

Article

Free Space Intra-Datacenter Interconnects Based on 2D Optical Beam Steering Enabled by Photonic Integrated Circuits

Charidimos Chaintoutis ^{1,*} , Behnam Shariati ² , Adonis Bogris ³, Paul V. Dijk ⁴,
Chris G. H. Roeloffzen ⁴, Jerome Bourderionnet ⁵, Ioannis Tomkos ⁶  and Dimitris Syvridis ¹

¹ Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Panepistimiopolis, Ilissia, 15784 Athens, Greece; dsyvridi@di.uoa.gr

² Department of Computer Architecture, Polytechnic University of Catalonia, C/Jordi Girona 1-3, 08034 Barcelona, Spain; behnam.sh@ieee.org

³ Department of Informatics and Computer Engineering, University of West Attica, 12243 Aghiou Spiridonos, Greece; abogris@teiath.gr

⁴ Lionix BV, P.O. Box 456, 7500 AL Enschede, The Netherlands; p.w.l.vandijk@lionix-int.com (P.V.D.); c.g.h.roeloffzen@lionix-int.com (C.G.H.R.)

⁵ Thales Research & Technology, 1 Avenue Augustin Fresnel, 91767 Palaiseau CEDEX, France; jerome.bourderionnet@thalesgroup.com

⁶ Athens Information Technology (AIT), Monumental Plaza, 44 Kifisias Ave., Marousi, 15125 Athens, Greece; itom@ait.gr

* Correspondence: cchaintoutis@di.uoa.gr; Tel.: +30-210-7275347

Received: 27 June 2018; Accepted: 26 July 2018; Published: 1 August 2018



Abstract: Data centers are continuously growing in scale and can contain more than one million servers spreading across thousands of racks; requiring a large-scale switching network to provide broadband and reconfigurable interconnections of low latency. Traditional data center network architectures, through the use of electrical packet switches in a multi-tier topology, has fundamental weaknesses such as oversubscription and cabling complexity. Wireless intra-data center interconnection solutions have been proposed to deal with the cabling problem and can simultaneously address the over-provisioning problem by offering efficient topology re-configurability. In this work we introduce a novel free space optical interconnect solution for intra-data center networks that utilizes 2D optical beam steering for the transmitter, and high bandwidth wide-area photodiode arrays for the receiver. This new breed of free space optical interconnects can be developed on a photonic integrated circuit; offering ns switching at sub- μ W consumption. The proposed interconnects together with a networking architecture that is suitable for utilizing those devices could support next generation intra-data center networks, fulfilling the requirements of seamless operation, high connectivity, and agility in terms of the reconfiguration time.

Keywords: optical communications; optical beam steering; data center interconnects

1. Introduction

Warehouse scale data centers (DCs) host several hundreds of thousands of servers and other equipment, spreading across thousands of racks that are interconnected by an underlying intra-data center network (intra-DCN) architecture. Such architectures should be highly scalable, upgradable, energy efficient, and extremely cost-effective [1]. Nowadays, most intra-DCNs are based on scale-up hierarchical designs (e.g., fat-tree [2] and butterfly [3]) which suffer from oversubscription, lack of fault tolerance, expensive layer-2/3 switches/routers, weak re-configurability, and poor scalability when

traffic increases. New architectures (e.g., based on modular designs) have been recently proposed addressing these issues [4]. However, the proposed solutions impose other challenges that demand low-cost/low-power dynamic re-configurability and seamless scalability [5].

Optical fibers and fiber-pigtailed transceivers are the key elements in constructing traditional DCNs. Millions of meters of fiber are required, forming an extremely complex cabling system that is difficult to manage [6], with the choice of cabling architecture greatly influencing the throughput, scalability, and energy efficiency management of a DC [6,7]. Although the cabling problem is still underestimated, when the scaling and cost of large infrastructures is under consideration, recent studies have started paying attention to its side effects. For instance, it was shown that the mass of long cables between the switches causes great difficulties in system maintenance [5]. When the network connection changes or line failure occurs, the system upgrading becomes extremely complicated. IBM also reported that the dense cabling will affect the heat dissipation due to the fact that airflow is obstructed [8], while HP reported that mapping a logical topology with servers, switches, and links onto a physical space with racks and cable trays such that the cable costs are minimized is a non-deterministic polynomial-time hard (NP-hard) problem [9].

