Implementation of a Media Aware Network Element for Content Aware Networks

Dragoş S. Niculescu, Mihai Stanciu, Marius Vochin, Eugen Borcoci
Telecommunication Dept., ETTI,
University Politehnica of Bucharest
Bucharest, Romania
e-mail:{dniculescu,ms,mvochin,eugenbo}@elcom.pub.ro

Nikolaos Zotos
Institute of Informatics and Telecommunications
NCSR Demokritos
Athens, Greece
e-mail: nzotos@iit.demokritos.gr

***Abstract*—ALICANTE is a recently proposed architecture that enables a lightweight form of virtualization for the purpose of offering QoS to media streams across the internet. A content aware network (CAN) is a cross domain overlay, which is provisioned in advance in order to provide preferential treatment to media streams. While it uses legacy infrastructure, such as core IP/MPLS routers and provisioned links, it relies on a special border router, called MANE (Media Aware Network Element). This paper presents a modular implementation of such a network element using off-the-shelf hardware and open source software. The implementation uses Click modular router to implement flow classification, MPLS encapsulation and decapsulation, separation between virtual CANs, and enforcement of separation between networks. Based on incipient measurements in a virtual testbed, we show that the implementation imposes minor overheads over existing routing infrastructure.**

Keywords—Content-Aware Networking; Network Aware Applications; Quality of Services; Multimedia distribution, Future Internet; Media Aware Network Element

I. INTRODUCTION

One of the new paradigms of the Future Internet (FI) is "content orientation", which is supposed to improve the user experience related to the new digital multimedia services and networked media content. This trend is recognized also by the European commission, which defined the "Objective ICT-2009.1.5: Networked Media and 3D Internet" in the FP7 Call 4 [1][2]. In this call, new directions are defined as content-aware networks (CAN) and network-aware applications (NAA). This approach breaks (partially) the classic TCP/IP and OSI stack network neutrality and application-transport separation concepts. The challenge is to get better performance without losing modularity of the architecture. CAN-NAA means the capability of the overall system to adjust network resource allocation based on limited examination of the nature of the content, while network-awareness means to process and distribute the content, based on limited understanding of the network conditions. Dynamic optimization is desired, with policies taking into account the content and adaptation needs, the user contexts, requirements and social relational networks. The FI should enable multiple user roles, e.g., as content producer, user, or manager.

The work of this paper is a part of the effort inside of an European FP7 ICT research project, "MediA Ecosystem Deployment Through Ubiquitous Content-Aware Network Environments", ALICANTE, [3][4]. An innovative architecture is proposed, in order to deploy a new type of "Networked Media Ecosystem", which allows flexible cooperation between providers, operators, and users. Three interworking environments are defined: Network Environment (NE) modeling Network Providers, Service Environment (SE) related to Content and Service Providers and User Environment (UE) including all End-Users (see Figure 1). The validation of the project architecture and results will be done in a large-scale international pilot.

The above environments are nowadays present in real deployments, but actually the collaboration between them is weak or non-existent. The current architectures do not exchange content-based and network-based information between the network layers and upper layers. This neutral network was considered many years as a strong principle governing the Internet, however it begins to show some disadvantages taking into account the multimedia-intensive aspects of the FI.

The paper is organized as follows: Section II presents some of the related work existent in the field. The ALICANTE architecture and its main concepts are defined in Section III. Section IV is focused on the MANE implementation, while some measurements are presented in Section V. Conclusions, open issues, and future work are presented in Section VI.

II. RELATED WORK

The content-aware networking (CAN) and network-aware applications (NAA) approach is a new mode to design the layered architecture, with a running debate about the benefits of better interactions as opposed to the penalty of losing modularity of the architecture.

