

Can open-loop control algorithms solve the data center switching problem?

C. POLLOCK,^{1*} F. PARDO,² M. IMBODEN,³ AND D.J. BISHOP^{1,4,5,6,7}

¹*Department of Mechanical Engineering, Boston University, Boston, MA, USA*

²*Nokia Bell Labs, Murray Hill, NJ, USA*

³*Institute of Microengineering, École Polytechnique Fédérale de Lausanne (EPFL), Neuchâtel, Switzerland*

⁴*Division of Material Science, Boston University, Boston, MA, USA*

⁵*Department of Physics, Boston University, Boston, MA, USA*

⁶*Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA*

⁷*Department of Biomedical Engineering, Boston University, Boston, MA, USA*

*cpollock@bu.edu

Abstract: There is a world-wide push to create the next generation all-optical transmission and switching technologies for exascale data centers. In this paper we focus on the switching fabrics. Many different types of 2D architectures are being explored including MEMS/waveguides and semiconductor optical amplifiers. However, these tend to suffer from high, path dependent losses and crosstalk issues. The technologies with the best optical properties demonstrated to date in large fabrics (>100 ports) are 3D MEMS beam steering approaches. These have low average insertion losses, and equally important, a narrow loss distribution. However, 3D MEMS fabrics are generally dismissed from serious consideration for this application because of their slow switching speeds (~few milliseconds) and costs (\$100/port). In this paper we show how novel feedforward open loop controls can solve both problems by improving switching speeds by two orders of magnitude and costs by one order of magnitude. With these improvements in hand, we believe 3D MEMS fabrics can become the technology of choice for data centers.

1. Introduction

The electrical power consumed in exascale data centers is large and growing. The largest data center today uses ~ 150 MW [1] of power and 100 MW centers are not uncommon [2]. As the world attempts to reduce its carbon footprint and use electrical power more efficiently, the size and number of data centers continues to grow exponentially [3]. In response to this threat, there is a world-wide effort to develop all-optical technologies to replace electronic transmission and switching/routing, two major consumers of power in modern data centers [4, 5].

In this paper we focus on the switching/routing aspect of the problem. In the telecom world 20 years ago, there were circuit switches and packet switches. The circuit switches brought connections up for long periods of time and the temporal requirements on those switches were driven only by the desired restoration times, typically 50 milliseconds, the SONET specification [6]. Packet switches were electronic and had switching time requirements of microseconds to nanoseconds. The desired optical switch fabric for today's data center application is a hybrid between these two traditional categories [7].

The typical architecture used today comprises thousands to tens of thousands of racks with multiple servers per rack. The top of rack (ToR) is typically a node in the switching network [8–10]. This drives the requirement of thousands of ports per fabric. As an example, the DARPA PIPES specification is 1000 ports [11]. The traffic between racks consists of packets of varying lengths, from nanoseconds to days. The longest of these packets are "elephant flows" and while they make up a small percentage of the total number of flows, they can contribute a relatively large portion of the total data traffic volume [12, 13]. The desire is to offload this traffic to an all-optical

network which can route and carry the information much more efficiently than can be done electronically [2, 4, 5, 7]. While there is no hard temporal specification, faster optical switching leads to a larger portion of flows being routed through the all-optical network and greater energy savings. The DARPA PIPES program, for example, has a specification for a 10 μ s switching time [11].

In the envisioned architecture, the switching fabric would be surrounded by optical elements of varying powers ranging from short reach (SR), long reach multimode (LRM), long reach (LR), extended reach (ER, ZR). These elements are high speed and expensive with a steep financial penalty for more optical power. This drives the need for low average losses in the fabric and equally importantly, a narrow distribution of loss for all connections. High speed optical systems suffer from limited dynamic range and a path length dependent loss is difficult and expensive to deal with. This drives the DARPA PIPES specification of 3 dB insertion loss [11]. Finally, cost drives everything in the data center. In the telecom world of yore, per port costs of \$100 - \$1000 were acceptable; whereas the target cost today is closer to \$10/port [4].