At the same time, the traditional way of designing DCNs, through the use of electrical packet switches in a multitier topology, has a fundamental weakness, namely that it has to be decided in advance whether to provide overprovisioned or oversubscribed interconnects between top of rack (ToR) switches. The former choice is quite expensive, whilst the latter leads to poor performance when congestion emerges. Recent work suggests the augmentation of oversubscribed networks with reconfigurable schemes on top of the electrically-switched network, using 60 GHz RF wireless [10], optical switching [11] and Free Space Optics (FSO) [12].

Wireless intra-DCN interconnection solutions (based either on radio or optical frequencies), have been proposed to deal with the cabling problem [13], and can simultaneously address the over-provisioning problem by offering topology re-configurability. One of the most promising wireless approaches is free space optics, which can support the dynamic change of capacity between pairs of ToRs in the short-term and to enable the realization of a DCN with all its inter-rack links being flexible and wireless in the mid-term. In addition to the above-mentioned advantages, FSO can benefit from vast license-free spectrum ($\cong 400$ THz) with extremely small footprint, while guaranteeing high data rates compared to RF links [12,14]. Additionally, by proper design of DCs and circumventing intermediate routers and removing cables, power consumption can be substantially reduced [12].

In [15], a novel FSO solution for intra-DCNs, utilizing 2D optical beam steering (OBS) for the transmitter, and high bandwidth wide-area photodiode arrays for the receiver, was proposed and analyzed. In the present work, we extend the contributions presented in [15], providing more results and implementation details. More specifically we provide a possible implementations for a lens optical system that could increase the optical power reaching the receiver. We include more results regarding the system's power budget and include more information about potential networking schemes that could incorporate the proposed transceivers, and explore intra-cluster connectivity scenarios that could provide increased inter-rack coverage, and finally we explore a scenario for direct, single-hop blade-to-blade communication among different racks that is a significant step towards disaggregation.

The proposed interconnection solution can support radical change in intra-DCN topologies, enabling random graph topologies. Such an architecture can be supported by a new-breed of FSO nodes utilizing novel transceivers, which can be developed on a photonic integrated circuit (PIC), offering a significant reduction in the cost, footprint, and switching speed. Thus, the proposed FSO interconnects could form the basis of a reconfigurable intra-DCN that could fulfil the properties of seamless operation, high connectivity (direct communication from a rack to many others) and agility in terms of the reconfiguration time.

2. AI Properties of the Proposed Optical Beam Steering

Current systems that incorporate beam steering include the use of micro-electromechanical mirrors and liquid-crystal-based spatial light modulators, with the tuning speed limited to the millisecond range [16]. Conceptually, the most straightforward approach to providing beam steering in two dimensions is to use a 2D array of coherent emitters so that each can be controlled in terms of the output phase. Initial work in this field focused on 1D arrays fabricated in InP and silicon photonic technology [17]. Even though thermo-optic phase tuning was used, its bandwidth still outperformed liquid-crystal technology by two orders of magnitude. A more elegant approach, in our opinion, is to realize the full 2D emitter array on a single PIC. High contrast silicon photonics is well suited for this, making use of vertical grating couplers. A 64×64 array consisting of a total of 4096 emitters was developed where no active beam shaping and steering was possible [18]. A second category of 2D OBS machines is the one utilizing tunable lasers. This configuration severely decreases the complexity of the control electronics required for 2D steering. Using the hybrid integration of silicon and InP technology, the laser source can be integrated on the same PIC. Full 2D OBS was achieved with a 32-channel hybrid silicon PIC that includes a tunable laser among other components [19].

Here, we propose a new class of transmitters relying on PICs offering a fast and large field-of-view OBS for intra-DCN. The transmitter PIC can be based on the hybrid integration of InP elements (i.e., lasers, modulators, amplifiers) on a TriPleX™ platform [17], as can be seen in Figure 1. TriPleX waveguides form a new class of integrated-optical planar lightwave circuits using low-cost, CMOS-compatible fabrication equipment.




Figure 1. Conceptual illustration of the envisioned transmitter.

2.1. Optical Beam Steering Technologies at the Transmitter Side

The OBS is accomplished using two mechanisms. The first mechanism is wavelength tuning. According to calculations, a wide angle of 20° could be achieved using 80 nm wavelength tuning, incorporating one or more tunable laser sources in the vertical direction. In the horizontal direction, the technology could provide similar range with the use of a high number of pixels and phase shifters relying on piezoelectric transducers (PZT) with sub- μ W power consumption and ns response time [20], thus complying with the low energy consumption and low latency requirements of DCNs. The available technology promises the fabrication of up to 1000 pixels with an inter-pixel distance of 20 μ m. By using stress-optic tuning with PZT the power consumption will be reduced by a factor of 1000 compared to thermo-optic technologies.