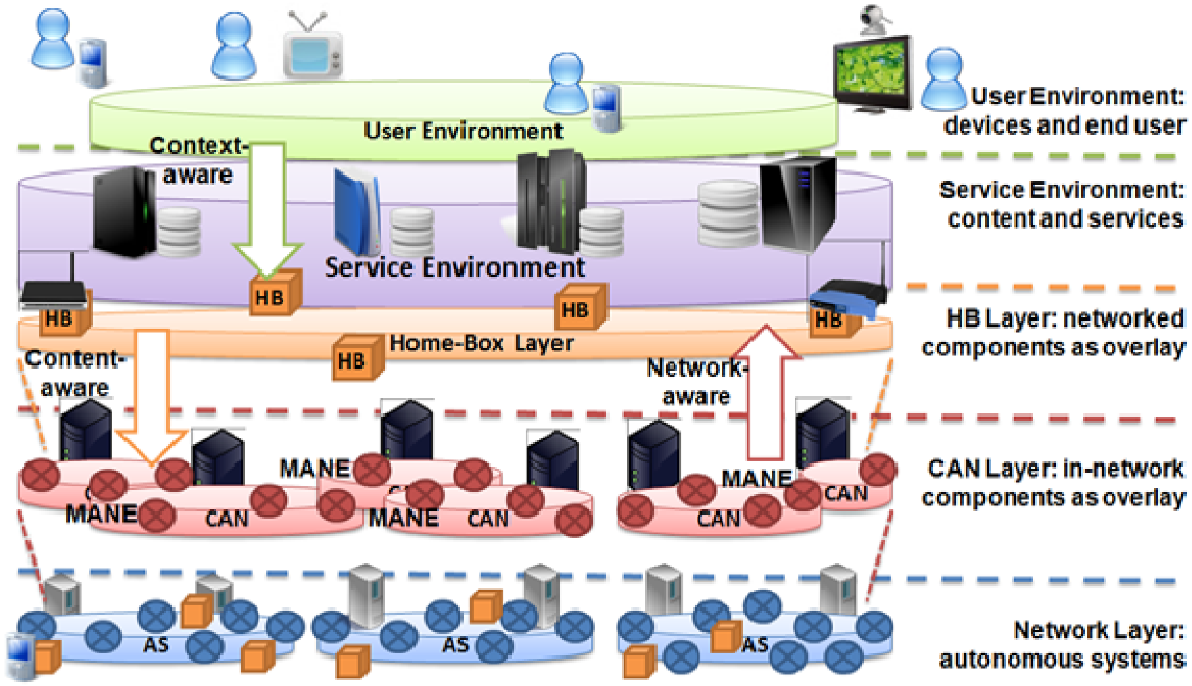

Figure 1. The ALICANTE Architecture

In such a context, both CAN and NAA are of interest both for research communities and industry, in the process of re-thinking the architecture of the FI.

The capability of content-adaptive network awareness to offer optimization for video transmission is analyzed in [5]. In [6], it is considered that CAN and NAA can offer a way for evolution of networks beyond IP. In [7], it is discussed how the CAN/NAA approach can lead to a user-centric FI and telecommunication services. The content adaptation issues in the FI as a component of CAN/NAA approach is discussed in [8]. The better QoE/QoS capabilities of the CAN/NAA architecture is analyzed in [9][10]. Further gains are obtained if context awareness is also considered [11][12].

Conversely, packet header processing time in the CAN routers raises concerns similar to the deep packet inspection techniques problems [13]. The application layer traffic optimization (ALTO) problem defined by the IETF can be solved by the cooperation between the CAN layer and the upper layer, as in [14][15].

However, no complete and open architecture currently exists, able to support multimedia distribution according to the CAN principles and scalable over sizeable networks and heterogeneous networking technologies. Therefore, an open field for research in this domain exists.

III. ALICANTE System Architecture

A. *Layers and entities*

The ALICANTE architecture, as shown in Figure 1, promotes concepts such as content-awareness to the network environment, user context-awareness to the service environment, and adapted services/content to the end-user for his/her best service experience while being either a consumer and/or producer.

Two new virtual layers are proposed on top of the traditional network layer: the CAN layer for network level packet processing and a Home-Box (HB) layer for the actual content delivery.

