Performance Evaluation of a Planar Layout of Data Vortex Optical Interconnection Network

Qimin Yang
Engineering Department, Harvey Mudd College
Claremont, California, USA
Qimin_yang@hmc.edu

***Abstract*-Optical interconnection networks provide solutions for high performance computing and communication networks which demand high throughput and low latency as well as high scalability. Data Vortex switching network has been studied previously for such purpose. In this paper, we focus on implementation issues, specifically on exploration of an alternative layout that allows the routing paths arrangement in such network to be implemented as multiple parallel planes. The original multiple cylinder three dimensional architecture is converted to planar structures. Since the new layout is designed to be functionally equivalent to the original Data Vortex, the routing performance in throughput and latency is shown to be very similar to that of the original network under same network operating conditions. The effect of injection at different angles is also investigated for optimized performance. The proposed planar architecture provides much more flexible configuration for multiple network levels, and the study validates that it provides comparable routing performance as in original design. Because of the flexibility of the new planar architecture, future work may explore further optimization in routing resources and performance.**

Keywords- Optical Interconnection Network; Data Vortex; Packet Switched; Routing; Planar Network.

I. INTRODUCTION

High performance multiprocessor supercomputers and multiple I/O communication systems require high throughput and low latency packet switched interconnection networks. As optical fiber technology matured with the optical communication industry, there are more research efforts recently dedicated to developing optically implemented interconnection networks for packet switched operation [1-2]. Considering the tremendous amount of routing decisions such networks have to make, recent researches in silicon photonics devices and platforms are important technology developments for seamless integration of the two domains [3-4]. Instead of converting existing electronic interconnection networks, the renewed interests in the area also bring in more innovative designs in network architectures that can efficiently utilize both electrical and optical domains. Photonic approaches are attractive solutions because they can not only easily achieve the bandwidth requirement, but also provide more promising potential in power consumption and scalability challenges that current electronic interconnection networks encounter [5-6]. A good network design should utilize the broad bandwidth of optical domain while avoiding extensive logic processing or optical buffering due to the lack of the mature technology.

Data Vortex network architecture is designed for localized high performance interconnection purpose, and it is scalable to a very large number of communication ports operated in packet switched mode [7-9]. With reasonable expense in routing resource redundancy, it is able to achieve very high traffic throughput while maintaining low latency and narrow latency distribution, both of which are extremely important for guarantee of packet's signal quality at the physical layer [10-11]. Data Vortex optical interconnection network relies on a three dimensional cylindrical topology to efficiently move the data flow from input to output ports [12-14]. Considering the benefit of planar structure for physical implementation, this work focuses on converting the three dimensional topology to multiple planes of routing levels to facilitate construction and integration.

The rest of paper is organized as follows: Section II explains the original Data Vortex switching topology, and difficulty with large scale system construction in cylinders. Section III presents the proposed planar layout of each cylinder level which allows for an equivalent routing topology. Section IV presents the routing performance comparison with the original layout and confirms the feasibility of the proposed system, and Section V concludes the study and discuss future works relevant to the area.

II. ORIGINAL DATA VORTEX DESIGN

The original Data Vortex network uses a cylindrical layout with multiple cylinders of routing nodes along angle and height dimensions. As an example, Figure 1 (a) shows the two outer cylinders' routing paths in a network of A=4 angles and H=4 in height. To allow for clear view, intra-cylinder paths patterns are also individually shown in Figure 1(b) for each of the three cylinders. The number of

 ISBN: 978-1-61208-226-4

cylinders required is given by $\log_2 H + 1$ due to the binary decoding nature of the routing process. The additional last cylinder is typically added, as shown in cylinder c=2 (=C) in Figure 1 (b), and it maintains the height position; but, it allows for angular resolution if each angle connects to a different sets of I/O ports. In addition, the last cylinder provides optical buffering in addition to the electrical buffers situated at I/O ports.