Each pixel is a grating coupler that emits elliptical light beams, with a therefore different divergence in the “vertical” direction (wavelength steering direction) and in the “horizontal” direction (pixel steering direction). Even if the divergence is small, it affects the propagation properties of the beams, thus beam collimation is a prerequisite for the implementation of a practical link with reach exceeding 10 m. A beam size (waist) in the order of a few mm is required to preserve the beam collimated for more than 20 m. In the vertical direction, the divergence is defined by the length and strength of the grating. A length of 2 mm guarantees sub-degree divergence. To control the divergence in the horizontal direction, a lensing system that relies on the use of a cylindrical lens in front of the chip could be incorporated, as is depicted in Figure 2. In the reception, even if collimated, the beam will have a diameter of a few mm. Hence, a 2D photodiode array should be utilized in order to

collect as much as possible of the optical power of the emitted light, as described in detail in the following section.

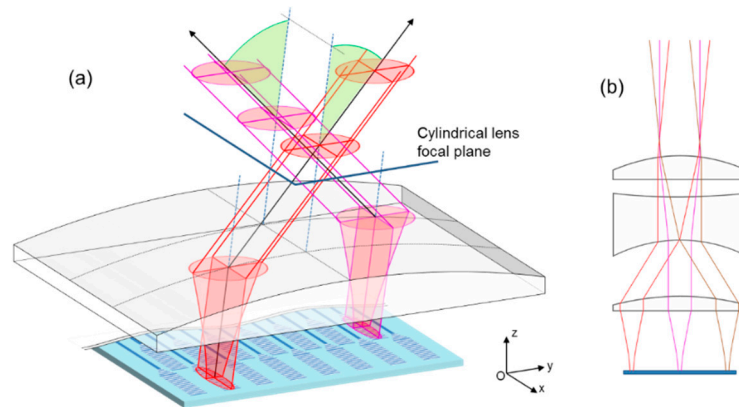


Figure 2. (a) Schematic of the use of a cylindrical lens to compensate exit mode ellipticity and obtain sub-degree divergence in both the horizontal (x) and vertical (z) directions; (b) Example of a three-element optical system implementation.

The radius of the free space beam will vary according to the well-known expression [21]:

$$\omega(z) = \omega_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2} \quad (1)$$

where $z_R = \frac{\pi\omega_0^2}{\lambda}$ is the Rayleigh length, λ is the wavelength, and ω_0 is the beam waist. A typical value of the waist of the beam emitted by the grating coupler is 10 μm . Thus, if a focal length (distance between the grating couplers and the cylindrical lens) of 80 mm is considered, the waist of the collimated beam will be close to 3 mm for a transmitted wavelength of 1.3 μm . Based on (1) and considering that the radius of the beam emitted by the transmitter will be in the order of a few mm, we estimate the radius of the beam as a function of distance for three beam waist values (Figure 3). It seems that a beam of 3 mm waist could be preserved relatively collimated for more than 30 m. In order to support longer distances, a larger initial beam waist is required.




Figure 3. Beam radius as a function of distance for three initial beam waists.

2.2. Optical Beam Steering Technologies on the Receiver Side

On the receiver side, the solution of a 2D photo diode (PD) array that will combine the high rate performance and the large field of view is proposed. For the bandwidth of 20–25 GHz that is targeted for modern intra-DCN interconnects targeting 100 Gb/s speeds, each single photodiode aperture is as small as 20–25 μm , thus making efficient optical coupling to the largely expanded beam substantially

challenging. The PD array at the receiver site will have to contain a large number of single photodiode elements in order to detect a reasonable amount of power. Based on the specifications of commercially available PD arrays [22], a linear 25 GHz photodiode array featuring integrated lenses, so as to enlarge the aperture per PD to 50 μm , could be used in order to substantially extend the detection area. The lens pitch center-to-center can be as small as 250 μm . With 16 apertures lined up, this would result in a linear PD array chip width of 4 mm (see Figure 4) corresponding more or less to the beam diameter at the receiver site for moderate distances (<40 m) (see Figure 3).




