shade control, exterior shading, architectural features	Dynamic solar control	Dynamic glazing

III. METHODOLOGY

The research was conducted in three phases: 1) modeling and simulation, 2) analysis of total annual energy consumption and heat gains through glazing and 3) comparison of results for the three best performing glazing-sunshade systems.

A. Modeling and Simulation

The COMFEN 5 [7] software was used for energy and visual modeling and simulation. This is a single-zone façade analysis tool based on EnergyPlus software and it is used to evaluate energy, thermal and visual performance of commercial building façades using different design scenarios. COMFEN provides comparative perimeter zone performance results between façade design options.

The base model used in this research was an office building (80' x 50') with an area of 4000ft² located in Miami, Florida. The model was designed for the south façade consisting of a curtain wall system (glazing/façade ratio=0.67). The curtain wall was simulated for 40 different glazing-sunshade systems (four glass types, nine shading systems and a base case of glazing with no sunshade). The base model was simulated using nine different types of shading systems for three possible locations relative to the window (exterior, between glass and interior)(see Table II).

B. Analysis of total annual energy consumption and heat gains

The building model was simulated multiple times for different glazing-sunshade systems. The results were analyzed in terms of total annual energy consumption and heat gains through glazing. The three best performing glazing-sunshade systems were then selected to compare their overall performance in terms of energy and indoor comfort.

C. Comparison of the three best performing glazing-sunshade systems

The results obtained through analysis were then compared to find an optimum glazing-sunshade system for the south façade with the least energy consumption, maximum indoor thermal and visual comfort

TABLE II. CATEGORIES OF SHADING DEVICES USED IN THE BASE MODEL

Categories of shading devices	Types				
	Venetian blinds (45°)	Venetian blinds (90°)	Screen	Rolling shades	Overhangs
Exterior	×	×	×	×	×
Between glass	×			×	
Interior	×			×	

The comparison was made under two categories:

- Energy consumption: Annual energy and peak demand impacts
- Indoor comfort analysis: Thermal comfort, daylighting and glare

IV. INPUT DATA FOR MODELING AND SIMULATION

The following input was used for modeling and simulation:

- Geographical location: Miami FL; 25 49' 26" N 80 17' 59" W
- IECC climate zone= 1
- Heating and cooling degree days [10]: HDD=149; CDD=4361
- Building dimensions (LxW)= 80' x 50'
- Weather data file used= TMY3
- Required EnergyPlus file types= *.epw, *.stat, and *.ddy

The climate of Miami is essentially subtropical, characterized by a long and warm summer, with abundant rainfall, followed by a mild, dry winter. The annual temperature profile shows high temperatures during summer (above 90°F) with similar peaks of direct and diffused solar radiations. Due to high outside temperature, Miami requires both sensible and latent cooling most of the year.

The ASHRAE standard 90.1 was used to determine envelope insulation requirements for the Miami climate. The lighting and cooling load values were used as suggested in the ASHRAE guide for energy efficient small office buildings[2]. The outdoor air flow rates used for ventilation were based on the area of the building (flow/area: cfm/ft2) (ASHRAE 90.1). Average carbon emissions per unit of electricity (generated by utility and nonutility electric generators) and gas values were taken from data provided by the EIA [23]. The selected glazing systems had U-values ranging between 0.1 and 0.3 and SHGC below 0.5 with double and triple glass types (Table III). These systems were comprised of multiple glass-gas layers and their thermal and optical properties like U-values, T_{vis} (visible transmission) and SHGC (Solar Heat Gain Coefficient) were calculated using WINDOW 7 software.

TABLE III. SELECTED GLAZING SYSTEMS FOR SIMULATION

ID	Glazing type	U-value (Btu/h-ft ² -F)	SHGC	T_{vis}	Thickness (in)
G1	Double glass low solar low-E clear (Argon)	0.23	0.37	0.7	0.95
G2	Double glass low T_{vis} low-E (Argon)	0.203	0.241	0.371	0.95
G3	Triple w/suspended film; dual low-E	0.144	0.467	0.631	1.45
G4	Triple, dual low-e; pyrolytic	0.145	0.3	0.541	1.67

V. ENERGY CONSUMPTION AND HEAT GAIN ANALYSIS

The energy consumption and heat gain analysis was performed for four sets of glazing-sunshade systems with each set comprised of one glazing system with all ten selected shadings. The energy consumption was measured for four energy usage categories: heating, cooling, fans and lighting. For the first set of glazing-sunshade systems, double glass (G1- low solar low-E clear (Argon)) with different types of shading (S1-S10) was simulated keeping all the other design and space parameters the same in each simulation. It was observed that the least amount of total energy (for heating, cooling, fans and lighting) was consumed when overhangs (10) were used whereas exterior roller shades (4) were the most efficient ones in reducing cooling loads (Fig. 1(a)) due to the least heat gains through windows (Fig. 1 (b)).

For the second set double glass (G2 - low T_{vis} low-E clear (Argon)) with different types of shadings (S1-S10) was simulated again keeping all the other design and space parameters the same in each simulation. The results in this case showed a decrease in the total energy consumption and window heat gains for each of the glazing-sunshade systems. It was further observed that the least amount of total energy (for heating, cooling and electricity) was consumed when no sunshade system was used. Similarly, the third and fourth sets of glazing-sunshade systems were analyzed and the best options were selected for the final comparison in terms of energy and indoor comfort.

VI. COMPARISON OF ENERGY PERFORMANCE AND INDOOR COMFORT

Three best performing glazing-sunshade systems were selected from the four sets after analyzing their energy performance for the south facing curtain wall. These systems were:

- A1: Double glass low VT low-e (Argon) (G2) with no sunshade system (S1)
- A2: Double glass low solar low-e (Argon) (G1) with overhangs (S10)
- A3: Triple glass, dual low-e; pyrolytic (G4) with external roller shades (S4).




Figure 1. Analysis of the first set of glazing-sunshade (G1 and S1-S10): (a) Energy consumption (b) Annual heat gains through glazing. Horizontal axis: 1-no sunshade; 2-External venetian blind 45°; 3-External venetian blind 90°; 4-External roller shade; 5-External screen; 6-Between glass venetian blind; 7-Between glass roller shade; 8-Internal venetian blind; 9-Internal roller shade; 10- Overhangs

These selected systems were further compared for 1) energy performance and 2) indoor comfort

A. Energy Consumption

The total energy usage was calculated as the sum of the three energy use types (heating, cooling and electricity (fans and lighting)). Based on the total annual energy use values, double glass low solar low-e with overhangs (A2) was the most efficient assembly in the current scenario. From the monthly energy consumption profile, it was observed that from March through September assembly A2 performed better than the other two (A1, A3) but from Jan-Feb and Oct-Dec all three systems were performing nearly in the same manner (Fig. 2) because the direction of conductive heat flow is from the inside to the outside of the building during these months in Miami.

B. Indoor Comfort

Thermal comfort. Three selected systems were analyzed in terms of thermal comfort and results were obtained as a percentage of people satisfied which is a direct output of the software used (Table IV).

The A3 assembly provided the best thermal comfort as it had the highest percentage of people satisfied, and least number of hours in a year when hourly temperature set points were not met (Fig. 3).

