Handling Rack Vibrations in FSO-based Data Center Architectures

Max Curran, Kai Zheng, Himanshu Gupta, Jon Longtin
Stony Brook University

***Abstract*—**
To overcome the shortcomings of traditional static (wired) data center (DC) architectures, there have been recent proposals of fully-wireless and reconfigurable architectures, e.g., FireFly, ProjecTor, based on Free-Space Optical (FSO) wireless links. However, there are significant challenges that need to be addressed to make the vision of FSO-based DC architectures a compelling reality. While some of these scientific challenges have been addressed in recent works, one key challenge that has yet to be addressed is how to handle FSO link misalignments due to DC rack vibrations (where the FSO transceivers are placed). The focus of this paper is to comprehensively address this challenge. Particularly, in this work, we present measurement results of a thorough study conducted over a live DC to characterize rack vibrations, and design a novel tracking and pointing (TP) system that is based on received power feedback and has a zero exposed-footprint on the deployment platform (racks). We develop and test a reconfigurable FSO link with our designed TP system and evaluate it over expected rack vibrations; our evaluation results demonstrate the effectiveness of our TP system.

I. Introduction

Data centers (DCs) are a critical piece of today's networked applications in both private and public sectors (e.g., [5], [6], [11], [12], [14]). A robust *datacenter network fabric* is fundamental to the success of DCs and to ensure that the network does not become a bottleneck for high-performance applications [30]. In this context, DC network design must satisfy several goals: high performance [16], [27], low equipment and management cost [16], [37], robustness to dynamic traffic patterns [38], [43], [28], [41], incremental expandability to add new servers or racks [22], [39], and other practical concerns such as cabling complexity [35], and power and cooling costs [24], [36].

Traditional data center architectures have been based on wired networks; being *static* in nature, these networks have either been (i) *overprovisioned* to account for worst-case traffic patterns, and thus incur high cost (e.g., fat-trees or Clos [19], [27], [16]), or (ii) *oversubscribed* (e.g., simple trees or leaf-spine architectures [17]) which incur low cost but offer poor performance due to congested links. Recent works have tried to overcome the above limitations by augmenting a static (wired) "core" with some flexible links (RF-wireless [28], [43] or optical [41], [18]). These *augmented* architectures show promise, but have offered only incremental improvement in performance due to various limiting factors. Furthermore, all the above architectures incur high cabling cost and complexity [35].

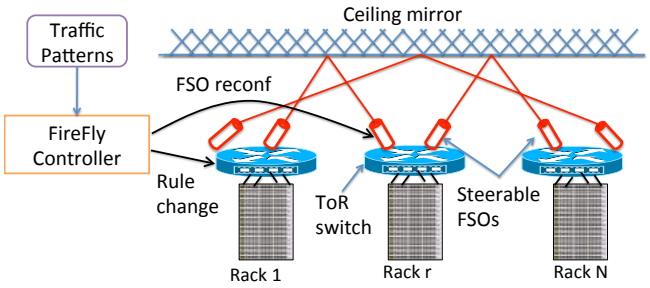

Fig. 1: High-level view of FireFly.

FSO-Based DC Architectures. To overcome the above cost-performance trade-offs and rigidity of past DC architectures, recent work [29], [26] has proposed an *extreme* design point—a *fully flexible*, *all-wireless* inter-rack fabric using *Free-Space Optics (FSO)* communication links. The FSO communication technology is particularly well-suited as it can offer very high data rates (tens of Gbps) over long ranges (>100m) using low transmission power and with small interference footprint [33].

Fig. 1 shows a conceptual overview of *FireFly*, the original FSO-based DC architecture which was proposed by our research group [29]. A number of FSO devices are placed on top of each rack and are connected to the top-of-the-rack switch. Each FSO device assembly is capable of precise and fast steering to connect to FSO devices on other racks. The controller intelligently reconfigures these devices in real-time to adapt to changing network requirements. Since the FSO beams may be obstructed by other devices in the system, FireFly proposes use of a ceiling mirror for beam redirection to ensure clear line-of-sight; in our recent work, we developed alternate line-of-sight techniques that preclude use of a ceiling mirror [21].