Figure 4. Beam footprint on the two-dimensional PD array.

With the use of a cylindrical lens in front of the 2D PD array, the amount of power concentrated on the receiver could be substantially increased. If no lens is considered in front of the 2D PD array, then the amount of the detected optical power is the ratio of the detection area over the beam cross-section. Almost 86% of the emitted optical power is contained within a circle of radius $r = \omega(z)$ if a Gaussian beam is considered. Each photodiode element is hence considered to be a circular detection area with a diameter of around $d_{PD} = 50 \mu\text{m}$. The percentage of power reaching its element is equal to (d_{PD}^2 / d_{beam}^2) . Figure 5 shows the optical losses due to the partial coverage of the PD array by the transmitted beam.




Figure 5. Losses due to the beam divergence as a function of the transmission distance (beam waist: 2 mm).

It can be seen that the 2D PD arrays could detect the transmitted signal with losses, due to the high beam cross-section compared to the PD array area, varying from 15 to 25 dB as a function of the number of PDs in the array. The losses also vary with distance due to the broadening of the beam waist, as depicted in Figure 3. The losses increase by 20 dB after 100 m of propagation and could be less than 25 dB, provided that the distances are shorter than 20 m, the initial beam waist is above 2 mm, and 2D arrays of at least 64 PDs are considered. The amount of 25 dB losses is marginally acceptable as it will be explained in the next paragraph.

2.3. Power Budget

The transmitter could emit 5 dBm out of the chip in order to support longer reach transmission. Assuming an extra lens in front of the 2D array that will roughly improve the detectivity by 5–10 dB (not the same effect for all steering angles), the detected power could scale with distance according to the formula $P_{REC} = P_{TRANS} + G_{CL} - L_{BD}$, where P_{REC} is the overall received power, P_{TRANS} is the transmitted power, G_{CL} is the gain offered by the use of the lens in front of the 2D array, (~ 10 dB), L_{BD} is the loss due to partial coverage and beam divergence already presented in Figure 5. The results are depicted in Figure 6 for an initial beam waist equal to 3 mm. We can see that for moderate distances (< 50 m), the system could support high bit rate operation (50 Gb/s), and for all distances 10 Gb/s (even 25 Gb/s PAM-4) is feasible especially if avalanche PDs with a sensitivity lower than -20 dBm are considered in the array. Additionally, the distance and the baud rate could be increased if a powerful Forward Error Correction (FEC) is employed, however with a proportional increase of the cost and the complexity. Another way to increase both figures is the utilization of shorter wavelengths (800 nm), which have a higher Rayleigh length.

The transmission losses due to atmospheric absorption are not expected to be more than 1 dB, as the attenuation coefficient for infrared wavelengths in clear weather conditions is less than 0.5 dB/km. An open issue is the effect of DCN conditions on the transmission properties of the optical wireless links, however, this would require detailed investigation and is beyond the scope of this paper. The proposed FSO node has enhanced accuracy and reliability compared to [12,14], as it does not include moving objects/mirrors to re-direct the beam.

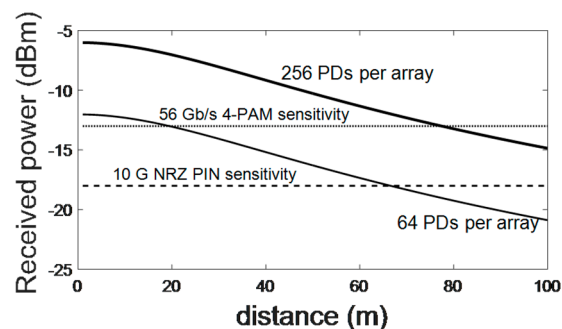


Figure 6. Transmission performance of the Optical Beam Steering systems (received power as a function of distance). The transmitted power out of chip is 5 dBm at 1300 nm, and the receiver contains 64 PDs or 256 PDs. In the plot, the sensitivity of the state of the art 4-PAM and 10 G systems using PIN PDs is also depicted (beam waist = 3 mm).