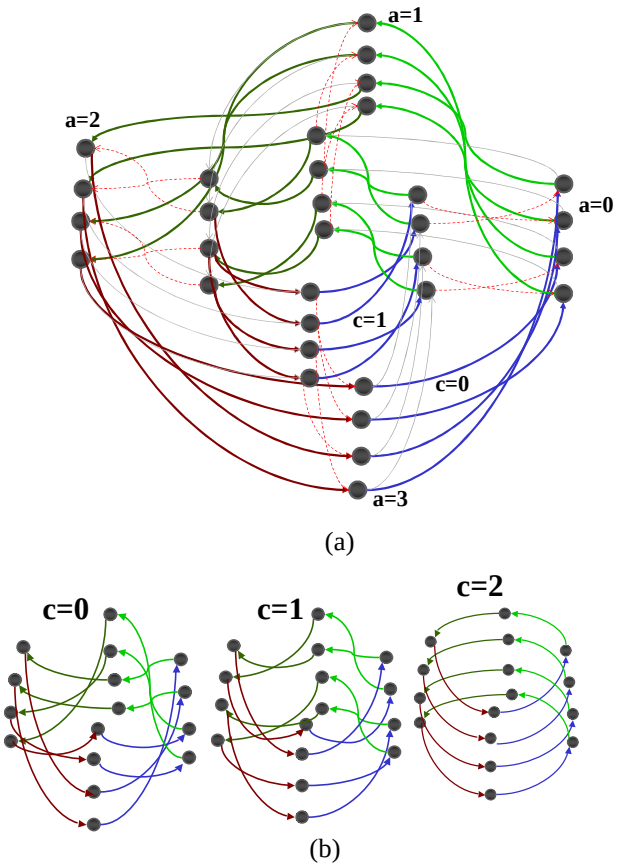

Figure 1. (a) Two outer cylinders of Data Vortex topology for A=4, H=4. (b) Intra-cylinder routing paths for each of the cylinders.

The number of active angles A_{in} are connected to I/O ports, so that the ratio A_{in}/A controls the redundancy in network operation. The choice of A_{in} needs to balance between the support of I/O ports for the given network cost and the routing performance. The smaller A_{in}/A results in better routing performance, but also means supporting of smaller I/O ports or more expensive implementation as the required number of routing nodes and optical switches is proportional to the total number of angle A. The typical choice of A is small number around 5 because an ideal operation with $A_{in}/A=1/5$ results in the optimum routing performance [7]. A much larger angle introduces much longer delay due to the latency associated with the angular resolution at the last cylinder.

The routing process starts with packets injected at the outermost cylinder, and after through each of the cylinders exit to output ports at the innermost cylinder. As shown in Figure 1, at the specific cylinder, the semi-twisted routing path patterns repeat from angle to angle which forms a cylinder by connecting the last angle to the first angle, allowing the packet to switch between two groups of height where the corresponding binary bit of the height position flip back and forth between "1" and "0". This not only allows the packets to quickly reach the correct height group, but also provide multiple open paths to reach their destination. Once the position group matches the desired group at the specific cylinder, the packet is forwarded to the inner cylinder by inter-cylinder paths (gray lines shown) that simply maintain the current height position. Such routing process continues until the packet reaches the innermost cylinder and exits the network. Another important design of the network is to guarantee single packet routing for each routing node through a traffic control mechanism. This greatly simplifies the node implementation and routing decisions, where electronic processing will not impose serious speed limitation. These traffic control signals are distributed throughout the network, and they are used between a pair of relevant nodes, sent from inner cylinder nodes to inform their outer cylinder neighbors its traffic condition. In case both have packets arriving, the inner node is always given higher priority and the outer cylinder traffic is deflected by staying at the outer cylinder instead. These control lines are shown as dash lines between the pair of nodes in Figure 1.