There are significant challenges that need to be addressed to make the vision of FSO-based DC architectures a compelling reality. While some of these scientific challenges have been addressed in recent works (e.g., [29], [21], [26]), one key challenge that has yet to be addressed is how to handle FSO link misalignments due to vibrations of DC racks where the FSO transceivers are placed. The focus of this paper is to comprehensively address this challenge.

Handling Rack Vibrations in FSO-Based DCs. Proposed FSO-based DC architectures place FSO transceivers on top of DC racks. Since the racks can vibrate for a variety of reasons (e.g., moving parts in the servers, building and other external

978-3-903176-07-2 © 2018 IFIP

vibrations such as HVAC sysems, humans), the FSO links may fail temporarily due to misalignments as they require very precise alignment for operation. One of the ways to handle this challenge is to use an active tracking and pointing (TP) system which actively corrects for such misalignments based on some feedback. However, in our context wherein tens of FSO devices need to be placed on top of racks with limited space, the key challenge in designing a viable TP system is to ensure that it has minimal/zero physical footprint. In our work, we show that favorable factors in the DC context such as indoor environment and relatively short link ranges make such a TP design feasible.

Paper Contributions and Organization. In the above context, this paper makes the following contributions:

- *Rack Vibration Measurements (§III).* To characterize typical DC rack vibrations, we conduct a measurement study and analysis via motion data collection using accelerometers and IMUs placed on racks in a live DC in our university.
- *Tracking and Pointing System (§IV).* We design a novel tracking and pointing system based on tracking-feedback from received power strength and with zero exposed-footprint, to handle any misalignments due to rack vibrations.
- *Testbed and Evaluation (§V).* We build a link testbed with the proposed TP system and evaluate it for expected misalignments based on our rack-vibration measurement study and analysis.

II. Background and Related Work

In this section, we give an overview of an FSO-based DC architecture with more detailed description of SFP-based FSO link design, and discuss related work.

FSO-based DC Networks. As mentioned in the previous section, FSO-based DC architectures are fully-flexible all-wireless and are based on the key insight that flexibility can facilitate near-optimal performance when done right. The FSO-based DCs are comprised of the following key components, viz., the FSO devices, link steering mechanisms, and the network management techniques. FSO devices needed to create FSO communication links in these architectures need to have a small form factor (so a few tens of them can fit on the top of a rack) and provide high data rates at ranges of up to 50-100m. Prior works [29] demonstrated a design of an FSO link prototype based on SFPs that satisfies the desired requirements; we discuss more details of such SFP-based FSO links in the next paragraph.

For network reconfigurability, the FSO devices are equipped with a mechanism to steer the laser beam from one receiver to the next; for viable performance, this steering must incur very low latency, i.e., on the order of a few milliseconds. In [29], two types of steering mechanisms, viz., switchable mirrors (SMs) and Galvo mirrors (GMs), were explored and their feasibility for FireFly demonstrated, while [26] has explored the feasibility of Digital Micromirror Device (DMD) as a steerable mechanism. Network management of these FSO-based DC architectures involves network design at two different timescales: (i) Preconfiguration of the network with an appropriate number of FSO devices with appropriately pre-oriented steering mechanisms; this preconfiguration is done at coarse (e.g., weekly) timescales and determines possible topologies that can be activated in real-time. (ii) Runtime reconfiguration of the pre-configured network which selects a *runtime topology* and engineers real-time traffic, based on the prevailing network state. These networks offer unprecedented benefits, such as reduced infrastructure cost, increased flexibility, and decreased cabling complexity, and have been shown to perform nearly the same as optimal wired networks.

SFP-based FSO Link Design. Prior works [29], [26], [23], [20] have designed and prototyped SFP-based FSO links to demonstrate feasibility of small-form factor FSO devices suited for FSO-based DC architectures. Below, we discuss this in more detail, as our testbed and TP system builds upon this design. An SFP (small form-factor pluggable) transceiver is a small ($1/2'' \times 1/2'' \times 2''$) and compact commodity optical transceiver [4], widely used to interface optical fibers with network switches. An SFP contains a laser source and a photodetector, for transmitting and receiving respectively. SFPs are available with a variety of laser sources, varying in the wavelength (typically, between 800nm to 1550nm) of the emitted beam as well as the supported data rate (anywhere from 10Mbps to recent variants called CFPs with 100Gpbs [1]). SFP+ refers to an enhanced version of the SFP that supports data rates up to 16Gbps. To create an FSO link using SFPs, the beam emanating from the transmitter SFP is channeled into a short optical fiber which feeds into a collimator. The collimated beam is then launched into free space towards the receiving SFP, where it is captured by another collimating lens and focused back into an optical fiber connected to the receiving SFP.