3D MEMS switching architectures are a mature, commercially available technology, developed almost 20 years ago, which meet all the relevant optical specifications. Systems have been built with an average loss of 1.3 dB for all connections in a 238x238 port system with a narrow loss distribution ($\pm .3$ dB), low polarization dependent loss of <40 dB, a spectral bandwidth of >100 nm, low (unmeasurable) crosstalk between channels, fully optically transparent supporting both classically and quantum coherent transmission [14–16]. This same technology has been expanded to create switching fabrics with 1,100 inputs and outputs [17]. The dual Achilles Heels for this technology have been speed and cost. Traditionally, 3D MEMS devices have had millisecond switching times [18–21] and cost $\sim$ \$100/port [4]. A major contributor to the high cost are the electronics to control each mirror. The typical 3D MEMS fabric is an analog device, requiring fast analog control of high voltages on the order of 100 V. This is usually done with the combination of a high speed DAC and a high speed, high voltage linear amplifier (HVamp); however, achieving high enough slew rates is a challenge. Each port in a 3D MEMS fabric requires 8 of these DAC/HVamp channels. At $\sim$ \$10/channel these systems cost \$80/port, well above the \$10/port target. Therefore, any solution must address the electronics cost per channel.

In this paper we discuss how 3D MEMS beam steering technology augmented with modern open loop controls can improve the step and settle response times by orders of magnitude. The theory for these methods was laid out by Imboden *et al.* [22] who demonstrates the feedforward double step drive in capacitive, electromagnetic, and thermally actuated MEMS mirrors. This technique improves the settling time of a high quality factor (Q) system from $\sim T_0 Q$ (ringdown time) to $T_0/2$, where T_0 is the period associated with the system's natural frequency, $\omega_0 = 2\pi/T_0$. Here we expand on that work by experimentally demonstrating the previously predicted overdrive methods using a LambdaRouter electrostatic tip-tilt mirror. The LambdaRouter is a well known device with good optical properties, allowing us to focus on improving the switching speed [14, 15, 17, 23]. We then use simulations to explore the impact of lower Q values on the response. From a combination of experimental and simulated results, we discuss how a micromirror could achieve a 10 μ s switching time. We also demonstrate the use of pulse width modulation (PWM) methods to drive these devices, which can reduce the control costs by a significant margin. We have previously shown how one can use PWM to control MEMS in an analog way, which allows the DAC/HVamp to be replaced with a simple switching circuit [24]. Timing replaces the need for analog control. Our thesis is that modern control theory approaches for driving 3D MEMS devices can eliminate the two bottlenecks that limit their general application in data centers. We believe control theory, appropriately applied to a mature, commercially available technology, can solve the data center switching problem.




Fig. 1. Test setup. (a) Double step drive. τ_1 is the torque applied in the first step and τ_0 is the final torque applied at time t_1 . (b) Illustration of the overdrive. τ_0 is the final torque. τ_{max} and τ_{min} are the maximum and minimum torques applied. t_1 is the time when the torque is switched from τ_{max} to τ_{min} . t_2 is the time the torque is switched from τ_{min} to τ_0 . (c) Colorized SEM image of a LambdaRouter mirror [23]. (d) Schematic of electrostatic tip-tilt mirror. Applying a voltage to V_1 will accelerate the mirror in one direction, whereas applying a voltage to V_2 will deflect it in the other direction.

2. Methods

2.1. Theory

In many cases, MEMS mirrors can be modeled as an underdamped system defined by the linear differential equation

$$I\ddot{\theta} + \gamma\dot{\theta} + k_t\theta = \tau(t), \quad (1)$$

where I , γ , k_t , $\tau(t)$, and θ are the system's moment of inertia, damping coefficient, torsional spring coefficient, applied torque and the mirror's angle respectively. Such a device will ring when driven by a step input and the settling time is related to the Q of the device. However, we have shown that by applying precisely timed steps or pulses to a high Q system, one can improve the settle time by orders of magnitude [22]. There are three open loop techniques discussed in this work: (1) double step (DS), (2) unipolar overdrive (UOD) and (3) bipolar overdrive (BOD), all of which can replace a single step (SS). Intuitively, all three work off the idea that an underdamped system will overshoot when given a step input. The precisely timed pulses/steps manipulate the mirror's momentum such that it reaches the desired angle, θ_0 , with zero velocity. At that point, the torque required to counter the restoring spring torque ($\tau_0 = k_t\theta_0$) is applied resulting in zero motion and zero net torque on the mirror. The double step, illustrated in Fig. 1(a), consists of a