Daylighting and glare analysis. A daylighting analysis was performed for the selected systems and daylight illuminance maps were generated for a summer day (June 21st at 11:00AM). These maps display work surface illuminances, calculated at 2'-6" (0.762 m) above the floor (default value), for the entire space in the form of a 10 x 10 grid (the grid is scaled to fit the space within the software). The maps showed high illuminance values for systems A1 and A2 near the façade area inside the office, whereas a low, but uniform illuminance level was observed when using assembly A3 because the roller shades were automatically positioned (e.g. closed either fully or partially) at that time of the day (Fig. 4)




Figure 2. Energy consumption: (a) Annual profile, (b) Monthly profile

The selected systems were further compared to study the glare during a clear summer day from the South side. The occupant's position (X=9.3, Y=15.6) and angle of view (X=6, Y=-9.6) were defined and point-in-time simulations were run for June 21st at 9:00AM, 12:00PM and 3:00PM (Fig. 5). It was observed that use of systems A1 and A2 caused very high values of glare (>185cd/ft2 (2000cd/m2)) during the morning and afternoon which is uncomfortable for the occupants whereas use of assembly A3 caused low glare values (average was less than 69cd/ft2 (750cd/m2)) (~55% less than A1 and 61% less than A2 at noon) during most the daytime.

VII. CONCLUSIONS

This case study was conducted using energy simulation software to look at the environmental of different glazing-sunshade systems for a glazed office building in the hot and humid climate of Miami, Florida. It was observed from the analysis that sunshade system behave differently (in terms of overall efficiency) with different glazing systems. For south facades, exterior shading such as roller shades and overhangs are the most efficient options when combined with glazing systems having low U-value and SHGC (<0.3). The analysis also showed that although the least amount of energy was consumed annually when overhangs were used, they are not the best option in terms of providing indoor comfort for the occupants.

TABLE IV. THERMAL COMFORT ANALYSIS

Thermal Comfort Factors	Window Shading Assemblies				
	A1	A2	A3	A3 vs. A1	A3 vs. A2
Average thermal comfort (PPS*)	86.37	85.09	88	>2%	>3.3%
Hourly temperature set points unmet (hours)	1173	1223	875	-289	-348

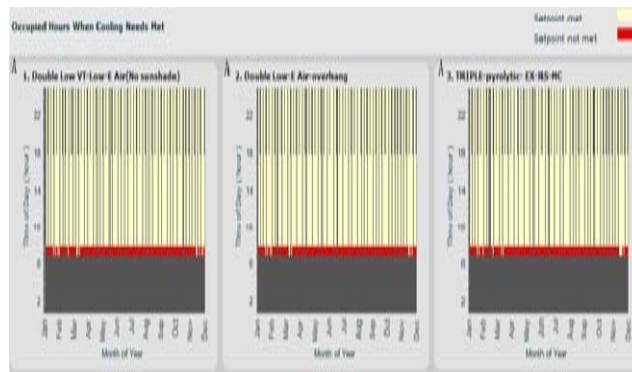


Figure 3. Occupied hours in the building when cooling needs are met.




Figure 4. Daylight analysis- Top: Daylight illuminance maps; Bottom: Perspective view illuminance contour lines.




Figure 5. Rendered images from Radiance showing luminance ranges for the selected systems (A1, A2, and A3) during different times of a summer day (June 21st).

More specifically the following conclusions were reached for the south façade glazing-sunshade system design:

- Glazing-sunshade systems with very low thermal and visual properties (visual transmittance (<0.3), U-value (<0.25) and SHGC (<0.25)) like system A1 used in this study can provide some degree of sun control without any sunshade system but can also increase glare and thus do not provide a comfortable indoor environment for the occupants. These systems can help in decreasing electrical loads but compromise the indoor comfort.
- Glazing-sunshade systems with relatively high visual transmittance (<0.8), low U-value (<0.25) and moderate SHGC (<0.4) (like system A2 used in this study) can work efficiently with fixed horizontal shading such as overhangs. Because of the low initial cost, this may be the preferred system; however, indoor comfort is compromised due to high glare during the morning and afternoon hours.

- Glazing-sunshade systems with a moderate visual transmittance (<0.6), low U-value (<0.2) and low SHGC (<0.3)) (like system A3 used in this study) worked efficiently with external shades, such as roller shades and venetian blinds, to reduce cooling loads and permit filtered views. Because of less conduction of direct sunlight (automated control), glare is reduced at times when the horizontal solar angle is low.

The conclusion of this case study is that for the occupant’s comfort relative to the standard thermal and visual set points, multiple-pane glazing systems with external sunshade systems such as architectural components (overhangs) and roller shades are an efficient strategy for south facing facades installed on office buildings in Miami. Furthermore, the results indicate early building performance studies using computer simulation tools are helpful to practitioners in achieving the goals of reduced energy consumption and increased indoor comfort.

REFERENCES

- [1] ASHRAE. (1985). *Physiological Principles for Comfort and Health. Fundamental Handbook*, American Society of Heating, Refrigerating and Air Conditioning Engineers, Atlanta, GA.
- [2] ASHRAE (2004). *Advanced Energy Design Guide for Small to Medium Office Buildings*. Accessed December 12, 2013. <https://www.ashrae.org/standards-research--technology/advanced-energy-design-guides/50-percent-aedg-free-download>
- [3] AUTODESK. (2013). *Autodesk : Sustainability Workshop*, Accessed December 20, 2014. <http://sustainabilityworkshop.autodesk.com/buildings/revit-whole-building-energy-analysis>
- [4] Ali, A., and Ahmed, T. (2012). “Evaluating The Impact Of Shading Devices On The Indoor Thermal Comfort Of Residential Buildings In Egypt”. In *Proceedings of the fifth international conference of IBPSA-USA*, 603-612. Madison, Wisconsin.
- [5] Bourg, J. (2008). “Sun Control and Shading Devices”. *Whole Building Design Guide*, National Institute of Building Sciences. Accessed January 23, 2014 <http://www.wbdg.org/resources/suncontrol.php>
- [6] Cazes, B. (2011). ”Windows and glazed area technologies and materials in Europe”. IEA (International Energy Agency) *Building Envelope Technologies and Policies Workshop*, November 17, Neuilly-sur-Seine, France.
- [7] COMFEN manual. (2012). Accessed January 12, 2014. http://windows.lbl.gov/software/comfen/5/comfen_docs.htm
- [8] Crawley, D. B., Hand, J. W., Kummert, M., and Griffith, B. T. (2005).”Contrasting the Capabilities of Building Energy Performance Simulation Programs.” *Proceedings of Building Simulation 2005*, Montreal, Quebec, Canada, IBPSA, 231-238.
- [9] Dubois, M. C. (1997). “Solar Shading and Building Energy Use: A Literature Review”. Report TABK—97/3049, Lund Institute of Technology, Dept. of Building Science, Lund, Sweden.
- [10] Fanger, O.P. (1986). “Radiation and Discomfort”. *ASHRAE Journal*, 33-34.
- [11] Florida Building Code (FBC). (2010). *Commercial Energy Efficiency*, Chapter 5. Accessed December 21. <http://www.flaseia.org/Documents/FlyerEnergyJanuary2012.pdf>