Related Work. To the best of our knowledge, the only work on measurement of rack vibrations is [40]; however, the focus of [40] is on the impact of vibrations on server hard drive failure, by characterizing vibration level. In contrast, we wish to evaluate impact of rack vibrations on FSO link alignments, and thus, our focus is on measuring motion related characteristics of rack vibrations.

Typical TP mechanisms [31], [32] include a fast steering mirror or gimbal controlled by digital servos [34], [15], [42] for pointing (recent project [7] uses a high-cost SLM), along with some tracking detectors to track the target or beam. Common tracking detectors include positioning sensing diodes [15] (e.g., CCDs [34], photodiode arrays [25]), accelerometers, cameras, GPS [42], etc. Our recent work [20] used a TP mechanism based on GMs and photodiodes in the context of outdoor picocell networks. In our context, we are most interested in developing a TP system that has minimal "exposed" footprint to allow placement of a large number (50+) of devices on top of each DC rack. Thus, we propose a novel tracking mechanism based on (i) the RSSI (relative

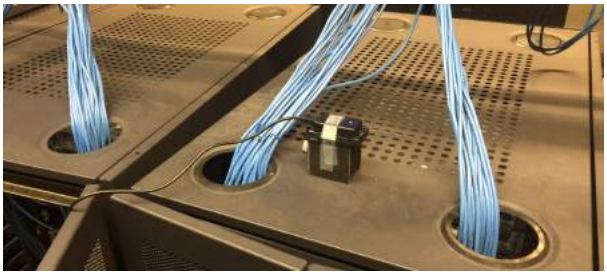

Fig. 2: Inertial measurement unit (IMU) on top of a DC rack.

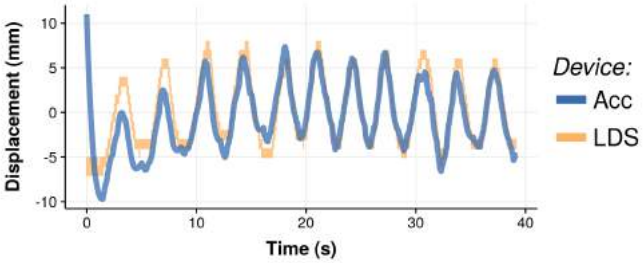

Fig. 3: Difference between the displacement as measured by integration of accelerometer's acceleration and the displacement measurement from the laser displacement sensor.

received signal strength) feedback for tracking (without using any exposed hardware), and (ii) already available steering mechanisms, e.g., GMs, in the FSO links.

III. Vibration Analysis

In this section, we present results of our measurement study conducted over racks in a live DC in our university, to better understand and characterize the vibrations experienced by devices placed on top of the DC racks. We will use the results of this study to evaluate the effectiveness of our proposed TP system.

Measurement Study Setup. To collect rack vibration measurements, we install an accelerometer [2] and an IMU (inertial measurement unit) [8] over various positions on or near racks in a live data center housed in the CEWIT center [3] in our university. The accelerometer measures linear acceleration (and hence, linear displacement via integration) in the three spatial dimensions, while the IMU measures the three angular/rotational accelerations, viz. pitch, yaw, and roll. Thus, together these measurement devices cover all possible motions/vibrations experienced by an object on the rack. The accelerometer and IMU gather measurements at a rate of 512 Hz and 250 Hz respectively, which is sufficient for our purposes.[1] We gather measurements with these devices placed on top and side of three different racks for a continuous 24-hour period.

Data Analysis in Frequency Domain. Once all of the data was collected, the Fourier transform was applied to both the accelerometer and IMU data for data analysis in the frequency domain and to filter out low-frequency noise. In particular, the data in the frequency domain is sent through an appropriately designed Butterworth bandpass filter to remove any low-frequency noise. Then, we integrate the filtered data from both devices to get velocity and displacement measurements.