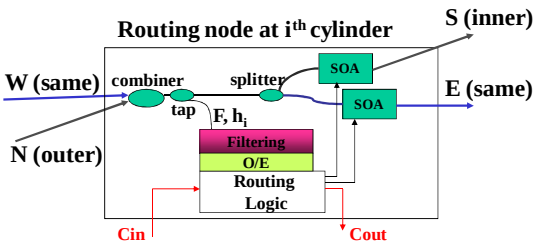

Figure 2. Routing node photonic implementation

For the physical implementation, the switches within the routing node are based on semiconductor optical amplifier (SOA). It not only allows for fast switching required for packet switching, but also provide broadband operation which allows for utilization of wavelength division multiplexing (WDM) techniques. Specifically, payload data is modulated onto WDM channels to keep short packet length while maintaining high data bandwidth, and header bits that contain the destination information are also modulated onto different wavelength channels for simple decoding within the routing nodes. In addition, any power loss in routing nodes can be easily compensated by the gain provided by semiconductor optical amplifier [8].

A detailed routing node implementation using optical components is shown in Figure 2. There are two input paths,

 ISBN: 978-1-61208-226-4

where N (North) connects to the outer cylinder node, and W (West) connects to the same cylinder node. Since the control mechanism mentioned above guarantees that at most one packet enters the node, the input paths from N and W are combined before the header bit extraction and decoding. The binary header bits of the target address are encoded using distinct wavelength channels, therefore simple passive filters and low packet rate optoelectronic (O/E) detector are used to extract such information from the optical packet. The single packet is then split to two potential output paths, one to S (South) to inner cylinder or one to E (East) to the same cylinder. The routing decision not only examine the header bit and correct height group, but also depend on an input control signal C_{in} so that ensure it only enters the inner cylinder node when there is no potential conflict for single packet condition. Similarly, the routing logic generates a new control signal C_{out} for its outer cylinder node for the same purpose. As a result of control mechanism, packet deflection can happen in Data Vortex network without packet loss. Packets that are deflected to stay on the outer cylinder can take advantage of the multiple open paths, and this causes a rather small latency penalty in comparison to other existing networks. SOA switches are used for their fast sub-nanoseconds switching speeds and internal gain for power compensation occurring at taps and splitters. More details of signaling and path connections can be found in the references [7][9].

Previous researches on Data Vortex network have been mainly focused on network's routing performance and physical layer system performances. As more applications call for such optical interconnection networks, making them more flexible and easier for construction are of great importance. So far only a 12x12 small testbed has been built mainly because it heavily relies on discrete components. On one hand, recent researches in either SOA switching fabric or novel silicon photonic devices provide more potential for integration at the device and routing node level. On the other hand, network architectures may not have the same upgrades for easier implementation. For example, a small scale Data Vortex can be easily implemented using fiber waveguides and individual node modules in three dimensions. But, such arrangement could be very complex and cumbersome for a much larger network. The difficulties also include keeping every level aligned and synchronized necessary for the routing operation. For all paths physically the same length at all levels, inner levels must somehow wind up paths to occupy a smaller physical space than its outer level in cylindrical arrangement. In addition, due to close coupling of the electrical layer for routing logic and traffic control with the optical layer, a fully three dimensional fabrication is not compatible to integration solutions. One solution is if each of the routing levels or cylinders can be integrated as a subsystem on a plane as the electrical circuits naturally arrange on planes, it becomes a much more manageable overall system which simply interconnects the planes of subsystems. The complexity would not grow drastically as the size of the system. Because connections between the levels are parallel links, either fiber based on other type of waveguides in more integrated form can be used.

III. PLANAR LAYOUT DESIGN

To eliminate the incompatibility of integration of electrical layer in the cylindrical arrangement, and make Data Vortex easier to construct in large scale networks, this study explores an alternative layout of the Data Vortex architecture to allow for planar construction of the multiple routing levels.