- [12] Galasiu, A. D. and J. A. Veitch. (2006). "Occupant Preferences and Satisfaction with the Luminous Environment and Control Systems in Daylit Offices: A Literature Review." *Energy and Buildings*, 38: 728–742.
- [13] Hensen, J., and Lamberts, R. (2011). "Building Performance Simulation for Design and Operation." Routledge, UK
- [14] International Energy Agency (IEA). (2013). "Technology Roadmap Energy efficient building envelopes". Accessed December 25, 2013. <http://www.iea.org/publications/freepublications/publication/TechnologyRoadmapEnergyEfficientBuildingEnvelopes.pdf>
- [15] ISO. (1983). "Determination of the PMV and PPD Indices and Specification of the Conditions for Thermal Comfort". DIS 7730, Moderate Thermal Environment.
- [16] Lee, E., Selkowitz, S., Levi, M., Blanc, S., McConahey, E., McClintock, M., Arup, O., Hakkarainen, P., Sbar, N., and Myser, M. (2002). "Active Load Management with Advanced Window Wall Systems: Research and Industry Perspectives." ACEEE 2002 Summer Study on Energy Efficiency in Buildings: Teaming for Efficiency, August 18-23, 2002, Asilomar, Pacific Grove, CA. Washington, D.C.: American Council for an Energy-Efficient Economy.
- [17] Lawrence Berkeley National Laboratory (LBNL). (2004). Daylighting: Shading Strategies, Section 5. Accessed January 13, 2014. <http://windows.lbl.gov/daylighting/designguide/section5.pdf>
- [18] Littlefair, P.J., (1995). "Light shelves: Computer Assessment of daylighting performance". *Lighting Research and Technology*, 27(2); 79-91.
- [19] Maile, T., Fischer, M., and Bazjanac, V. (2007). "Building Energy Performance Simulation Tools - a Life-Cycle and Interoperable Perspective." CIFE Working Paper #WP107
- [20] U.S. Census Bureau. (2009). Table 396. Cloudiness, Average Wind Speed, Heating and Cooling Degree Days, and Average Relative Humidity -- Selected Cities. Accessed December 21. www.census.gov/compendia/statab/2012/tables/12s0396.xls
- [21] U.S. Department of Energy (U.S. DOE). (2011). "Building Energy Data Book". Accessed January 1, 2014. <http://buildingsdatabook.eren.doe.gov/ChapterIntro3.aspx>
- [22] U.S. Environmental Protection Agency (EPA). (1997). "An Office Building Occupant's Guide to Indoor Air Quality". Report EPA-402—97-03. Indoor Environments Division, Washington, DC.
- [23] U.S. Energy Information Administration (EIA). (2002). Table 1. 1998-2000 Average State-level Carbon Dioxide Emissions Coefficients for Electric Power. Updated State-level Greenhouse Gas Emission Coefficients for Electricity Generation, 1998-2000. Accessed February 26, 2014. <http://www.eia.gov/FTP/ROOT/environment/e-supdoc-u.pdf>
- [24] U.S. Energy Information Administration (EIA). (2010). Energy Consumption by End-Use Sector, Ranked by State. Accessed March 24, 2014. http://www.eia.gov/state/seds/data.cfm?incfile=/state/seds/sep_sum/html/rank_use.html&sid=US
- [25] Wienold, J., and Christoffersen, J. (2006). "Evaluation Methods and Development of a New Glare Prediction Model for Daylight Environments with the Use of CCD Cameras." *Energy and Buildings*, 38: 743–75.

Comparison of Artificial Neural Networks and Support Vector Machines for Weigh-In-Motion Based Truck Type Classification

Yueren Wang

Rinker School, College of Design, Construction and Planning,
University of Florida
Gainesville, United States
ywang016@ufl.edu

Ian Flood

Rinker School, College of Design, Construction and Planning,
University of Florida
Gainesville, United States
flood@ufl.edu

Raja R.A. Issa,

Rinker School, College of Design, Construction and Planning,
University of Florida,
Gainesville, United States
raymond-issa@ufl.edu

Abstract – the paper develops and compares a comprehensive range of configurations of artificial neural networks and support vector machines for solving the truck classification by weigh-in-motion problem. A local scatter point smoothing schema is also demonstrated as a means of selecting an optimal set of design parameters for each model type. Three main model formats are considered: (i) a monolithic structure with a one versus all strategy for selecting truck type; (ii) an array of sub-models each dedicated to one truck type with a one versus all truck type selection strategy; and (iii) an array of sub-models each dedicated to selecting between pairs of trucks. Overall, the SVM approach was found to outperform the ANN based models. The paper concludes with some suggestions for extending the work to a broader scope of problems.

Keywords – artificial neural network; empirical modeling; support vector machine; truck weigh-in-motion.

I. INTRODUCTION

Empirical modeling is concerned with the development of a representation of some aspect of a system based on data observed from that system or from an analog of that system. While empirical modeling is widely used in fields such as business, engineering, and science, it poses many challenges that need to be overcome before its full potential can be realized [1]. The ability to classify moving trucks based on the strain envelopes they induce on bridge girders (termed weigh-in-motion) has been identified as an example problem rich in the issues challenging empirical modeling, and thus provides a good point of reference in developing and evaluating this modeling technique [1].

To date a variety of empirical modeling techniques have been applied to the problem of truck-type classification from bridge weigh-in-motion data. Supervised learning methods such as artificial neural networks (ANNs) have been studied extensively in this regard. In a comprehensive study by

Gagarin et al. [2] an ANN was used to map directly from a stream of strain readings measured on a bridge girder resulting from a truck crossing event to an output array where each element represented a truck type; the output element with the strongest response represented the ANN's determination of the truck-type. This approach demonstrated reasonable accuracy in classifying trucks, around 857% correct classifications, but it did not perform well for truck types with similar wheel configurations. A later study [3] attempted to improve performance by using a type of Hamming Network (a binary classifier) with novel presentation formats. While both types of ANN showed some promise as classifiers, they suffered from a more fundamental problem associated with direct-mapping solutions, namely that each classifier can only work for a single bridge configuration. Each new bridge requires collection of a new set of training patterns followed by training of a new classifier. Moreover, it is not possible to extend the scope of application of these classifiers by including additional input variables to describe a bridge's configuration since this would lead to a geometric explosion in the number of training patterns required.

In an attempt to get around this problem, a radically different approach was considered by Vala et al. [4] based on genetic algorithms (GAs). In this study the GA was used to evolve a truck-type configuration that could best explain the strain envelope, and a numeric structural model of the bridge was used to evaluate the evolving solution. While the GA approach was, in principle, more flexible than the ANN classifier in terms of its scope of application, the study was only preliminary and its ability to estimate truck-type satisfactorily was not conclusive.

Work is on-going at the University of Florida developing more flexible methods of classifying truck-types from weigh-in-motion data as part of a broader line of study

developing empirical modeling methods. However, to conclude the work on direct-mapping it was decided to compare the use of ANNs to that of support vector machines (SVMs), an empirically based classification device that has developed a growing interest since the mid-1990's with success in a diverse range of applications including facial detection [5], CT image recognition [6], and market power estimation [7].

This paper is specifically concerned with comparing the performance of SVM and ANN based approaches to weigh-in-motion based truck-type classification.

II. MODELING APPROACH

A. Truck Types and Bridge Properties

A total of nine different truck types were considered in this study, representing those most frequently operating on US highways, as adopted in earlier research [8] and outlined in Figure 1. For each truck type there is a range of values defining its axle loadings and axle spacings as summarized in Table I.




Figure 1. Nine truck types used in this paper (adapted from Gagarine et al. [1])

TABLE I: AXLE LOAD AND SPACING RANGE OF NINE TRUCK TYPES (adapted from Gagarine et al. [1])

Truck Type	Axle Loads (KN)					Axle Spacings (m)				
	1	2	3	4	5	6 1 and 2	2 and 3	3 and 4	4 and 5	5 and 6
1	13.3-53.4	8.8-80.1				2.74-6.10				
2	13.3-53.4	8.8-80.1	8.8-80.1			2.74-6.10	1.22			
3	13.3-53.4	8.8-80.1	8.8-80.1			2.74-4.98	5.49-11.6			
4	13.3-53.4	8.8-80.1	8.8-80.1	8.8-80.1		2.74-5.49	1.22	1.22		
5	13.3-62.3	8.8-71.2	8.8-71.2	8.8-80.1		2.74-6.10	1.22	6.10-11.6		
6	13.3-53.4	8.8-80.1	8.8-80.1	8.8-80.1		2.74-5.49	6.10-11.6	1.22		
7	13.3-53.4	8.8-71.2	8.8-71.2	8.8-80.1	8.8-80.1	2.74-6.10	1.22	6.10-11.6	1.22	
8	13.3-53.4	8.8-71.2	8.8-71.2	8.8-80.1	8.8-80.1	2.74-6.10	1.22	6.10-11.6	1.22	1.22
9	13.3-53.4	8.8-80.1	8.8-80.1	8.8-80.1	8.8-80.1	2.74-5.49	5.49	3.05	5.49	


The bridge type considered was 100 meters in length, single span, simply supported, with a single lane. The bridge was treated as a rigid beam and the study assumed no dynamic effects on the structure. Single truck crossing events only were considered.