Integration Validation. In order to measure the accuracy of our integration process, we conduct an experiment in the lab using the accelerometer and a laser displacement sensor (LDS) [10]. We move the accelerometer by hand and measure its actual displacement by the LDS which uses a laser to accurately measure displacement. Then, we compare the displacement measured by LDS with the displacement as computed by integration of the accelerometer measurements; see Figure 3. We observe that the difference between the two displacement results is minimal. This shows that the errors in the displacement result from integrating accelerometer's acceleration data do not accumulate over time, thus validating our integration analysis.

Vibration Results. Below, we summarize our vibration measurement results as related to FSO link's intrinsic movement tolerance.

- Linear displacement (see Figure 4a) was found to be minimal (maximum 0.25mm). This amount of linear displacement can be easily handled by a link's intrinsic movement tolerance (a few millimeters [29]). Moreover, this is subsumed by the angular displacement for links longer than one meter. Thus, we don't analyze linear speed of displacement.
- Angular displacement (see Figure 4b) was observed to be at most 1.5 mrad with an average of 0.9 mrad. This is more than sufficient to disconnect a link with no TP system. For example, on a 10m link, a 1.5 mrad angular deviation of the transmitter would cause a 1.5cm linear displacement at the receiver. This demonstrates a need to have an effective TP system to keep an FSO link continuously operational on DC racks.
- To evaluate the effectiveness of a TP mechanism, we must focus on the angular *speed* of movement, since the TP system incurs a non-zero latency and can therefore only be effective up to some angular speed [20]. Thus, we analyze the angular speed of movement caused by rack vibrations (see Figure 4c); the average rotation speed was found to be 3.30 mrad/s with a maximum being 6.98 mrad/s.

IV. Tracking and Pointing (TP) System

FSO links require precise alignment to function properly, and link misalignment can cause the entire link to stop functioning. Link misalignment is fundamentally caused by movement of the beam at the receiver plane, which is caused by the movement of the FSO transceivers. In a DC environment, rack vibrations can cause the FSO transceivers to deviate from their original positions and thus cause link disconnection. To counter such link misalignments due to rack vibrations, we use an active tracking and pointing (TP) system which uses a *tracking* mechanism to track the beam movement at the

[1] Higher frequency vibrations, if any, will have negligible displacements.

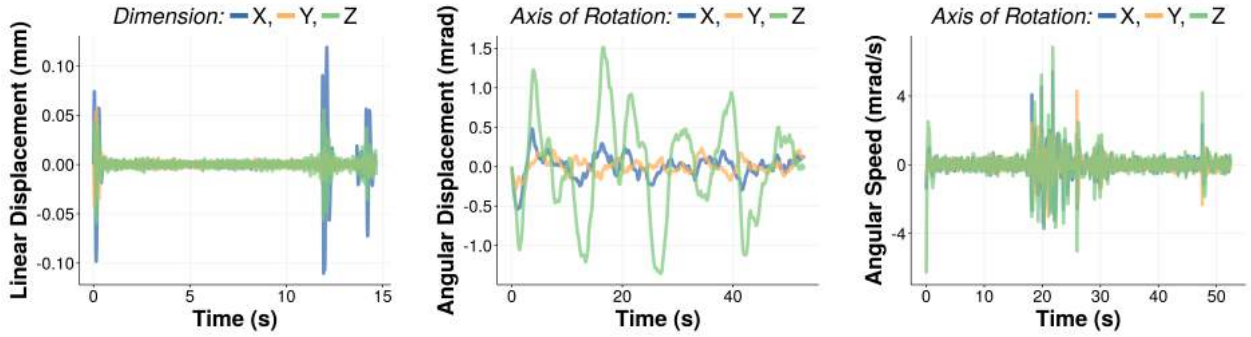

(a) (b) (c)

Fig. 4: Vibration measurements from the accelerometer and IMU placed on top of a DC rack. Here, each plot shows only a small slice of time.

RX and provides feedback to the *pointing* mechanism which corrects any misalignments.