In the logical level, we want to maintain the same principle of minimizing deflection probability by arranging the same semi-twisted routing patterns, while allowing for parallel planes for easy layered integration. To achieve this, instead of connecting the last angle of routing nodes to the first angle in the original Data Vortex, we added paths along the same angle at the first and the last angle (green paths as shown in Figure 3 and Figure 4), while changing half of the routing paths in the opposite traveling direction (red paths vs. blue paths). As a result, traveling on the same plane now forms a looping pattern similar to the ones by staying on the same cylinder in the original Data Vortex network. As examples, Figure 3 and Figure 4 show a network with H=4 at the first two routing levels (in comparison to the two outermost cylinders in the original network layout in Figure 1) for A=4 and A=5, respectively. As shown, the direction of the same angle green paths at the first and last angle depends on whether the angle A is even or odd, and the connection pattern follows the same pattern between the nodes as in between angles. The same cylinder paths' direction also depends on whether A is even or odd accordingly as shown.

The new layout now allows for parallel planes of different routing levels that correspond to the original cylinders. Between different planes, only parallel routing paths are needed as that in the original Data Vortex networks. The control signal paths are not shown for a clear view, but they should be set up similarly between the pair of nodes from inner cylinder to the outer cylinder whenever two nodes try to send packets to the same node. These control paths can be integrated with the optical routing path on another plane that would be perpendicular to the planes of routing levels. The control signals apply to the edge angle nodes in a similar fashion even though the output node could be located on the same angle. The detailed organization of nodal circuits on the plane is beyond this study. Either all nodes of the same plane are fabricated on a same platform/board which requires planar optical waveguide technology, or nodes of the same angle can be fabricated on a same board if angles are interconnected by more flexible fiber waveguides. Three dimensional arrangements are still open, but without the difficulty of occupying a different size space as in cylindrical networks.

 ISBN: 978-1-61208-226-4

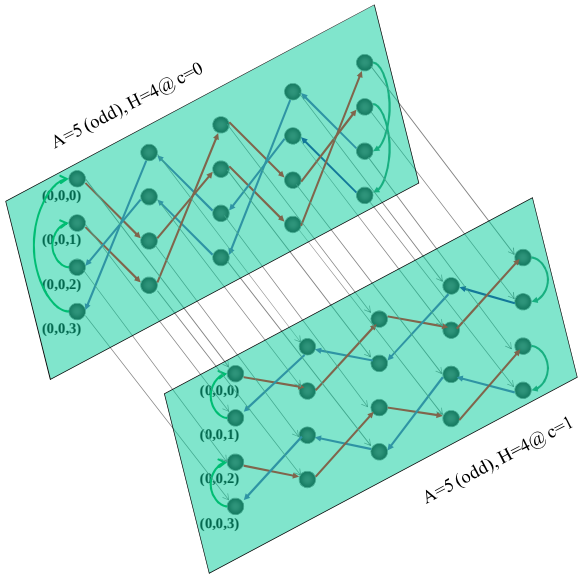

Figure 3. Routing paths in planar layout for A=5 (odd) and H=4 at the first two levels

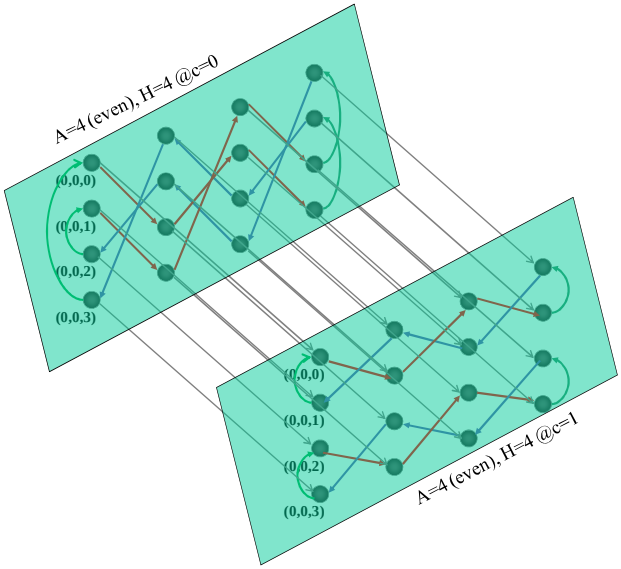

Figure 4. Routing paths in planar layout for A=4 (even) and H=4 at the first two levels