B. Model Structure


A truck crossing event was represented as an array of bending moments induced at mid-span, while the type of truck inducing the bending moments was indicated across an array of outputs.

Three different model formats were considered as illustrated in Figure 2. The first model format comprised a


monolithic model that mapped directly from the input array of bending moments to a set of 9 outputs representing the different truck types. Each output represented a different truck type and was capable of generating a value between 0.0 and 1.0. The output that generated the closest value to 1.0 in response to a set of bending moments was assumed to identify the truck type. This format was only adopted for the ANN model since SVMs cannot include more than one output.



Model Format 1



Model Format 2



Model Format 3

Figure 2. Three model formats adopted for the study

The second model format shown in Figure 2 comprised a set of 9 sub-models, each dedicated to a single truck type. A single array of bending moments was shared as input, and each sub-model had a single output capable of generating a value between 0.0 and 1.0. As for the first model format, the output that generated the closest value to 1.0 was assumed to identify the truck type. This format was adopted for both the ANN and SVM models.

The third model format shown in Figure 2 comprised 36 sub-models, each dedicated to selecting between a pair of truck types (there being 36 permutations of truck pairs in total). A single array of bending moments was shared as input. Each output would select between a pair of trucks. For example, sub-model 1 was dedicated to comparing truck

types 1 and 2; an output of 0 would indicate truck type 1 and an output of 1 would indicate truck type 2. The truck type with the most selections across the output array was assumed to be the truck type crossing the bridge.

C. Truck Crossing Simulation

The data used for training and validation of the models was based on a random selection of truck configurations (based on Table I). Data was generated by simulating the passage of a truck over the bridge. The bending moment induced at the mid-span of the bridge, m , was calculated during the truck crossing event using a 50 Hz sample rate as indicated in Figure 3.




Figure 3. The simulation of a type 3S truck passing a bridge and its bending moment envelope

Each simulated truck crossing event was used to generate a single input to output pattern to be used for training or validation of the models/sub-models. For model formats 2 and 3, each sub-model was trained independently. Each pattern comprised 626 inputs representing the bending moments induced by the truck crossing event, and an array of binary outputs used to indicate truck type. A total of 900 input patterns were generated (100 for each truck type) and the corresponding outputs were tailored to match the operation of each model/sub-model. For each pattern, the axle loads and spacings were selected using a uniformly distributed random variate with values ranged between the limits listed in Table I.

III. PRELIMINARY MODEL DEVELOPMENT

Training of both the ANN and SVM models requires preselection of certain training parameters, the values of which can significantly affect the performance of the model. In addition, since the initial input arrays have a high dimension (626 values) Principal Component Analysis (PCA) was used to prune this number down to something more manageable.

The architecture of the ANNs adopted for this study was the popular feedforward layout. Two ANN variants were considered, one with a single hidden layer of 600 neurons and a second with two hidden layers of 300 hidden neurons each. This provided a total of four ANNs, two using format 1 (Figure 2) and two using format 2. All ANNs used the sigmoidal activation function. Future work may use a sensitivity analysis to assess the dependence of model performance on the number of hidden neurons. The training algorithm used for the ANN was error backpropagation, a gradient descent technique, and was implemented within the MATLAB R2012a environment.

Since the ANN backpropagation approach requires a careful selection of the step size to ensure convergence and

acceptable training quality, this study tested a range of learning rates from 0.01 to 0.1 in intervals of 0.01.

For the SVM, the kernel function adopted was the Radial Basis Function due to its popularity. The SVM’s kernel function also requires a careful selection of its scaling value, and so a range of values were tested from 3.60 to 3.70 with intervals of 0.01.

For pruning the number of input variables, a range in the array size was considered from 10 to 55 in steps of 5, using principal component analysis (PCA) to select the most significant inputs in each case.

A. ANN Development, Model Format 1

Training of an ANN was allowed to progress until 300 training epochs had been completed. Training used a random selection of 80% of the 900 pattern data set. The remaining 20% of the patterns were used for validating the resultant ANN. Model development was repeated for the range of learning rates and input vector sizes outlined above, providing 100 training trials. These experiments were repeated 10 times, each occasion using a different set of 900 patterns. The performance of the ANNs was measured as the portion of the validation patterns correctly classified. This was averaged over the 10 repetitions of the experiment. Local scatterplot smoothing (LOESS) was adopted (with a span value of 0.15) to find the location of the peak performance and thus the optimal values for the number of input variables and the learning rate. Figure 4 shows the results of these experiments, plotting the proportion of correct classifications (from 0.0 to 1.0, color coded) against the number of inputs (PCA derived) and the learning rate. The optimal values were found to be 32 for the number of inputs and 0.088 for the learning rate. The corresponding R square value was 0.8249, indicating an acceptable description of parameter relationship using the smoothing method.




Figure 4. LOESS regression on the performance of ANN Format 1 with one hidden layer

The experiment was repeated this time using an ANN with two hidden layers, with the resultant performance surface shown in Figure 5. The optimum set-up was found to be 35 inputs and a learning rate of 0.0676. The R squared value was again acceptable at 0.8926.




Figure 5. LOESS regression on the performance of ANN Format 1 with two hidden layers

B. ANN Development, Model Format 2

The set of experiments described in section A above were repeated but this time using the model format 2 shown in Figure 2, that is, the system comprising 9 sub-models.

For the one hidden layer ANN, the number of hidden neurons in each sub-model was 67 for the one hidden layer ANN, giving 603 hidden neurons in total. For the two hidden layer ANN, 33 hidden neurons were included in each layer of each sub-model providing a total of 594 hidden neurons.

Figure 6 shows the results of these experiments, as before plotting the proportion of correct classifications against the number of PCA selected inputs and the learning rate.




Figure 6. LOESS regression on the performance of ANN Format 2 with one hidden layer

The optimal values were found to be 39 for the number of inputs and 0.06520.088 for the learning rate. The corresponding R square value was acceptable at 0.8304.

Similarly Figure 7 plots the LOESS smoothed surface representing the proportion of correctly classified validation patterns for the two hidden layer ANN based on model format 2. In this case the optimal values were found to be 40 for the number of PCA selected inputs and 0.0736 for the learning rate, with an R square value of 0.8304.




Figure 7. LOESS regression on the performance of ANN Format 2 with two hidden layers

C. SVM Development, Model Formats 2 and 3

The next set of experiments concerned development of the SVM models, the first using the one versus all strategy (model format 2, Figure 2) and the second using the one versus one strategy (model format 3, Figure 2).

Figure 8 shows the LOESS smoothed performance surface for the SVM based on model format 2. This plots the proportion of correctly classified trucks in the validation data set versus the number of PCA selected inputs and the Radial Basis Function kernel scaling value. The optimal values were found to be 19 for the number of inputs and 3.6677 for the kernel scaling value, with an R square equal to 0.9984.

Similarly the SVM system based on model format 3 is plotted in Figure 9. The optimal values were 18 for the number of PCA selected inputs and 3.6960 for the kernel scaling value, with an R square value of 0.9999.