The CAN layer offers an enhanced support for packet payload inspection, processing and caching in network equipment. It is developed over traditional IP network/transport layer. It will improve data delivery via classifying and controlling messages in terms of content, application and individual subscribers; it improves QoS assurance via content-based routing and increases network security level via content-based monitoring and filtering. In such a way, content- and application-aware networks are created to provide high levels of performance, end-user experience, and to enable application and subscriber-specific data forwarding. The specific components in charge of creating this CAN layers are the Media-Aware Network Elements (MANE), i.e., the new CAN routers, and the CAN managers.

The Home-Box layer is an upper layer, using CAN services and taking into account network-aware information delivered upward by the CAN layer. Thanks to this layer, inter-working with the User, Service, and Network Environments, one can elaborate network and context-aware applications and deliver the necessary inputs to create content-aware networks. The Home-Box (HB) is a physical and logical entity located at end-user's premises. The adaptation, service mobility, security, and overall management of services and content are being assured at this layer through a new specific middleware

 ISBN: 978-1-61208-126-7

proposed by the project, working in conjunction with the other layers.

The interactions between the above mentioned two layers establish together a powerful cross-layer optimization loop providing end-users with the best possible service experience and optimizing the resource usage.

The upper SE layer uses information delivered by the CAN layer and enforces network-aware applications procedures necessary to perform the adaptation of the media resources to the user's preferences.

The main management and control entity in the CAN layer is the CAN Manager (CANMng). Corresponding to its roles, we distinguish the following interfaces of CAN Manager with the Virtual Home-Box layer: to advertise CANs and negotiate their usage and to help the establishing of connectivity relationships at Virtual HB layer based on, e.g., network related distance information. The CANMng has also interfaces to the lower network layer in order to negotiate CANs and request their installation.

Each AS has one CANMng, playing the following roles: to (re)define the CANs (according to the enhanced connectivity service targeted) and perform all related actions to configure, maintain and update CANs; to advertise and negotiate the CAN usage with upper layers, using Service Level Agreements/Specifications (SLA/SLS) contracts; to communicate with other CAN managers in order to establish multi-domain chains, again, using SLA/SLS contracts; to communicate with its own intra-domain network resource managers (IntraNRM). The IntraNRMs have the ultimate authority upon the network provider resources, thus conserving each domain’s independency.

B. The Content-Aware Network Router

The MANE, a content-aware network router, is an intelligent network node. It performs appropriate processing (routing, filtering, adaptation, security operations, etc.) taking into account the content type, the content properties (described by metadata or extracted by protocol field analysis) and also depending on network properties and network status. The results of the content related information analysis provide metrics, which help deciding the best strategy to adopt for the best content repurposing and publishing methods. The MANE basic set of functions are:

Content-aware intelligent routing: the MANE will decouple the higher level routing process from the lower level forwarding and perform intelligent routing, based on results extracted from packet fields analysis or content description metadata

Content-aware QoS and resource allocation: the MANEs will be able to deduct the QoS requirements of different flows based on the flows content. The CAN layer will load-wise monitor the current status of the CANs. The MANE will maintain an aggregated image of flows that they forward, and for every recognized flow type, an instance of CAN (VCAN – Virtual CAN) will be assigned depending on the level of QoS guarantees and network status. This will optimize resource allocation in the network depending on traffic types and QoS requirements The CAN level will interact with the domain network resource management in order to perform mapping onto different L2/L3 QoS-aware technologies (e.g., MPLS/Diffserv or Carrier Ethernet). Some amount of relatively infrequent dynamic re-allocation of the network resources between different CANs is possible, optimizing resource usage. The MANE has also an adaptation role, deployed at different points in the delivery chain: at the service creation, during the transport by the CAN routers, and at the Home-Box site;

Specific Security issues: increase usage of encryption technologies (such as IPSec) at network level has the direct consequence that content type can become hidden and packet inspection becomes ineffective; to mitigate this, we will exploit the possibility to include content related information in dedicated fields. Thus, the end-to-end communication remains encrypted and private, while the content-aware network concept can still function.