3. Inter-Data Center Networking Scheme and Results

In this section we demonstrate a DCN architecture that could efficiently incorporate the aforementioned devices. With proper DC rack orientation, the transmitter and receiver can have direct line of sight (LoS). We propose a cylindrically shaped cluster with the racks placed on the cluster's border, as shown in Figure 7a,b, thus increasing the number of direct inter-rack FSO links. Figure 7b shows a cluster hosting 48 racks (universal standard racks with 50 cm width are considered). The coverage of OBS elements is determined by the steering angle (θ) and the diameter of the cylinder (R). Increasing the number of elements or the tuning range would increase the steering angle. ToR switches will be tier-2 for this architecture. On the other hand, tier-1 switches equipped with FSO transceivers, placed in the center of each cluster (specified with black dot), are utilized for inter-cluster connectivity, as shown in Figure 8b.




Figure 7. (a) Data Center cluster hosting 48 racks and OBS element with steering angle θ ; (b) The coverage of the FSO node for inter-cluster connection.




Figure 8. Proposed DC clusters and server orientations. (a) From top; (b) Inter-cluster and inter-rack FSO links.

Tier-1 FSO nodes will be placed at a greater height compared to the ToR tier-2 FSO nodes, in order to have direct LoS. Due to 2D OBS capability, tier-2 FSO nodes can communicate using certain wavelengths, while different wavelengths can be exploited for inter-cluster connectivity. The proposed solution will heavily reduce the number of required switches/routers compared to fiber-based solutions [7]. We present the results on the coverage potential of the proposed FSO nodes considering different practically possible values for the steering angle (10° – 20°) and the distance the beam can travel (25 m–100 m). Regarding the cluster size, we assumed two alternatives, hosting 36 or 48 racks. Assuming realistic size racks, we obtain the number of direct inter-cluster and intra-cluster connections that the FSO nodes can establish. The numbers obtained rely on the fact that there is direct LoS between each pair of tier-1 (for inter-cluster) and each pair of tier-2 (inter-rack) FSO nodes. In Table 1, we present the coverage for inter-cluster FSO nodes, while the inter-rack coverage is summarized in Table 2. The presented numbers can be increased if spatial multiplexing is employed by using more than one transceiver, positioned in a proper configuration so as to enlarge the field of view of each node. For instance, three transmitters positioned in a row in the inter-cluster communication scenario depicted in Figure 7b, equivalently to what mobile communications traditionally do using multiple directional antennas to achieve 360-degree coverage, would increase the steering angle to 30° and would allow all-to-all single-hop communication at the tier-1 level.

Table 1. The number of possible direct inter-cluster links one FSO node can establish.

R (Meter)/ Θ (Degree)	Clusters of 36 Racks				Clusters of 48 Racks			
	25	50	75	100	25	50	75	100
10	4	13	29	52	2	7	14	25
15	5	20	44	77	3	10	21	37
20	7	26	58	103	4	13	28	49

Table 2. The number of direct inter-rack links one FSO node can establish.

	10°	15°	20°
36 racks	3	4	6
48 racks	4	5	7

Although the inter-cluster coverage is remarkable, the intra-cluster connectivity suffers from the limited steering angle, which can be alleviated exploiting proper mirrors [12] as depicted in Figure 9a. Moreover, the 2D steering capability can be exploited so as to allow for direct, single-hop blade-to-blade communication among different racks. The transceiver of each blade selects the rack using horizontal steering and can select the blade belonging to this rack through vertical steering (see Figure 9b), enabling flat disaggregated DCN architectures [23]. Due to the high level of direct inter- and intra-cluster connectivity and ultra-fast optical switching capability envisioned in the solution, the DCN topology can be intelligently re-configured, adapting to rapid traffic changes. Note that the specific devices are able to support radical change in the intra-DCN architecture, enabling random graph topologies which have proved to be more scalable and flexible [24].