It can be seen from Figures 8 and 9 that the scaling value did not play an important role in determining the performance of the resultant SVM. However, the number of PCA selected inputs was clearly very important for both SVM modeling approaches.




Figure 8. LOESS regression on the performance of SVM Format 2 with a one versus all strategy

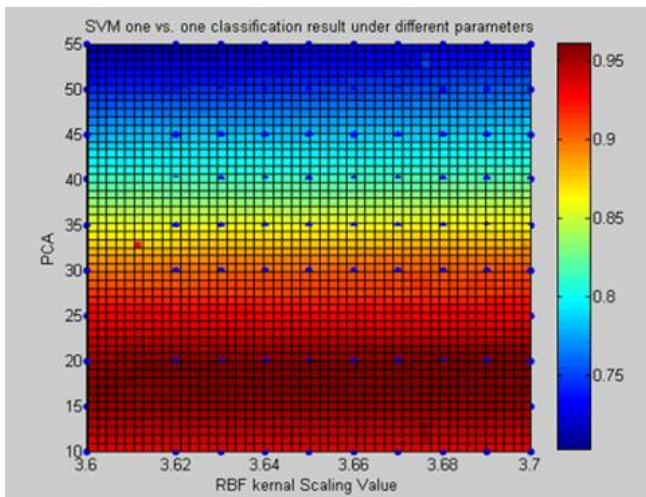


Figure 9. LOESS regression on the performance of SVM Format 2 with a one versus one strategy

IV. MODEL EVALUATION

The optimal values determined for the number of inputs and the learning rate or kernel value were used to develop the final versions of each of the 6 model forms. The performance of each of these models is compared in Figure 10 in terms of their ability to correctly classify the validation patterns. All 1,800 validation patterns generated across the 10 data sets were used for this purpose.

Clearly, the results demonstrate that the SVM models outperform the ANN models. Of the SVM models, the one versus all strategy was found to slightly outperform the one versus one strategy.

For the ANN models, the structure comprising 9 sub-models significantly outperformed the monolithic ANN structure. Having individual sub-models may allow the system more flexibility in learning the pattern of a specific

truck type and therefore improve the accuracy of the entire model.

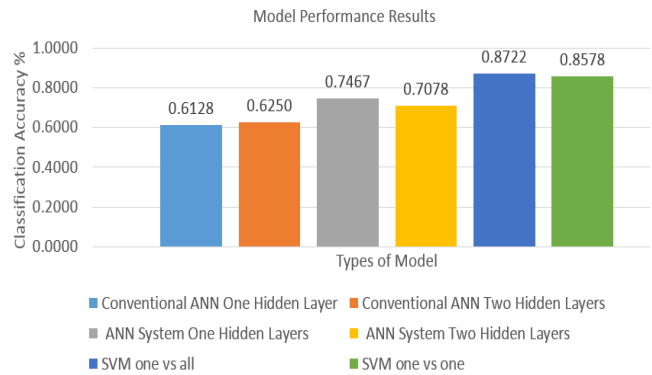


Figure 10. Comparison of optimal model performances for 1,800 validation patterns

For the ANN models, the number of hidden layers did not appear to have a significant impact on classification performance.

Figure 11 provides an analysis of the misclassified truck patterns for the single hidden layer ANN, model format 1, for the validation patterns. The blue arrows indicate the number and direction of the misclassifications. It can be seen from this figure that the misclassifications between truck type 7, 8 and 9 and truck type 1, 2 and truck type 3, 5 and 5, 6 contributed to the majority of the misclassification instances. As might be expected, it is also apparent from this that the misclassifications tended to occur between trucks with similar axle configurations.




Figure 11. Analysis of truck misclassifications for the single hidden layer ANN, model format 1 (see Figure 2)

V. CONCLUSION AND FUTURE WORK

The study developed and compared the performances of 6 ANN and SVM based truck classifiers using weigh-in-motion data. The optimal versions of each model were determined using a LOESS based empirical modelling parameter selection schema. The results indicated that the SVM models significantly outperformed the ANN models in terms of the number of correct truck classifications.

Future work should be concerned with developing models that are extendable to a wider range of problems, including bridges of different lengths, span configurations, and numbers of lanes, as well as situations involving multiple truck crossing events. Such models should also be able to estimate truck parameters such as axle loadings and spacing. A challenge is to achieve this while circumventing the problem of a geometric increase in the number of training patterns with respect to the number of variables required to describe the problem.

ACKNOWLEDGMENT

This work was funded by the University of Florida student fellowships program.

REFERENCES

- [1] I. Flood, and RRA. Issa, "Empirical Modeling Methodologies for Construction," *J. Constr. Eng. Manage.*, vol. 136, pp. 36-48, August, 2009.
- [2] N. Gagarin, I. Flood, and P. Albrecht, "Weighing trucks in motion using Gaussian-based neural networks," *Proc. Int. Joint Conf. on Neural Networks*, Baltimore, Md., 2, 1992, pp. 484-489.
- [3] I. Flood, "Developments in Weigh-in-Motion Using Neural Nets," *Proc. Comput. Civ. And building Eng.*, 2000, pp. 1133-1140, doi: 10.1061/40513(279)148.
- [4] G. Vala, I. Flood and E. Obonyo, "Truck Weigh-in-motion using reverse modeling and genetic algorithms," *Proc. Comput. Civ. Eng.*, 2011, pp. 219-226, doi: 10.1061/41182(416)27.
- [5] E. Osuna, R. Freund and F. Girosi, "Train support vector machines: an application to face detection," *Proc. IEEE Computer Vision and Patten Recognition*, June, 1997, pp. 130-136, doi: 10.1109/CVPR.1997.609310.
- [6] Q. Wang, W. Zhu and B. Wang, "Three-Dimensional SVM with latent variable: an application for detection of lung lesions in CT images," *J. Med. Syst.*, 2015, pp. 39-171.
- [7] SP. Karthikeyan, IJ. Raglend and KS. Kumar, "Application of SVM as classifier in estimating market power under deregulated electricity market," *Proc. ICPERES*, Nov 2014 pp. 1309-1317, doi: 10.1007/978-81-332-2119-7_127.
- [8] N. Gargarin, I. Flood, and P. Albrecht, "Computing truck attributes with artificial neural networks," *J. Comput. Civ. Eng.*, vol. 8, pp. 179-200, April 1994.

Modeling Interactive Digital TV Users Behavior

Samuel da Costa Alves Basilio

Centro Federal de Educação Tecnológica de Minas Gerais (CEFET-MG)

Leopoldina - MG - Brazil

Email: samuelbasilio@leopoldina.cefetmg.br

Abstract—The performance evaluation required in the proposition of new large scale distributed systems usually faces the challenge of correct characterization of the load that is imposed on them. In the case of proposals in the area of Digital Television (TV), including terrestrial, cable, Satellite and IPTV, obtaining such a characterization from real deployment scenarios has proved to be a very difficult accomplishment due to the impediment experimental access to these distribution networks in operation. Thus, a researcher usually uses simulations that impose work loads crudely approximated, or even fictitious oversized to your system, leading uncertainty to potential service providers regarding the optimal sizing of the equipment required. Here, we present a mathematical model of simple implementation, able to represent the behavior of users of Digital TV. The model can be parameterized to represent different states of behavior about the system to be simulated, and thus adapt to various scenarios of interest. We also show how this model was used in the performance analysis of a proposed service provider.

Keywords—Digital TV; Model; User Behavior; Interaction; Simulation.