Requirements. The key requirement for our TP systems is that it should be able to keep an FSO link aligned in response to expected rack vibrations, as characterized in the previous section. Any TP system can only be effective up to a certain speed of beam movement in the receiver plane, due to non-zero latency. To relate the beam movement in the receiver plane to motion caused by rack vibrations, we focus on the most dominating component of the vibrations, viz., angular speed and displacement of the TX assembly due to rack vibrations. Thus, we require that the TP system be able to handle vibrations that have an angular speed of at most 7 mrad/s, with an angular displacement of at most 2 mrad. In addition, for a viable TP system in our context, the TP hardware should have a minimal exposed footprint on the top of the rack.

A. Native RSSI-based TP

We propose the following lightweight TP system for use in DC environments:

Tracking Beam Position/Displacement. The goal of the tracking mechanism is to estimate the beam displacement, and use it to provide an appropriate feedback to the pointing mechanism which then corrects the beam alignment. To estimate the beam displacement, we use the RSSI (relative received signal strength) value corresponding to the received power of the beam; this RSSI value is available by "querying" the SFP+ module via the available I^2C interface. In our experiments, we used an SFP+ evaluation board [13] to interface with the SFP+ module. Since the evaluation board doesn't need to be exposed (or even placed) on the rack, it has no exposed footprint on the top of the rack. Note that due to the symmetry of the Guassian beam, an RSSI value (or the RX power) corresponding to the current beam "position" is not sufficient to uniquely determine beam's position/displacement. To uniquely determine beam's current position C, we gather RSSI values at *multiple positions around* C, and use these values to uniquely determine C. For example, we can gather RSSI values at C, and at beam positions slightly "north" of C, "south" of C, "east" of C, and "west" of C. (In fact, only 3 of these 5 values are needed to uniquely determine the beam's position/displacement, but more values may be needed to compensate for noise.). To acquire RSSI values at the position C' around C, we "intentionally" move the steering mechanism (GM, in our case) to the desired position C' and read the RSSI value from the SFP+ module. The beam shape and the noise level will help determine the exact positions relative to C that are most effective for our purposes. Once the RSSI values at these nearby positions around C have been recorded, we use either a precomputed-table or a gradient based control algorithm (see below) to determine the correction to be applied to the beam to move it back to the original (aligned) position.

Correction Algorithms. The goal of a correction algorithm is to take the RSSI values from multiple positions around the current beam position C, compute a correction to be applied to the beam, and then apply the correction via the steering mechanism to align the beam back to the original (aligned) position (this is the *pointing mechanism*). This overall process ensures continuous operation of the link at runtime. More formally, a correction algorithm is given a list of $n+1$ tuples, viz. (x_i, y_i, r_i) for $0 \leq i \leq n$, where (x_i, y_i) is the relative position to C, the current beam location, and r_i is the RSSI value at this relative location. We assume that $i = 0$ refers to the current beam location C (i.e., x_0 and y_0 are both 0). Below, we describe two correction algorithms that use these $n+1$ tuples to estimate a beam correction; one of them uses a training phase to precompute a table with (position, RSSI) values, while the other uses the relative differences in RSSI values to compute the beam correction.

Table-Based Correction Algorithm. In this approach, we first have a training phase wherein we query and store (in a table) RSSI values for a wide range of beam locations. The locations queried during the training phase should encompass the entire beam profile, and preferably should be at the finest granularity possible. At runtime, the input $(n+1)$ tuples are then "compared" with the table rows as below to find the row in the table that most closely represents C. In particular, the best match returned by the algorithm is the absolute location

(x_c, y_c) that minimizes the following quantity:

$$\left\| \begin{pmatrix} r_0 \\ \vdots \\ r_n \end{pmatrix} - \begin{pmatrix} T(x_c + x_0, y_c + y_0) \\ \vdots \\ T(x_c + x_n, y_c + y_n) \end{pmatrix} \right\| \quad (1)$$

where $T(x, y)$ is the RSSI value in the table at location (x, y). As (x_c, y_c) is approximately the current location of the beam (with original aligned position being (0,0)), we supply a value of $(-x_c, -y_c)$ to the steering mechanism to move the beam back to the original aligned position. Note that, if the beam is continuously moving, the correction may result in beam moving back to some position *close* to the original aligned location.