Logically, the new layout forms a very similar connection as that of the original Data Vortex network, and the rest designs such as parallel inter-level forwarding paths, control mechanism and additional last cylinder will all maintain the same. Therefore, we should expect very similar routing performance when comparing the new layout to the original Data Vortex cylindrical layout. On the other hand, the new same angle paths may affect the routing performance because nodes at different angles may carry slightly different traffic load and may result in different traffic distribution within the network. Packet injection at different angle should also be investigated with further details to make sure the planar layout can achieve the same routing performance as desired.

IV. PERFORMANCE EVALUATION

A custom written C/C++ event simulator is used to evaluate the new layout Data Vortex architecture. Same network size is chosen for two layouts while different traffic load and network redundancy are included in the study to ensure the performance comparison at various operation conditions. We choose the network angle to be A=5 as in earlier studies. A less redundant condition with increased A_{in}>1 for the same given A is also studied [7]. The delay performance examines the average number of hops packets experience in the network over a long period of simulation time after a steady state is reached. The throughput performance is presented by the successful injection rate at the input ports over the same simulation period. Because it is a non-blocking network, the successful injection rate reflects the overall data capacity that the network can handle. A more congested or saturated network essentially deflects more traffic at outer levels and creates traffic backpressure up to the injection ports. We always choose the same angle parameter for comparison to the original network layout, so angular resolution is not included for this study.

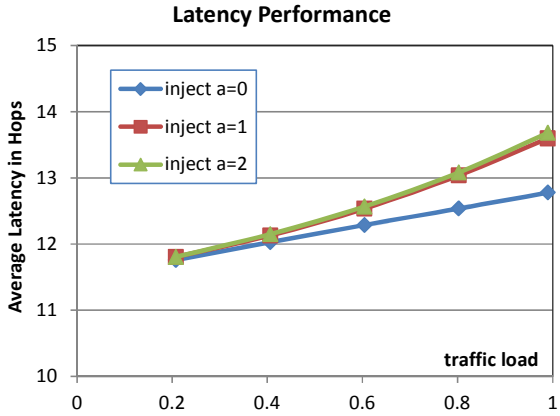

Figure 5. Latency Performance Comparison for injection at different angles for A=5, A_{in}=1, H=128

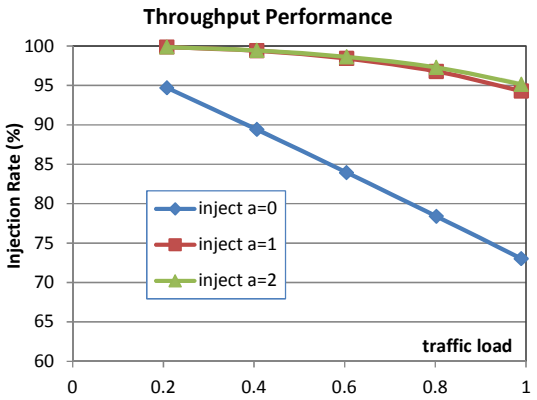

Figure 6. Throughput Performance Comparison for injection at different angles for A=5, Ain=1, H=128

First, we study the angular dependence of injection port in the new planar layout. Due to the symmetry of the layout, in a network of A=5, we only need to compare routing performance with injection occur at a=0 (end angle), a=1 and a=2 (middle angle) respectively. Figure 5 and Figure 6 present the latency and throughput performance for A=5 and H=128 with one injection angle A_{in}=1 at the specified angle a. The ratio of A_{in}/A is chosen for the optimized throughput performance as suggested in previous study. As shown, we should avoid inject at the end angles where we loop the packets back in the opposite direction on the same angle because the injection rate and throughput performance is shown to be significantly lower. The performance difference between the two middle angles however is shown to be negligible. More case studies with multiple injection angles also show similar results, which indicates that edge angles should be generally avoided if possible. The middle angle or angles close to the middle should be preferred to achieve the best overall performance in throughput and latency in the new layout. In case the network operates in much less redundant condition such as injecting at all A angles, then the edge angles have to be used as well.