Another issue, beyond privacy, which is addressed by using special fields and/or metadata to describe the content, is the processing time required by deep packet inspection; eliminating the need for this procedure will significantly improve the performance of the CAN-enabled routers.

IV. MANE HIGH LEVEL ARCHITECTURE

The main interactions of the MANE are presented in Figure 2. The control plane interactions are indicated with thick shaded arrows, and the data path is indicated with thin arrows. From above, the Intradomain NRM has the role of providing the means to create FEC associations for entry in the MPLS domain. Each packet is marked with a VCAN header by the generating HB/SB, and a path is decided through CAN Manager – IntraNRM collaboration. The type of VCAN and the entry into the MPLS tunnel are then provided for each entry MANE router. The core of the mane is a classifier/router module which identifies incoming traffic based on its VCAN header, encapsulates it into the MPLS header, and sends it to the appropriate LSP.

 ISBN: 978-1-61208-126-7

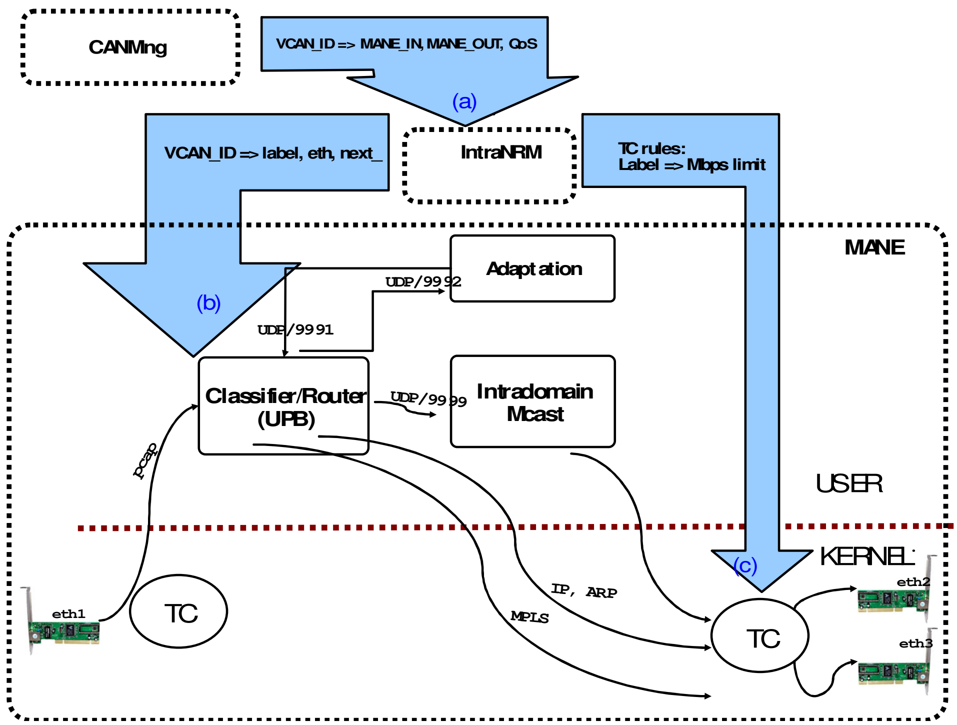

Figure 2. The ALICANTE Architecture: details on Virtual CAN Layer

The architecture is completely modular so that functionality can be developed in parallel. Modules all run in user space and are interconnected using UDP/IP so they can run on either the same, or on different machines. Another advantage is that it is easy to reconfigure the architecture by simply changing the UDP ports, or by inserting new modules.

The central module is the Classifier/Router, which harvests packets from the incoming interface and distributes them to the other modules, or encapsulates them into the MPLS paths. The classifier needs to know the association between VCAN IDs present in all incoming packets so that it can route traffic to the appropriate VCAN. VCANs are assumed to be configured in advance by the CAN Manager and provisioned through the IntraNRM. In particular, MANE needs to be instructed explicitly by the IntraNRM on the association between the VCAN_ID and an MPLS label to be used (shaded arrow marked b)). This module also decapsulates MPLS traffic that comes from the domain, before forwarding it to the appropriate HB/SB.