Gradient-Based Correction Algorithm. The above table-based correction algorithm's performance depends on the training phase, and thus, may suffer if there is noise or imperfections in the system that changes the RSSI values from what was observed in the training phase. Thus, we also experimented with a gradient-based correction approach, wherein the estimation of correction is (largely) based only on the input tuples. In particular, the algorithm computes the "gradient" within the input tuples, and uses this to move the beam by a constant/input value. Implicitly, the algorithm uses the fact that the beam profile is approximately Gaussian, i.e., like a "hill." Thus, the relative difference between the positions around C can be used to correct the beam from C back to the original position.

More formally, given the $n+1$ tuples (x_i, y_i, r_i), as defined above, we first estimate a gradient $G(i)$ for each $1 \leq i \leq n$ as:

$$G(i) = \frac{r_0 - r_i}{\sqrt{x_i^2 + y_i^2}} \quad (2)$$

We then use these $G(i)$ values to compute the overall correction as follows. First, to account for noise, we only use values of $G(i)$ that are large enough; in particular, for each $G(i)$, we calculate its contribution to the correction $C(i)$ as:

$$C(i) = \begin{cases} V, & \text{if } G(i) \geq K \\ 0, & \text{if } |G(i)| < K \\ -V, & \text{if } G(i) \leq -K \end{cases} \quad (3)$$

Above, V and K are appropriately picked constants. Now, the overall correction (x_c, y_c) is calculated as follows:

$$\begin{pmatrix} x_c \\ y_c \end{pmatrix} = \sum_{i=1}^{n} \frac{C(i)}{\sqrt{x_i^2 + y_i^2}} \begin{pmatrix} -x_i \\ -y_i \end{pmatrix} \quad (4)$$

Note that unlike the previous Table-based algorithm, the goal of the Gradient-based algorithm is not to correct the beam all the way back to the original aligned position (as it doesn't have sufficient information), but to simply move the beam back in the direction of the aligned position.

V. Results

We now describe and evaluate the TP system using a SFP-based FSO link prototype.

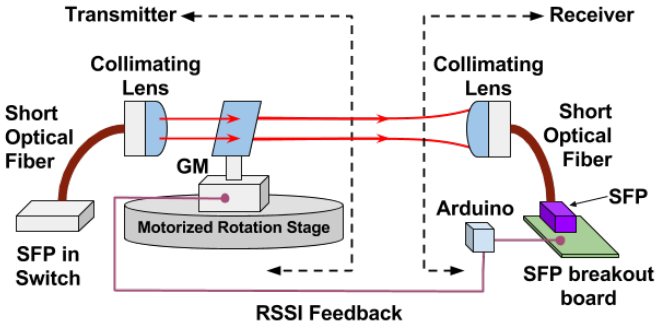

Fig. 5: Schematic of the FSO link tested with the TP system.

FSO Link with TP System Testbed. We setup our FSO link prototype (shown in Figure 5) in a controlled environment with nearly no natural movement, so that we can artificially recreate the movements caused by rack vibrations. We use 10G 1550nm SFP+ with 10 GBASE-ZR interface. We create a 10m unidirectional FSO link, with the other direction connected via a long fiber cable for simplicity. The transmitter assembly is equipped with a GM that is used as the pointing mechanism of the TP system. The entire transmitter assembly is placed on a motorized rotational stage, which allows us to simulate rotational rack vibrations. To access the pins of the SFP+ directly (and query the RSSI values), we use a Timbercon SFP+ evaluation board [13] at the receiver. We connect to the evaluation board a custom Arduino microcontroller, which uses the I^2C protocol to fetch the RSSI from the receiver SFP+.

Experimental Results. For our experiments, we used two nearby positions around any current beam position C as input to the correction algorithms; in particular, the nearby positions used were north and east at an "angular distance" of 0.2 mrad from C. To demonstrate the link operation during continuous terminal movement due to rack vibrations, we compute the link's TCP throughput with TP active and TX assembly rotating at varying angular speeds of 0-7 mrad/sec, using a motorized rotational stage, with an amplitude of a few mrads. In particular, we measure the average link throughput (data rate) every second, over a ten minute period using the iPerf3 tool [9].