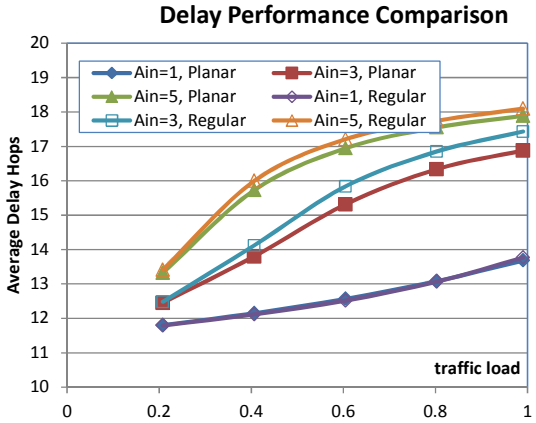

Figure 7. Delay performance comparison for two layouts for different A_{in} with A=5, H=128

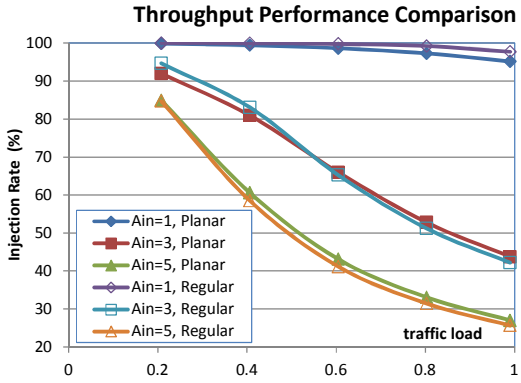

Figure 8. Throughput comparison for two layouts for different A_{in} with A=5, H=128

Next, we compare the planar layout performance with the original Data Vortex network as shown in Figure 7 and Figure 8, respectively. Here, all planar layouts use middle angles for injection angles to avoid unnecessary performance degradation, while A_{in}=5 case all angles including edge angles will be used for injection. The planar layout results are marked by solid shapes while regular layout results use hollow shapes for all three redundant operations. As shown, the routing performances in all cases are very close between the two layouts. We also notice that for Ain=1, the regular network achieves a little better throughput, while the planar layout achieves slight benefit over regular network in less redundant network conditions, such as Ain=3 and Ain=5. However, overall, there is very small difference between the two layouts as long as the injections are carried out through the middle angles.

Finally, to show the performance evaluation is valid for all network sizes, we also examine the comparison for different network heights. It has been confirmed that only small discrepancies have been observed between the two layouts and very similar trends are seen, as shown in Figure 7 and Figure 8 for different network sizes.

V. CONCLUSION

To summarize, an alternative planar layout of the Data Vortex architecture was proposed to facilitate the three dimensional integration of the network nodes. This is accomplished by dividing routing paths along the cylinder to two different directions, while adding same angle paths in the first and last angle to form similar routing loops. As a result, all nodes on the same routing level can be implemented on a plane, and overall architecture appears to be multiple planes interconnected. The optical paths and control links between routing levels along the same angle can further be put on a plane perpendicular to the routing level planes, and easily form a three dimensional structure that facilitates the modular design and synchronization of

 ISBN: 978-1-61208-226-4

the overall network. The routing performance is specifically evaluated and in comparison to the original network performance under various network conditions and network sizes. The detailed analysis has shown that the new planar layout achieves very similar routing performance as that of original Data Vortex network, and, in general, edge angle should be avoided for packet injection for optimum routing performance.

REFERENCES

[1] A. Wonfor, H.Wang, R.V.Penty, and I.H. White, "Large Port Count High Speed Optical Switch Fabric for Use Within Datacenters", *Journal of Optical Communications and Networks*, Vol. 3, No. 8, pp. A32-39, 2011.