The core router part is not represented here, but we assume it has complete MPLS support and is implemented either with specialized hardware, or with Linux machines requiring a specially patched kernel. The IntraNRM manages al the labels and the bandwidth provisioning for each path. In fact bandwidth provisions are sent down to both MANEs and core routers to be enforced, perhaps with `tc` functionality (marked as shaded arrow c).

The Classifier/Router module is implemented in user space and uses Click modular router [16], as shown in Figure 3. For close to Linux performance it could be moved down into the kernel space in the final phases of development. It performs MPLS and IP routing both in and out the domain. It performs FEC associations for incoming IP traffic, and MPLS decapsulation for traffic outgoing to HB. It dispatches traffic to local modules (Adaptation, multicast, etc), but also accepts traffic from them so they don’t need to handle routing or encapsulation tasks. The convention in implementing the MANE is that eth0 interface is used for testbed support and therefore not part of ALICANTE. Interfaces eth1, eth2, eth3 … are used either as ingress into, or egress from the MANE. The main classification task is performed by a dedicated classifier element, called `cl_ing` (for ingress traffic), that identifies the VCAN of the incoming packet and uses the appropriate MPLS label to decide the policies for forwarding, shaping and policing. The elements grouped in the `elementclass` Card handle all the bookkeeping necessary to IP and MPLS to exchange packets on the local network (for card2 and card3, the internal details are omitted).