I. INTRODUCTION

When dealing with research in Digital TV, we often face a challenge in time to validate the software developed. As this environment involves millions of users, the developed software have to take into account concepts such as scalability, availability and performance, but the effort to evaluate these features may not be trivial. Undoubtedly, the best way to improve, correct errors and check the software requirements is applying it in a real environment, even if with a restricted group of users. In the case of Digital TV environment, this done has proven to be difficult due to the fact that the real environment is also commercial. Additionally, experimental tests are not generally accepted. Although the use of a limited number of users is not ideal, in most cases, it is the resource available for researchers and software developers. This feature can be presented as a good solution when we want to evaluate interface, functionality, among others factors. However, when we want to evaluate such criteria as scalability and availability, it is not enough. In this case, a resource that can be used efficiently, cheaply and reliably is simulations. Nowadays, computational resources are relatively inexpensive and can be used to simulate environments with a large number of users. However, to use this resource, we need a reliable model, which should represent as closely as possible the environment behavior. The challenge to use this resource is in the development and implementation of the model to be simulated. It is necessary that the model developer closely observes the behavior of all the environment elements and abstract them in a simple model. The model must have a

balance between fidelity to the real environment and simplicity of implementation. Moreover, a behavior model must suit to its purpose. We should not use a desktop user behavior model to simulate a mobile phone environment. Similarly, we should not use a web user behavior model to simulate an interactive digital TV load.

Knowing these problems, we present in this paper a mathematical model that can easily be implemented and even so is faithful enough to reality. Furthermore, with the data from an experiment where 27 viewers had their interactions with TV capture along with a TV audience survey data from a local statistics research institute, we show a parameterization example of the proposed model. The presented model may also be used in other contexts, such as targeted advertising and social context analysis, audience measurement, among others.

This paper is structured as follows. Section II discusses some related work; Section III is dedicated to mathematical model of the behavior of users of Interactive Digital TV. Section IV exemplifies the model instantiation for a specific case. Finally, Section V shows the achieved objectives and discusses future work.

II. RELATED WORK

The attempt to model the behavior of media consumption system users is not a new work as seen in [1], where Branch et al. characterize and model the behavior of users of their video on demand system. But, as new ways of interaction appear, as well as new technologies and new systems of media consumption, existing models are often not adequate. Alvarez et al. [2] shows an architecture for audience measurement, a model for data consumption and some metrics to quantify the impact of consumer. This metric is calculated in a similar way in [3].

Along with the user behavior model, some work show a characterization of this behavior in a real environment. That is the case of [4], which characterizes the behavior from a system with millions of users. An interesting metric presented in this paper is the session time that has a paramount importance when simulating the behavior of various users over long periods of time.

An important point is cited in [5], where a synthetic load generator is shown. In this paper, Costa et al. cite the need for heterogeneity in load generators because many are reported in the literature, but most work only a group-specific data, such as educational. The work presented by Qiu et al. [6] also has a generator of synthetic load but only focusing on the Internet Protocol Television (IPTV) environment. Nevertheless, this work has an advantage because it used data from a real system with millions of users.

III. INTERACTIVE DIGITAL TV USERS BEHAVIOR MODEL

Aiming to support possible future simulations, we present in this session a Markov model of the Interactive Digital TV users interactions. This work comes to fill a gap in the digital TV research. It is really hard to evaluate points like scalability, availability and performance at this area because until now there is no model to represent the interactive digital TV user behavior.

To define the Interactive Digital TV user behavior, we must consider every possibility of viewer interaction. Below, we define the possible states its transitions. Figure 1 shows the states and the possible transitions between these states:




Figure 1. Markovian model of the users interactions.

- E_i , when an interactive application is running;
- E_n , when only the broadcast video is running;
- E_f , when the viewer turns off the digital receiver;
- E_g , when a native application is running.
- p_{if} , probability of the digital receiver to be unplugged since an interactive application was running;
- p_{ii} , probability of the viewer to continue the execution of an interactive application;
- p_{ig} , probability of the viewer start a native application as an interactive application was running (pausing the execution of the interactive application);
- p_{in} , probability of a running interactive application to be closed or the viewer change the channel (ending the execution of the application);
- p_{ni} , probability of an interactive application to be started since no other application was running;
- p_{nn} , probability of the viewer does not start any application;
- p_{nf} , probability of the receiver to be turned off without any application be running;
- p_{ng} , probability of the viewer start a native application;
- p_{gn} , probability of the viewer select a channel through a native application (such as the electronic program guide);
- p_{gg} , probability of the viewer to continue running a native application;

- p_{gf} , probability of the receiver be turned off with a native application running;
- p_{gi} , probability of a native application to be terminated after being initiated when an interactive application was running (recovering the state of the interactive application that was then paused);
- p_{ff} , probability of the receiver be off.

We differentiate the state where native applications (we consider native applications those that are specific to the digital receiver, from the factory or installed later as the electronic program guide.) are running from the state where interactive applications in general are because depending on data capturing approach used it may not be possible to capture the interactions of native applications [7]. As these native applications are a resource provided by the digital receiver, if the data capture approach by interactive applications is used, user interaction with these native applications can not be obtained. However, if the data capture approach used is by middleware extension, all interactions may be obtained, the ones with interactive applications such as the ones with native applications. In both approaches, the channel change interactions can be obtained.

Note that it is simple to specialize the presented model to any specific case. For example, if we wanted to specify at the initial model an arbitrary interactive application, it would only need to add to each state E_x of the application the probabilities p_{xy} and p_{yx} , where m is the number of states of the chain $E_x = \{E_{a1}, E_{a2}, \dots, E_{am}\}$, $p_{xy} = \{p_{xn}, p_{xg}, p_{xi}, p_{xf}\}$ and $p_{yx} = \{p_{nx}, p_{gx}, p_{ix}, p_{fx}\}$. We also have to remove the state E_f , which represents the final state of the application chain, as this state is reached at some point when any of the probabilities p_{xy} happen.

To illustrate the specialization of the model Figure 2(a) shows the model of an arbitrary interactive application that has two active states, E_{a1} and E_{a2} , and a final state E_f . To extend the original model we remove the state E_f from the application model and add the states E_{a1} e E_{a2} at the original model, keeping the probabilities that relate the states E_{a1} and E_{a2} and adding the probabilities that relate E_{a1} and E_{a2} with E_i , E_g , E_n e E_f . Note that the estate E_i that at the original model represented all interactive application, is now called E'_i and represente all others interactive applications. If all applications are represented individually, the state E'_i can be removed from the model. The same reasoning can be applied at the state E_g . Figure 2(b) shows the final result.

A. Expected Session Time


An interesting metric that we can obtain from this model is the expected session time. The session time of a viewer can be estimated by calculating the number of steps required to reach the state E_f . Below is the formalism to calculate this time, taken from [8].

Let $(X_n)_{n \geq 0}$ be a Markov chain with transition matrix P . The *hitting time* of a subset A of I is the random variable $H^A : \Omega \rightarrow \{0, 1, 2, \dots\} \cup \{\infty\}$ given by:

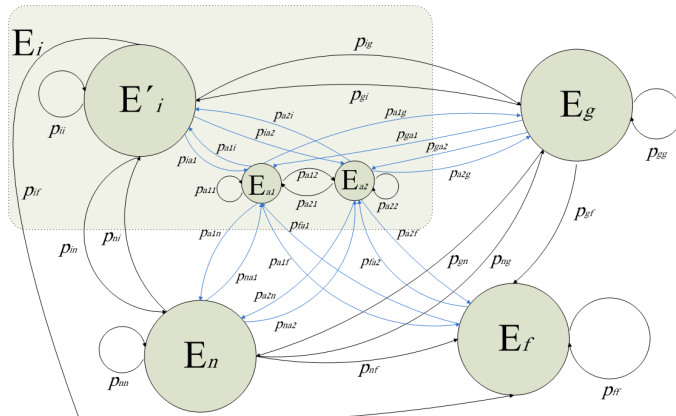
$$H^A(\omega) = \inf\{n \geq 0 : X_n(\omega) \in A\}$$

where the infimum of the empty set \emptyset is ∞ . The probability starting from i that $(X_n)_{n \geq 0}$ ever hits A is then:

$$h_i^A = P_i(H^A < \infty).$$



(a) Model of an arbitrary application



(b) Specialized model

Figure 2. Model specialization.