First, we observed (plots not shown for brevity) that the Table-based algorithm outperformed the Gradient-based algorithm, and thus, we focus on the Table-based algorithm below. Figure 6 shows the CDF of the FSO link's throughput with the TP system active and TX assembly rotating at varying angular speeds (0 to 7 mrad/s). The fixed link (i.e., for 0 mrad/sec speed) achieved an overall average throughput of 9.41 Gbps. We were also able to achieve nearly identical throughput CDF for angular speeds of up to 2.5 mrad/s. Throughput CDF begins to degrade very slightly for angular speed between 4-7 mrad/s, with the average throughput still being near-optimal at about 9.37 Gbps for angular speeds up to 6 mrad/s and about 8.5 Gbps at 7 mrad/s angular speed. To offer further perspective on these results, we analyze our vibration measurement data further and observe that even though the

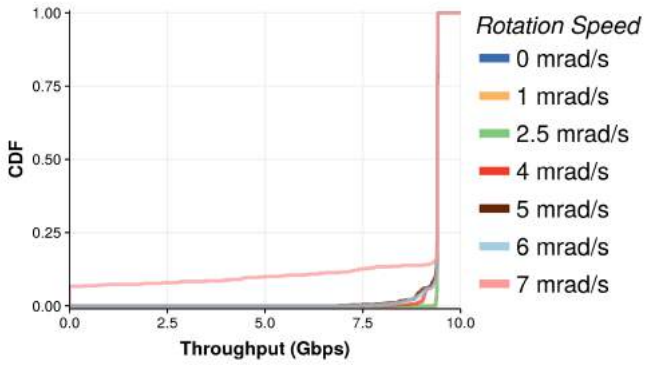

Fig. 6: CDF of the link throughput for various angular speeds of the TX assembly.

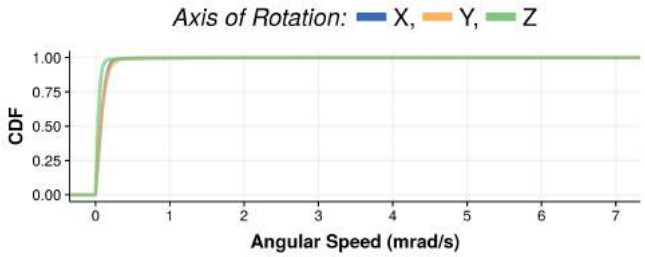

Fig. 7: CDF of the angular speeds observed in our vibration measurement study.

maximum angular speed observed was 6.98 mrad/s, the angular speed went above 6 mrad/s very rarely: only 0.00816% of the time. See Figure 7, which plots the CDF of the angular speeds observed. Thus, in summary, our TP system achieves near-optimal throughput at about 99.992% of the time, with a 10% throughput degradation at remaining times.

VI. Conclusion

In this work, we addressed one of the key challenges that arise in the context of recently proposed FSO-based data center architectures. In particular, we conducted a thorough measurement study of DC rack vibrations, and proposed a novel RSSI-based TP system with zero exposed-footprint that can handle expected rack vibrations. To the best of our knowledge, ours is the best work on handling DC rack vibrations for reliable operation of FSO links placed on DC racks.