step input of torque τ_1 for time t_1 , at which point the torque τ_0 is applied. A full derivation for the timing t_1 and torque τ_1 are presented in Imboden *et al.* [22]. They are originally derived for a translational system but the solution is general and applies to any linear, second order differential equation, such as Eq. (1). As the Q increases, t_1 and τ_1 approach $T_0/2$ and $\tau_0/2$. It should be noted that for an electrostatic mirror like the LambdaRouter, $\tau \sim V^2$ for small displacements, where τ is the torque and V is the applied voltage. The overdrive method, illustrated in Fig. 1(b), consists of a step input τ_{max} until time t_1 , then a step to τ_{min} until time t_2 , at which point the final torque τ_0 is applied. For brevity, we limit the discussion to only two versions of the overdrive (1) unipolar overdrive where τ_{min} is defined as 0 and (2) bipolar overdrive where τ_{min} is defined as $-\tau_{max}$. Similarly, the derivations for t_1 and t_2 for both overdrive methods are provided by Imboden *et al.* [22]. The analytic derivation assumes zero damping which significantly simplifies the derivation and solution. This is a good assumption when the Q factor is high ($Q > 100$). These three drive techniques, double step and the two overdrives, can be implemented with analog controls or digital controls using PWM [24, 25]. Using PWM, the torque discussed above is related to the PWM duty cycle. For small displacements, PWM linearizes the forcing function with respect to the duty cycle such that $\tau \sim R_{DC}\tau_{PWM}$ where R_{DC} is the duty cycle (0-1) and τ_{PWM} is the max torque being switched on and off [24]. There are twelve different drives discussed throughout this work and summarized in Table 1. Ten of the drives are analog and two are PWM. The normalized t_2 timing for a high Q system is shown in Table 1. The bipolar drive is not directly tested with the PWM method. However, previous work has shown that given a high enough frequency, PWM can be directly substituted for analog control and implies that the bipolar drive could be implemented with PWM [24]. Likewise, the t_2 values are based on the analog driving condition, but since PWM and analog are interchangeable given a high enough frequency, we assume the same analytic solution for both.

Table 1. Summary of the different testing conditions.

	τ_1/τ_0	Name	(t_2/T_0)	Analog			PWM	
				Theory	Simulation	Experiment	Theory	Experiment
Double Step (Inner Gimbal)	0.5:1	DS	0.500	x	-	x	x*	x
	τ_{max}/τ_0							
Unipolar (Inner Gimbal)	1:1	UOD 1x	0.333	x	x	x	x*	x
	2:1	UOD 2x	0.290	x	x	x	-	-
	3:1	UOD 3x	0.276	x	x	x	-	-
Bipolar (Outer Gimbal)	1:1	BOD 1x	0.290	x	x	x	-	
	2:1	BOD 2x	0.210	x	x	x		
	3:1	BOD 3x	0.174	x	x	x		
	1:4	BOD 4x	0.153	x	x	x		
	1:5	BOD 5x	0.137	x	x	x		
	1:9	BOD 9x	0.104	x	x	x		

* PWM theory is based on the timing and forcing of an analog signal with PWM duty cycle analogous to the force.

2.2. Experimental setup and device

This work focuses specifically on electrostatic mirrors since they are currently the most popular micromirror used by commercial optical switches as they are relatively fast and low power [4, 19]. We use the same two-axis gimbal design used in the LambdaRouter mirror array, shown in Fig. 1(c) and a simplified schematic of the setup is shown in Fig. 1(d). To generate the analog drives we use a combination of a pulse generator, summing amplifiers and high speed, HV amplifiers. To generate the PWM signal, we modulate the signal from the summing amplifier prior to

the HV amplifier. Although we use expensive components for the PWM drive in this work, these can be replaced with a relatively inexpensive controller and HV switching circuit [24, 25]. Generating a voltage difference (V_1) between the mirror and pad1 produces an electrostatic force and subsequently a torque on the mirror, tilting it in one direction. Similarly, applying a voltage difference (V_2) between the mirror and pad2 the mirror produce a torque in the other direction. We use two identical mirrors for the experiment. The first mirror uses the inner gimbal and experimentally demonstrates the double step and unipolar overdrive methods. The step response for a system defined by Eq. (1) is known. By fitting the response of our system to this known response, we can determine the natural frequency and Q factor. The first mirror has a natural frequency of 834 Hz ($T_0 = 1.2$ ms) and a Q factor of 2.48, well below the condition for which the previously derived overdrive theory is valid. The double step and the unipolar overdrive do not require negative torques, therefore they can be implemented using only pad1. The first mirror broke during the experiment so a second mirror was used for the bipolar overdrive. Unfortunately, the inner gimbal of the second mirror does not work so the outer gimbal is used. Due to the changes from the inner to the outer gimbal, the second mirror has slightly different dynamics with a natural frequency of 522 Hz ($T_0 = 1.9$ ms) and a Q factor of 1.27, again well below 100. Throughout this work the timing is normalized with respect to T_0 so the difference in frequency is not important, however the difference in Q does have a noticeable effect. The bipolar overdrive requires both positive and negative torques which is not possible with a single electrode. Pad2 is implemented to provide the negative torque required for the bipolar drive. We experimentally demonstrate all twelve open loop methods and Table 1 labels which mirror is used by specifying whether the inner or outer gimbal is used. To determine the optimal timing, we start with the theoretical t_1 and t_2 (τ_1 for double step), then iterate the drive parameters (i.e. t_1 , t_2 , τ_1) until we minimize the mirror's ringing.