When A is a close class, h_i^A is called the *absorption probability*. The mean time taken for $(X_n)_{n \geq 0}$ to reach A is given by:

$$k_i^A = E_i(H^A) = \sum_{n < \infty} nP(H^A = n) + \infty P(H^A = \infty).$$

For our case, we have $i = \{E_n\}$ and $A = \{E_f\}$. It is clear that $k_{E_f} = 0$. Starting at E_n and performing one step, with probability p_{ni} we reach E_i , with probability p_{ng} , we reach E_g , with probability p_{nf} we reach E_f , and with probability p_{nn} we keep at E_n . Then:

$$k_{E_n} = 1 + p_{ni}k_{E_i} + p_{ng}k_{E_g} + p_{nf}k_{E_f} + p_{nn}k_{E_n}$$

$$k_{E_n} = \frac{1 + p_{ni}k_{E_i} + p_{ng}k_{E_g}}{(1 - p_{nn})}$$

The 1 appears because we count the time for the first step. Similarly, for k_{E_g} and k_{E_i} , we have the following system that solved gives the expected session time:

$$\begin{cases} k_{E_n} = \frac{1 + p_{ni}k_{E_i} + p_{ng}k_{E_g}}{(1 - p_{nn})} & (1a) \\ k_{E_g} = \frac{1 + p_{gi}k_{E_i} + p_{gn}k_{E_n}}{(1 - p_{gg})} & (1b) \\ k_{E_i} = \frac{1 + p_{ig}k_{E_g} + p_{in}k_{E_n}}{(1 - p_{ii})} & (1c) \end{cases}$$

B. Expected Number of Interactions in a Given Period

It is important to note that the probabilities p_{ii} , p_{gg} and p_{nn} , represent the probability of the viewer continue in its current state, running an interactive application, a native application or not running any applications, respectively. If at some step in the execution of this model any of these probabilities happen, an event of interaction can be generated or not. If at the state E_i the probability p_{ii} happen, the viewer continues with the interactive application running. In the time period covered by this step, he may or may not have interacted with the application. There is, then, at this point, a probability of generating a user interaction event. We call this probability *interaction rate at the state*, being noted as pe_{ii} . The same reasoning can be applied to p_{gg} and p_{nn} , and we call pe_{gg} and pe_{nn} their respective interaction rates at the state.

With that, another important metric that we can calculate is the expected number of interactions that a viewer shall perform in a given period. Consider that whenever there is a transition between the states E_i , E_g and E_n an event I is generated. This event is an interaction. Also, when a transition where there is no state change occurs p_{ii} , p_{gg} and p_{nn} an event is generated with probability pe_{ii} , pe_{gg} and pe_{nn} respectively. Thus, the probability that for each state E_i , E_g and E_n , we generate an interaction event is:

$$I = \begin{cases} p_{ni} + p_{ng} + p_{nn}pe_{nn} & , \text{ if the current state is } E_n \\ p_{in} + p_{ig} + p_{ii}pe_{ii} & , \text{ if the current state is } E_i \\ p_{gi} + p_{gn} + p_{gg}pe_{gg} & , \text{ if the current state is } E_g \end{cases} \quad (2)$$

Observing function (2), we note that the probability of generate an event is dependent of the current state. Therefore, owe calculate for each step k the probability of being in each state E_i , E_g and E_n . Let $p_{E_i}^{(k)}$, $p_{E_n}^{(k)}$ and $p_{E_g}^{(k)}$ be the probabilities in step k of being at state E_i , E_n and E_g , respectively, these probabilities are given by the recursive functions system:

$$\begin{cases} p_{E_i}^{(k)} = p_{gi}p_{E_g}^{(k-1)} + p_{ni}p_{E_n}^{(k-1)} + p_{ii}p_{E_i}^{(k-1)} \\ p_{E_n}^{(k)} = p_{gn}p_{E_g}^{(k-1)} + p_{nn}p_{E_n}^{(k-1)} + p_{in}p_{E_i}^{(k-1)} \\ p_{E_g}^{(k)} = p_{gg}p_{E_g}^{(k-1)} + p_{ng}p_{E_n}^{(k-1)} + p_{ig}p_{E_i}^{(k-1)} \end{cases} \quad (3)$$

This way, to calculate the total amount of generated events $I^{(k)}$ in k steps, we must sum for each step k the probabilities of generate an event in each state. Formally:

$$I^{(k)} = \sum_{j=1}^k p_{E_n}^{(j)} I + p_{E_i}^{(j)} I + p_{E_g}^{(j)} I \quad (4)$$

C. Expected Number of Applications Execution

We also can estimate the expected number of applications execution. That metric is represented by the expected number of visits to a state. If we extend our model so that each interactive application is represented by an individual state, we even shall be able to estimate the expected execution number of an specific application.

Let,

$$N_j(k) = \sum_{m=1}^k I(X_m = j) \quad (5)$$

the number of visits to state j during times 1 to k . And let,

$$G_{ij}(k) = E(N_j(k)|X_0 = i) = \sum_{m=1}^k P_{ij}^{(m)} \quad (6)$$

the average time of visits to state j during times 1 to k , starting at i . If we consider j as our state E_i and i as our states E_n and E_g , the expected number of interactive applications execution in k time is given by

$$G_{E_n E_i} + G_{E_g E_i} \quad (7)$$

IV. MODEL INSTANTIATION

As part of the work [7], an experiment was conducted in which 27 volunteers watched television and had their interactions captured. All analyzed viewers watched the available programming for 15 minutes. Using these data we conducted an instantiation of the presented model in order to illustrate its use.

Table I shows the quantification of used data and Table II the probabilities found. In Table I we consider as bounce rate the percentage between the total of executions and total executions where after start the application, the user closed it without making any other interaction.

TABLE I. SUMMARY OF USED DATA.

USED DATA	VALUE
INTERACTIONS TOTAL	2241
CHANNEL CHANGE TOTAL	355
INTERACTIONS WITH INTERACTIVE APPLICATIONS	1886
TOTAL EXECUTIONS OF INTERACTIVE APPLICATIONS	139
AVERAGE TIME OF APPLICATIONS EXECUTION	45.02 secs
BOUNCE RATE	45.32%
BIGGEST NUMBER OF INTERACTIONS OF A SINGLE USER IN A MINUTE	57
USERS TOTAL	27
EXPERIMENT TOTAL TIME	6 hours and 45 mins

For definition of these data, the experiment time was discretized in seconds. From this we calculated the number of seconds that a viewer was in each state E_i , E_g and E_n . Knowing how long each viewer spent in each state and transitions between the states it was possible to get the data from Table II. In the experiment, every session was started in E_n state.

With the probabilities it is possible to calculate the expected session time replacing the obtained values in equations (1a), (1b) e (1c).

$$\left\{ \begin{array}{l} k_{E_n} = \frac{1 + \frac{118}{13895}k_{E_i} + \frac{35}{13895}k_{E_g}}{\left(1 - \frac{13728}{13895}\right)} \quad (8a) \\ k_{E_g} = \frac{1 + \frac{18}{2428}k_{E_i} + \frac{13}{2428}k_{E_n}}{\left(1 - \frac{2391}{2428}\right)} \quad (8b) \\ k_{E_i} = \frac{1 + \frac{2}{3477}k_{E_g} + \frac{133}{3477}k_{E_n}}{\left(1 - \frac{3340}{3477}\right)} \quad (8c) \end{array} \right.$$

Solving the system we find:

$$k_{E_n} = \frac{19903979}{22128} \approx 899.49290 \quad (9)$$

TABLE II. VALUES FOR THE PROBABILITIES OF OUR MODEL.