References

[1] 100G CFP Optical Transceivers. https://tinyurl.com/yczvt6x7.
[2] Accelerometer. http://www.gcdataconcepts.com/x2-1.html.
[3] CEWIT. http://www.cewit.org/.
[4] Cisco 10GBASE SFP+ Modules. https://tinyurl.com/y894eh9a.
[5] Cisco Global Cloud Index: Forecast and Methodology, 2012 to 2017. http://tinyurl.com/7gnfeeb.
[6] Data center survey. https://tinyurl.com/ycr77yvz.
[7] Hyperion Project, UK. http://projecthyperion.co.uk.
[8] Inertial Measurement Unit. https://tinyurl.com/ybw58762.
[9] iPerf. https://iperf.fr/.
[10] Laser displacement sensor. https://tinyurl.com/yac9fzkj.
[11] Magic quadrant for data center network infrastructure. http://tinyurl.com/mpo3jzt.
[12] NSA Utah Data Center. http://nsa.gov1.info/utah-data-center/.
[13] Timbercon - SFP+ Test Host Board. https://tinyurl.com/y84tvmx9.
[14] US government gives IBM cloud green light to serve agencies. http://tinyurl.com/mx2v2mt.
[15] M. K. Al-Akkoumi. A tracking system for mobile FSO. In *SPIE Proceedings Vol. 6877*, 2009.
[16] M. Al-Fares, A. Loukissas, and A. Vahdat. A scalable, commodity data center network architecture. In *SIGCOMM*. ACM, 2008.
[17] M. Alizadeh and T. Edsall. On the data path performance of leaf-spine datacenter fabrics. In *High-Performance Interconnects (HOTI)*. IEEE, 2013.
[18] K. Chen et al. Osa: An optical switching architecture for data center networks with unprecedented flexibility. *IEEE/ACM Transactions on Networking (TON)*, 2014.
[19] C. Clos. A study of non-blocking switching networks. *Bell Labs Technical Journal*, 1953.
[20] M. Curran et al. Fsonet: A wireless backhaul for multi-gigabit picocells using steerable free space optics. In *MobiCom*. ACM, 2017.
[21] M. Curran and H. Gupta. Providing line-of-sight in a free-space-optics based data center architecture. In *Communications (ICC), International Conference on*. IEEE, 2016.
[22] A. R. Curtis et al. Legup: Using heterogeneity to reduce the cost of data center network upgrades. In *Co-NEXT*. ACM, 2010.
[23] P. Deng et al. MEMS-based beam steerable free space optical communication link for reconfigurable wireless data center. In *Proc. of SPIE Vol*, volume 10128, pages 1012805–1, 2017.
[24] N. Farrington. *Optics in data center network architecture*. University of California, San Diego, 2012.
[25] M. S. Ferraro et al. InAlAs/InGaAs avalanche photodiode arrays for free space optical communication. *Appl. Opt.*, 2015.
[26] M. Ghobadi et al. Projector: Agile reconfigurable data center interconnect. In *SIGCOMM*. ACM, 2016.
[27] A. Greenberg et al. Vl2: a scalable and flexible data center network. In *SIGCOMM*. ACM, 2009.
[28] D. Halperin et al. Augmenting data center networks with multi-gigabit wireless links. In *SIGCOMM*. ACM, 2011.
[29] N. Hamedazimi et al. Firefly: A reconfigurable wireless data center fabric using free-space optics. In *SIGCOMM*. ACM, 2014.
[30] J. Hamilton et al. Data center networks are in my way. *Standford Clean Slate CTO Summit*, 2009.
[31] T.-H. Ho. *Pointing, Acquisition, and Tracking Systems for Free-Space Optical Communication Links*. PhD thesis, University of Maryland, College Park, 2007.
[32] S. V. Kartalopoulos. *Free Space Optical Networks for Ultra-Broad Band Services*. John Wiley and Sons, 2001.
[33] D. Kedar and S. Arnon. Urban optical wireless communication networks: the main challenges and possible solutions. *IEEE Communications Magazine*, 2004.
[34] C. Lv et al. Implementation of FTA with high bandwidth and tracking accuracy in FSO. In *International Conference on Consumer Electronics, Communications and Networks (CECNet)*, 2012.
[35] J. Mudigonda et al. Taming the flying cable monster: A topology design and optimization framework for data-center networks. In *USENIX Annual Technical Conference*, 2011.
[36] R. Niranjan Mysore et al. Portland: a scalable fault-tolerant layer 2 data center network fabric. In *SIGCOMM*. ACM, 2009.
[37] L. Popa et al. A cost comparison of datacenter network architectures. In *Co-NEXT*. ACM, 2010.
[38] A. Singla et al. Proteus: a topology malleable data center network. In *ACM SIGCOMM Workshop on Hot Topics in Networks*, 2010.
[39] A. Singla et al. Jellyfish: Networking data centers, randomly. In *NSDI*, 2012.
[40] X. Tan et al. An advanced rack server system design for rotational vibration (rv) performance. In *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Intersociety Conference on*. IEEE, 2016.
[41] G. Wang et al. c-through: Part-time optics in data centers. In *SIGCOMM*. ACM, 2010.
[42] T. Yamashita et al. The new tracking control system for free-space optical communications. In *2011 International Conference on Space Optical Systems and Applications (ICSOS)*, 2011.
[43] X. Zhou et al. Mirror mirror on the ceiling: Flexible wireless links for data centers. *SIGCOMM*, 2012.