 ISBN: 978-1-61208-126-7

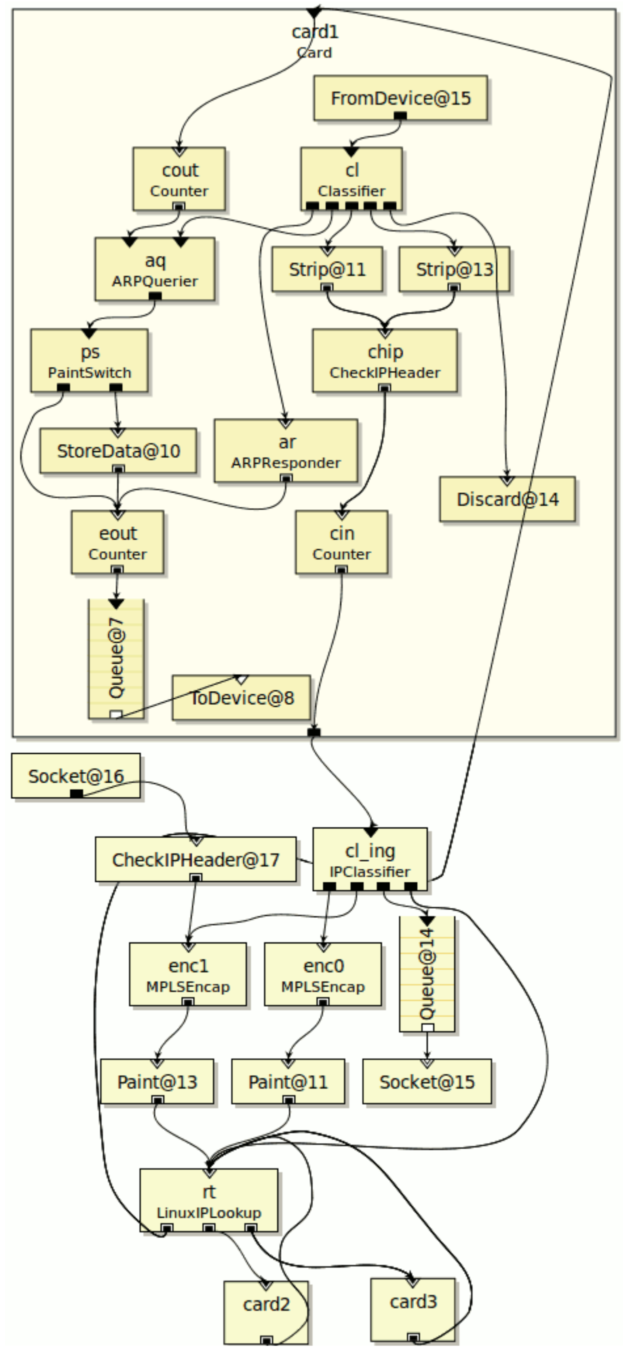

Figure 3. Userspace MANE classifier/router implementation using Click elements

V. EXPERIMENTS

Using the implementation of MANE described in the previous section, we built a topology comprising three HBs, three MANEs, and two MPLS core routers, as shown in Figure 4.

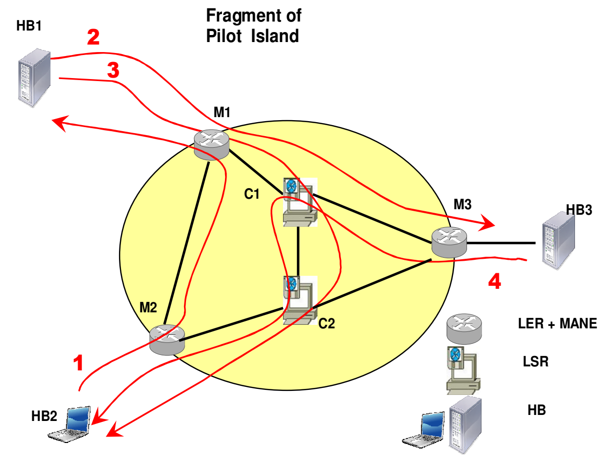

Figure 4. Pilot island using three MANE routers, two core routers and three home boxes (HB).

At this phase of the project, all the elements in the picture run in a virtualized fashion using VirtualBox on a server with quad CPU Intel Xeon X2250, 4GB memory, running Ubuntu Linux. We measured performance for three different routing configurations: using standard IP, using kernel MPLS, and using our implementation of MANE (currently in user space). The results for path 3 are summarized in Table I:

TABLE I. MEASUREMENT RESULTS

	IP	MPLS	MANE
ping 32 byte packets RTT/stddev [ms]	1.46/0.44	1.48/0.45	1.70/0.69
ping 1460 byte packets RTT/stddev [ms]	1.53/0.46	1.43/0.65	1.75/0.78
UDP Rate[Mbps]/Load	45.3/66%	45.2/66%	41.0/85%
TCP Rate[Mbps]/Load	61.0/33%	60.0/37%	37.8/62%

For the RTT tests, we used ping with large and small packets. The MANE implementation brings a minor increase in end to end transit time and a slight increase in the standard deviation of RTT. The data rates achieved with the MANE implementation are 10% less than the plain Linux data rates, but part of this difference can be accounted by the current implementation of all modules in user space. This also explains the increase in processing time in the networking elements, reported as 'Load' in the table.

VI. CONCLUSIONS AND FUTURE WORK

We presented the high level architecture of a media aware network, which aims at virtualizing network resources for the purpose of offering higher QoS to media flows. The MANE (Media Aware Network Element) is an

 ISBN: 978-1-61208-126-7

edge router that has a central role in implementing the separation between networks, by classifying incoming traffic and distributing it to appropriate MPLS paths inside each domain.

We implemented the MANE using off the shelf hardware, using Click modular router to interface the components: classification, routing, adaptation, multicast, MPLS FEC association, encapsulation and decapsulation. Our preliminary implementation on a virtual testbed shows a modest increase in processing overhead when compared with traditional IP/MPLS processing.

In the future we aim at developing the MANE in two directions: adding deep packet inspection functionality, to assist in classification of traffic not yet associated with a VCAN, and integration with high speed network processing cards, to target operation at line speed. Both these directions aim at creating a MANE that can be deployed in the field by service providers.