PROBABILITY	VALUE
p_{ii}	$\frac{3340}{3477}$
p_{in}	$\frac{133}{3477}$
p_{ig}	$\frac{2}{3477}$
p_{if}	$\frac{2}{3477}$
p_{gi}	$\frac{18}{2428}$
p_{gg}	$\frac{2391}{2428}$
p_{gn}	$\frac{13}{2428}$
p_{gf}	$\frac{6}{2428}$
p_{nn}	$\frac{13728}{13895}$
p_{ni}	$\frac{118}{13895}$
p_{ng}	$\frac{35}{13895}$
p_{nf}	$\frac{14}{13895}$
$p_{e_{nn}}$	$\frac{1036}{13895}$
$p_{e_{ii}}$	$\frac{1126}{3477}$
$p_{e_{gg}}$	$\frac{482}{2428}$

We are also able to calculate the expected number of interactions in 15 minutes. As every session started in E_n state, we have that $p_{E_i}^{(1)} = p_{E_g}^{(1)} = p_{E_f}^{(1)} = 0$ and $p_{E_n}^{(1)} = 1$. Calculating the probabilities for each step k of the system (3) and replacing Table II probabilities in the sum (4), we have:

$$\begin{aligned} I^{(900)} &= \sum_{j=1}^{900} p_{E_n}^{(j)} \left(\frac{118}{13895} + \frac{35}{13895} + \frac{13728}{13895} \frac{1036}{13895} \right) + \\ &\quad p_{E_i}^{(j)} \left(\frac{133}{3477} + \frac{2}{3477} + \frac{3340}{3477} \frac{1126}{3477} \right) + \\ &\quad p_{E_g}^{(j)} \left(\frac{18}{2428} + \frac{13}{2428} + \frac{2391}{2428} \frac{482}{2428} \right) \\ I^{(900)} &\approx 82.8720 \quad (10) \end{aligned}$$

We will also estimate the amount of interactive applications executions using (7), (3) and data from Table II. For this case, we have that

$$\begin{aligned} G_{E_n E_i}(k) &= \sum_{m=1}^k P_{E_n E_i}^{(m)} = \sum_{m=1}^k p_{E_n}^{(m-1)} p_{ni} \approx 3.45 \\ G_{E_g E_i}(k) &= \sum_{m=1}^k P_{E_g E_i}^{(m)} = \sum_{m=1}^k p_{E_g}^{(m-1)} p_{gi} \approx 0.51 \end{aligned}$$

so:

$$G_{E_n E_i} + G_{E_g E_i} \approx 3.96 \quad (11)$$

Observing the results (9) and (10), we verified the accuracy of our calculations and probabilities, as the experiment lasted

15 minutes (900 seconds), this was the average time and soon, the expected session time. Also the interactions total is 2241 and the users total 27, as seen at Table I, so the expected number of interactions is 83 interactions per user with a 15 minutes session. (11) shows a larger error. As the total executions of interactive applications is 139 and the users total 27, the expected amount of interactive applications executions per user should be closer to 5. This larger error happens because we ignore the chance of an application to be closed and opened at an interval shorter than 1 second, this way the current state was E_i and has not changed. For simplicity, this possibility is not considered in our calculations.

A. Model Use to Generate Synthetic Load

Also as part of the work presented in [7], we conducted a load test on a real server which implemented an audience and interaction analysis service provider. This service provider is constantly receiving data captured at the viewers digital receiver and store this data in a relational database. The purpose of this load test was to illustrate the use of the presented model. In this case to generate synthetic load for scalability and performance testing. The implementation of this server was made using a virtual machine with the settings shown in Table III.

TABLE III. SEVER CONFIGURATION.

SYSTEM PART	TECHNICAL DETAILS
CPU	Intel Xeon 2.0GHz
RAM	1GB
Cores	1 or 2
Swap	2GB
Operating System	Ubuntu 12.04 LTS
HD	14GB
Web Server	Apache Tomcat/6.0.35
Data Base	MySQL: 5.5.31

In this test, we used the probabilities presented in Table II. We have further increase the probabilities values pe_{ii} , pe_{gg} and pe_{nn} to 1. This way, we assure that every second an interaction would be generated and sent to our server. For the test we send 10000 requests to the server at each session. We started sending batches of 100 simultaneous requests and check the response time. We have been increasing the number of simultaneous requests to reach 500. To generate this amount of requests it takes about 150 to 800 model instances. We did this experiment twice using the same server, at the first time with a single core and at the second time with two cores. Figure 3 shows the result of this test.

As was to be expected, as we increase the number of simultaneous requests the response time for each request also increases. When we send more than 500 simultaneous requests our server starts to fail due to overload. It is also noteworthy that with a small increase in computational power, response time greatly improves on average 3.008 ms.

V. CONCLUSION

In this paper, we present a mathematical model of Interactive Digital TV user interactions. Despite being simple and easy to implement, this model is sufficiently faithful to reality and can be used for the most diverse simulations purposes. We also showed how it is possible to extend the model to more specific and detailed models.




Figure 3. Time taken to complete each request.

We still use the data obtained in a field experiment as input to the model. Thus, we exemplify the use of the model for the calculation of certain metrics and generating synthetic load.

As future work, we propose new extensions to the proposed model, increasing its specialization and complexity. Using data from validity statistical captures would also be very interesting, because with this, we could have a much more reliable numerical model. But, capturing these data is impossible considering the current audience and interactivity measurement approaches.

ACKNOWLEDGMENTS

The authors thank CEFET-MG (Brazil) for financial support for this work.

REFERENCES

- [1] P. Branch, G. Egan, and B. Tonkin, "Modeling interactive behaviour of a video based multimedia system," in Communications, 1999. ICC '99. 1999 IEEE International Conference on, vol. 2, 1999, pp. 978–982.
- [2] F. Alvarez, C. Martin, D. Alliez, P. Roc, P. Steckel, J. Menendez, G. Cisneros, and S. Jones, "Audience measurement modeling for convergent broadcasting and iptv networks," Broadcasting, IEEE Transactions on, vol. 55, no. 2, 2009, pp. 502–515.
- [3] F. Alvarez, D. Alliez, C. Martin, J. Menendez, and G. Cisneros, "Audience measurement technologies for user centric media," in Consumer Electronics, 2008. ISCE 2008. IEEE International Symposium on, 2008, pp. 1–4.
- [4] V. Gopalakrishnan, R. Jana, R. Knag, K. Ramakrishnan, D. Swayne, and V. Vaishampayan, "Characterizing interactive behavior in a large-scale operational iptv environment," in INFOCOM, 2010 Proceedings IEEE, 2010, pp. 1–5.
- [5] C. Costa, C. Ramos, I. Cunha, and J. Almeida, "Genius: a generator of interactive user media sessions," in Workload Characterization, 2004. WWC-7. 2004 IEEE International Workshop on, 2004, pp. 29–36.
- [6] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao, "Modeling user activities in a large iptv system," in Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference, ser. IMC '09. New York, NY, USA: ACM, 2009, pp. 430–441.
- [7] S. d. C. A. Basilio, M. F. Moreno, and E. Barrère, "Supporting interaction and audience analysis in interactive tv systems," in Proceedings of the 11th European Conference on Interactive TV and Video, ser. EuroITV '13. New York, NY, USA: ACM, 2013, pp. 23–30.
- [8] J. Norris, Markov Chains, ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.