[2] Shinji Nishimura, Kazunori Shinoda, Yong Lee, Fumio Yuki, Takashi Takemoto, Hiroki Yamashita, Shinji Tsuji, Masaaki Nido, Masahiko Namiwaka, Taro Kaneko, Kazuhiko Kurata, Shigeyuki Yanagimachi, and Naoya Ikeda, "Optical Interconnection for High-speed Routers", *Optical Fiber Conference*, OThH2, 2011.

[3] T.-Y. Liow, K. W. Ang, Q. Fang, M. B. Yu, F. F. Ren, S. Y. Zhu, J. Zhang, J. W. Ng, J. F. Song, Y. Z. Xiong, G. Q. Lo, and Dim-Lee Kwong, "Silicon Photonics Technologies for Monolithic Electronic-Photonic Integrated Circuit Applications", *Optical Fiber Conference*, OThV1, 2011.

[4] Noam Ophir1, Kishore Padmaraju1, Aleksandr Biberman1, Long Chen2, Kyle Preston2, Michal Lipson2, and Keren Bergman, "First Demonstration of Error-Free Operation of a Full Silicon On-Chip Photonic Link", *Optical Fiber Conference*, OWZ3, 2011.

[5] R.S.Tucker, "Green optical communications-part II: energy limitations in networks", *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. 17, No. 2, pp 261-274, 2011.

[6] Odile Liboiron-Ladouceur, Pier Giorgio Raponi, Nicola Andriolli, Isabella Cerutti, Mohammed Shafiqul Hai, and Piero Castoldi, "Scalable Space-Time Multi-plane Optical Interconnection Network Using Energy-Efficient Enabling Technologies", *Journal of Optical Communication and Networks*, Vol. 3, No. 8, pp. A1-A11, 2011.

[7] Q. Yang, K. Bergman, G. D. Hughes, and F. G. Johnson, "WDM packet routing for high-capacity data networks," *J. Lightw. Technol.*, vol. 19, pp. 1420–1426, Oct. 2001.

[8] C. Hawkins, B. A. Small, D. S. Wills, and K. Bergman, "The Data Vortex, an all optical path multicomputer interconnection network," *IEEE Transactions of Parallel and Distributed Systems*, vol. 18, No. 3, pp. 409-420, Mar 2007.

[9] Cory Hawkins, D. Scott Wills, Odile Liboriron-Ladouceur, and Keren Bergman, " Hiearchical clustering of the data vortex optical interconnection network", *Journal of Optical Networking*, Vol.6, No. 9, pp. 1179-1190, Sep. 2007.

[10] O.Liboiron-Ladouceur, B.A.Small and K.Bergman, "Physical Layer Scalability of WDM Optical Packet Interconnection Networks", *J. Lightwave Technology*, vol. 24, pp. 262-270, 2006.

[11] A. Shacham, B.A. Small, O. Liboiron-Ladouceur and K. Bergman, "A Fully Implemented 12x12 Data Vortex Optical Packet Switching Interconnection Network," *Journal of Lightwave Technology*, vol. 23, No. 10, pp. 3066-3075, 2005.

[12] Neha Sharma, D. Chadha, Vinod Chandra, "The augumented data vortex switch fabric: an all-optical packet switched interconnection network with enhanced fault tolerance", *Journal of Optical Switching and Networking*, Vol. 4, pp. 92-105, 2007.

[13] A. Shacham and K. Bergman, "Optimizing the performance of a data vortex interconnection network," *Journal. Optical Networking*, vol. 6, No. 4, pp. 369-374, April 2007.

[14] A. Shacham and K. Bergman, "On contention resolution in the data vortex optical interconnection network", *Journal of Optical Networking*, vol. 6, No. 6, pp. 777–788, June 2007.

[15] http://lightwave.ee.columbia.edu/downloads/acs_workshop/Columbia_GaTech_ACS_April16_2010_hw2.pdf (Last retrieved September 2012)

 ISBN: 978-1-61208-226-4