ACKNOWLEDGMENTS

This work was supported in part by project POSDRU/89/1.5/S/62557 and in part by the EC in the context of the ALICANTE project (FP7-ICT-248652)

REFERENCES

[1] Networked European Software and Sevices Initiative (NESSI) Strategic Research Agenda, Vol. 3. FP7-2.exec, NESSI Roadmap, Feb. 2008.

[2] European Commission, FP7 ICT Work Programme 2009-2010.

[3] Borcoci, E., Negru, D. and Timmerer, C., “A Novel Architecture for Multimedia Distribution based on Content-Aware Networking” Proc. of. CTRQ 2010, Athens, June 2010, pp. 162-168.

[4] FP7 ICT project, “MediA Ecosystem Deployment Through Ubiquitous Content-Aware Network Environments”, ALICANTE, No248652, http://www.ict-alicante.eu (last accessed: Mar. 2010).

[5] Martini, M.G., Mazzotti, M., Lamy-Bergot, C., Huusko, J., and Amon, P., “Content Adaptive Network Aware Joint Optimization of Wireless Video Transmission”, IEEE Communications Magazine, vol. 45, no. 1, Jan. 2007, pp. 84-90.

[6] Kourlas, T., “The Evolution of Networks beyond IP”, IEC Newsletter, vol. 1, Mar. 2007. Available at www.iec.org/newsletter/march07_1/broadband_1.html (last accessed: Mar. 2010).

[7] Baladrón, C., “User-Centric Future Internet and Telecommunication Services”, in: G. Tselentis, et. al. (eds.), Towards the Future Internet, IOS Press, 2009, pp. 217-226.

[8] Zahariadis, T., Lamy-Bergot, C., Schierl, T., Grüneberg, K., Celetto, L., and Timmerer, C., “Content Adaptation Issues in the Future Internet”, in: G. Tselentis, et al. (eds.), Towards the Future Internet, IOS Press, 2009, pp.283-292.

[9] Liberal, F., Fajardo, J.O., and Koumaras, H., “QoE and *-awareness in the Future Internet”, in: G. Tselentis, et al. (eds.), Towards the Future Internet, IOS Press, 2009, pp. 293-302.

[10] Huszák, Á. and Imre, S., “Content-aware Interface Selection Method for Multi-Path Video Streaming in Best-effort Networks”, Proc. of 16th International Conference on Telecommunications, Marrakech, Morocco, Jul. 2009, pp. 196-201.

[11] Baker, N., “Context-Aware Systems and Implications for Future Internet”, in: G. Tselentis et. al. (eds.), Towards the Future Internet, IOS Press, 2009, pp. 335-344.

[12] Kodeswaran, S. B. and Joshi A., “Content and Context Aware Networking Using Semantic Tagging”, Proc. of 22nd International Conference on Data Engineering Workshops (ICDEW'06), Atlanta, Georgia, USA, Apr. 2006, pp. 67–77.

[13] Rainge, E., “The Inevitable Failure of Content-Aware/DPI Network Devices and How to Mitigate the Risk”, Sept. 2008 (adapted from Worldwide Network Test and Measurement 2008.2012 Forecast and 2007 Market Shares, Available at http://www.breakingpointsystems.com/resources/white-papers/idc-white-paper/content-aware-testing.pdf (last accessed: Mar. 2010).

[14] Aggarwal, V. and Feldmann, A., “Can ISPs and P2P Users Cooperate for Improved Performance?”, ACM SIGCOMM Computer Communication Review, vol. 37, no. 3, Jul. 2007, pp. 29-40.

[15] Xie, H., Krishnamurthy, A., Silberschatz, A., and Yang, Y., “P4P: Explicit Communications for Cooperative Control Between P2P and Network Providers”, Available at http://www.dcia.info/documents/P4P_Overview.pdf (last accessed: mar. 2010)

[16] Kohler E., Morris, R.T., Chen B., Jannotti J., and Kaashoek, M. F., “The click modular router”, ACM Trans. Comput. Syst. 18, 3 (August 2000), pp. 263-297

 ISBN: 978-1-61208-126-7