2.3. Simulations

We use Euler's method to numerically solve the mirrors' responses to the overdrive conditions using phase plane trajectories [26, 27], assuming the mirror's motion is defined by Eq. (1). We start by solving the τ_{max} step response given the initial state, zero angle and zero velocity. Instead of plotting angle versus time however, we plot angle versus velocity, giving us the phase plane trajectory for τ_{max} . We then simulate the τ_{min} step in a similar way except that we start the system at the final state (θ_0 and zero velocity) and simulate backwards in time. Plotting the τ_{max} and τ_{min} trajectories together, we can find where the two curves cross and calculate the timing t_1 and t_2 necessary for the overdrive. Using this approach, we can think of moving the system from the initial state to the final state with our controls. We then verify the timing by simulating the overdrive response with τ_{max} , τ_{min} , and τ_0 . The main benefit of the simulation is that it provides insight into using the overdrive methods on systems with non-zero γ values. When simulating the conditions in Table 1, we use the same ω_0 and Q factor for the respective mirrors. As shown in Table 1, there are no simulations for the double step or the two PWM methods. The reason is that double step already has an analytic solution so the simulation does not provide additional insight. Similarly, we have previously shown a direct correlation between analog and PWM control, so we assume the same timing applies [24].

3. Results and discussion

3.1. Analog control

The experimental responses for the conditions in Table 1 are shown in Fig. 2. Figs. 2(a) and 2(b) show the angle θ/θ_0 versus the time t/T_0 for the first and second mirror respectively. Figs. 2(c) and 2(d) show the corresponding driving voltages V/V_0 for the first and second mirror, where V_0 is the final applied voltage associated with the steady-state angle θ_0 . For the experiments




Fig. 2. Experimental results. (a) The response of the first mirror to a single step, double step, and unipolar overdrives with different τ_{max} . (b) The response of the second mirror to single step and bipolar overdrives with different τ_{max} . (c) Voltages corresponding with the responses in Fig. (2a). d) Voltages corresponding to the 9x and 4x responses in Fig. (2b). The solid and dashed lines represent pad1 and pad2 voltages respectively. Pad2 voltages are shown as negative to more easily distinguish from the pad1 voltages and to better represent the torques acting on the mirror.

presented here, V_0 is 50 V and θ_0 is 5° optical. The single step response for the first mirror, Fig. 2(a), settles to within 5% of the final value at $t = 2.2T_0$. The results in Figs. 2(a) and 2(c) demonstrate not only the improvement in performance but also the simplicity of these drive techniques. The second mirror investigates the bipolar drive and similarly, the responses are shown in Fig. 2(b). We show the SS response, which settles at $t = 1.2T_0$, for comparison to the first mirror. The difference between the two SS responses is consistent with the difference in the Q factors. The driving voltages for BOD 4x and 9x are shown in Fig. 2(d), the rest of the drives are omitted to simplify the plot. Here the solid lines represent the voltages applied to pad1 and the dashed lines are the voltages to pad2. The pad2 voltages have been plotted as negative voltages to more easily distinguish the different curves and better represent the torques acting on the mirror. This plot demonstrates the improvement in speed over single step, with t_2 for BOD 9x at $0.11T_0$. One consequence of increasing the torque and speed of the pulses is that we start to excite a higher order mode of the mirror, shown in the BOD 9x response. Imboden *et al.* [22] showed that for the two step drive, adding a second step into the pulses can suppress such higher order modes. It is clear from Fig. 2 that both overdrive methods have a speed advantage over




Fig. 3. (a) Comparison of theoretical, simulated and experimental results for t_1 for different bipolar drive amplitudes. (b) Comparison of theoretical, simulated and experimental results for t_2 for different bipolar drive amplitudes. (c) Comparison of undamped and damped t_2 results at different bipolar amplitudes and Q values. (d) Simulated overshoot error vs timing error for different bipolar amplitudes using a Q factor of 100. (e) Simulated overshoot error vs timing error for different values use a bipolar with $\tau_{max}/\tau_0 = 10$.

the single step and double step. What is less obvious from Fig. 2 is the comparison between the unipolar and bipolar drives. If we continue to increase the τ_{max}/τ_0 the bipolar overdrive will continue to settle quicker; whereas, the analytic solution shows that for a high Q system the unipolar overdrive is limited to $0.25T_0$ as the braking torque is limited to that of the restoring torque of the spring [22].