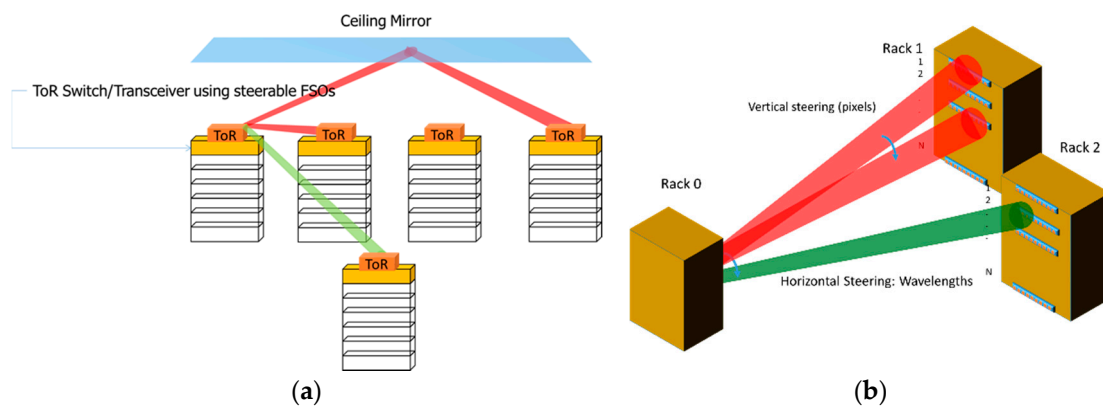


Figure 9. (a) Reconfigurable inter-rack communication of larger coverage with the use of mirrors (b) Blade-to-blade inter-rack communication that could enable intra-DCN disaggregation.

4. Conclusions

We presented a new class of FSO-integrated transceiver with optical beam steering that provides switching at sub- μ W consumption in order to support next generation intra-DCN architectures. A networking architecture that is suitable for the utilization of those devices was discussed and dimensioning results were presented. This architecture greatly reduces the cabling complexity and related maintenance issues, which are of concern for the current DCs. Potential alternative schemes that could further improve the connectivity and that could also enable flat architectures suitable for intra-DCN disaggregation were also discussed. Due to the higher number of direct links compared to the current architectures, the number of intermediate electrical switches for traffic routing decreases, resulting in a lower deployment cost as well as a reduced amount of power consumption.

Author Contributions: C.C. and B.S. contributed equally to this work, mainly focusing on the network topologies exploiting reconfigurable FSO links. A.B. and J.B. contributed to the theoretical calculations, I.T. contributed to networking aspects, C.G.H.R. and P.V.D. provided the description and technical specifications of the TriPleX chips, D.S. proposed the core idea and provided overall technical leadership. C.C. wrote the manuscript with contributions from all authors.

Funding: The research leading to these results was partially received funding from the Catalan Institution for Research and Advanced Studies (ICREA).

Acknowledgments: We acknowledge our partners from Albis, III-V Lab, Mellanox for fruitful discussions.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Singh, A.; Germano, P.; Kanagala, A.; Provost, J.; Simmons, J.; Tanda, E.; Wanderer, J.; Hölzle, U.; Stuart, S.; Vahdat, A.; et al. *Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network*; ACM Press: New York, NY, USA, 2015; pp. 183–197.
2. Al-Fares, M.; Loukissas, A.; Vahdat, A. A scalable, commodity data center network architecture. *ACM SIGCOMM Comput. Commun. Rev.* **2008**, *38*, 63–74. [[CrossRef](#)]
3. Kim, J.; Dally, W.J.; Abts, D. Flattened butterfly: A cost-efficient topology for high-radix networks. *ACM SIGARCH Comput. Architect. News* **2007**, *35*, 126–137. [[CrossRef](#)]
4. Vahdat, A.; Al-Fares, M.; Farrington, N.; Mysore, R.N.; Porter, G.; Radhakrishnan, S. Scale-Out Networking in the Data Center. *IEEE Micro* **2010**, *30*, 29–41. [[CrossRef](#)]
5. Farrington, N.; Andreyev, A. Facebook's Data Center Network Architecture. In Proceedings of the 2013 Optical Interconnects Conference, Santa Fe, NM, USA, 5–8 May 2013; pp. 49–50. [[CrossRef](#)]
6. Cisco. *Data Centre Top-of-Rack Architecture Design*; Cisco: San Jose, CA, USA, 2009.
7. Shariati, B.; Diamantopoulos, N.P.; Klonidis, D.; Comellas, J.; Tomkos, I. On the benefits of FMF based data center interconnection utilizing MIMO-less PAM-M transceivers. In Proceedings of the 2017 19th International Conference on Transparent Optical Networks (ICTON), Girona, Spain, 2–6 July 2017.
8. "Optimized airflow assessment for cabling", IBM, 2007. Available online: <https://www-935.ibm.com/services/us/its/pdf/oaac-ds-gtd01332-usen-00-053107.pdf> (accessed on 9 September 2017).
9. Agarwal, R.; Mudigonda, J.; Yalagandula, P.; Mogul, J.C. *An Algorithmic Approach to Datacenter Cabling*; Hewlett-Packard Development Company, L.P.: Palo Alto, CA, USA, 2015.
10. Zhang, W.; Zhou, X.; Yang, L.; Zhang, Z.; Zhao, B.Y.; Zheng, H. *3D Beamforming for Wireless Data Centers*; ACM Press: New York, NY, USA, 2011; pp. 1–6.
11. Cao, Z.; Proietti, R.; Clements, M.; Yoo, S.J.B. Experimental Demonstration of Flexible Bandwidth Optical Data Center Core Network with All-to-All Interconnectivity. *J. Lightwave Technol.* **2015**, *33*, 1578–1585. [[CrossRef](#)]
12. Hamedazimi, N.; Qazi, Z.; Gupta, H.; Sekar, V.; Das, S.R.; Longtin, J.P.; Shah, H.; Tanwer, A. *FireFly: A Reconfigurable Wireless Data Center Fabric Using Free-Space Optics*; ACM Press: New York, NY, USA, 2014; pp. 319–330.
13. Ramachandran, K.; Kokku, R.; Mahindra, R.; Rangarajan, S. *60 GHz Data-Center Networking: Wireless=>Worry Less*; NEC Technical Report; NEC Laboratories America: Princeton, NJ, USA, 2008; pp. 1–11.
14. Ghobadi, M.; Kilper, D.; Mahajan, R.; Phanishayee, A.; Devanur, N.; Kulkarni, J.; Ranade, G.; Blanche, P.-A.; Rastegarfar, H.; Glick, M. *ProjecToR: Agile Reconfigurable Data Center Interconnect*; ACM Press: New York, NY, USA, 2016; pp. 216–229.
15. Shariati, B.; Bogris, A.; Dijk, P.V.; Roeloffzen, C.G.; Tomkos, I.; Syvridis, D. Free Space Intra-Datacenter Interconnection Utilizing 2D Optical Beam Steering. In Proceedings of the Optical Fiber Communication Conference 2018, San Diego, CA, USA, 11–15 March 2018; p. Tu2I-7.
16. Mellette, W.M.; Ford, J.E. Scaling limits of MEMS beam-steering switches for data center networks. *J. Lightwave Technol.* **2015**, *33*, 3308–3318. [[CrossRef](#)]
17. Wörhoff, K.; Heideman, R.G.; Leinse, A.; Hoekman, M. TriPleX: A versatile dielectric photonic platform. *Adv. Opt. Technol.* **2015**, *4*. [[CrossRef](#)]
18. Mekis, A.; Gloeckner, S.; Masini, G.; Narasimha, A.; Pinguet, T.; Sahni, S.; De Dobbelaere, P. A Grating-Coupler-Enabled CMOS Photonics Platform. *IEEE J. Sel. Top. Quantum Electron.* **2011**, *17*, 597–608. [[CrossRef](#)]
19. Hulme, J.C.; Doylend, J.K.; Heck, M.J.R.; Peters, J.D.; Davenport, M.L.; Bovington, J.T.; Coldren, L.A.; Bowers, J.E. Fully integrated hybrid silicon two dimensional beam scanner. *Opt. Express* **2015**, *23*, 5861–5874. [[CrossRef](#)] [[PubMed](#)]
20. Hosseini, N.; Dekker, R.; Hoekman, M.; Dekkers, M.; Bos, J.; Leinse, A.; Heideman, R. Stress-optic modulator in TriPleX platform using a piezoelectric lead zirconate titanate (PZT) thin film. *Opt. Express* **2015**, *23*, 14018. [[CrossRef](#)] [[PubMed](#)]
21. Svelto, O.; Hanna, D.C. *Principles of Lasers*; Springer US: Boston, MA, USA, 1995; ISBN 9781461576679.
22. Albis Optoelectronics AG. Available online: <http://www.albisopto.com/> (accessed on 10 June 2018).

23. Yan, Y.; Shu, Y.; Saridis, G.M.; Rofoee, B.R.; Zervas, G.; Simeonidou, D. FPGA-based optical programmable switch and interface card for disaggregated OPS/OCS data centre networks. In Proceedings of the 2015 European Conference on Optical Communication (ECOC), Valencia, Spain, 27 September–1 October 2015; pp. 1–3.
24. Singla, A.; Hong, C.Y.; Popa, L.; Godfrey, P.B. *Jellyfish: Networking Data Centers, Randomly*; NSDI: San Jose, CA, USA, 2012; Volume 12, p. 17.



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).