Fig. 3 further investigates the bipolar drive results. Figs. 3(a) and 3(b) compare the bipolar drive t_1 and t_2 values of the experiments with theoretical and simulated values. Both plots show that when $V_{max}/V_0 = 1$, there is a significant difference between the theoretical and experimental timings; however, as V_{max}/V_0 increases, the theory and experimental results have better agreement. Similarly, Figs. 3(a) and 3(b) show that the Q has larger effect on t_1 than on t_2 , with the discrepancy between the experimental and theoretical t_1 greater than that of t_2 for all V_{max}/V_0 values. The disagreement between the data and theory is due to the high Q factor assumption discussed earlier. We see that the simulations, which account for the Q of the device, have good agreement with the experimental findings.

Using the simulation, we further investigate the effect of damping and the overshoot error due to errors in t_1 and t_2 for the bipolar overdrive. Fig. 3(c) plots the ratio of a damped versus undamped t_2 value for three different forcing conditions. This plot essentially quantifies the accuracy of an infinite Q assumption, but more importantly shows the effect of damping on the final settling time. At higher values of Q (≥ 100), the assumption is good with the simulated and theoretical timings agreeing. As the Q value decreases the timings diverge such that a critically damped system ($Q = 0.5$) will settle 34%, 13%, and 4% slower than an undamped system driven

by a bipolar overdrive with $\tau_{max}/\tau_0 = 2, 4, \text{ and } 10$ respectively. Fig. 3(d) investigates the effect of timing error, t_{error} , on the angular error, θ_{error} , for the same forcing conditions as Fig. 3(c). Here we define t_{error} as the deviation from the optimal t_1 and t_2 timing and θ_{error} as the normalized overshoot of the system ($\frac{\theta_{max}-\theta_0}{\theta_0}$). We use a Q factor of 100 for this study. As we expect, the larger the timing error, the larger the error in the response. Similarly, higher torques allow the mirror to move quicker, but at the expense of more error for a given timing error. For example, at $t_{error}/T_0 = 0.001$, the BOD 2x drive will have less than 2% error, whereas the BOD 10x drive has ~6% error. Fig. 3(e) explores the effect of the Q factor on θ_{error} . For this plot we use a BOD 10x drive. With more damping, the system is less sensitive to error in the timing; however, Fig. 3(c) shows that this also means a slower settling time.

Throughout this work we have normalized the settling time and displacements, therefore generalizing the results for other linear systems. One system in particular that is of interest is a micromirror capable of switching within 10 μs [11], and Figs. 3(b)-3(d) help build a roadmap for creating such a device. An electrostatic MEMS mirror with a $Q > 3$ driven by a bipolar overdrive with an τ_{max}/τ_0 of 10 ($V_{max}/V_0 = 3.16$) will settle within $0.17T_0$. Therefore, a micromirror with a natural frequency of 10 kHz can switch in under 10 μs . These kinds of frequencies are relatively common for a scanning electrostatic mirror [28, 29]. Figure 3(d) shows that in order to keep the error under 1% ($\theta_{error} = 0.01$) one needs timing precision of 15 ns ($0.00015T_0$). Again, this is well within the range of modern controllers.

3.2. PWM control

Switching time is an important parameter for these mirrors; however, as discussed previously, providing a cheaper method of control is equally important, and PWM is a more cost effective method. Each one of the analog drives in Table 1 can be implemented as a PWM drive by varying the duty cycle to control the torque. Fig. 4a and 4b illustrate how the analog drives for DS and UOD 1x drives are converted into a PWM signal. For the double step, instead of using a $\tau_0/2$ step, the PWM version uses a 50% duty cycle signal. The UOD 1x drive is more trivial because it is stepping up and down between 0% and 100% duty cycle. However, this same technique can be used to make smaller steps by adjusting the 100% duty cycle pulses with different duty cycles. Fig. 4(c) compares the response of the PWM double step and PWM 1x unipolar overdrive with their analog counter parts. The same analog response from Fig. 2(a) is included for comparison. Instead of controlling the electrostatic force by adjusting the voltage level, we adjust the duty cycle of a PWM signal. For example the drive for PWM DS, consists of a 68% duty cycle signal until $t = 0.51T_0$ at which point the duty cycle is increased to 100%. Both of these parameters agree with the theoretical value within 3%. It should be noted that the PWM method is implemented with no discernible performance difference compared to the analog method.

4. Conclusion

As data traffic continues to grow exponentially, engineers and scientists around the world work towards the next generation of all-optical transmission and switching technology. 3D MEMS beam steering is a mature, commercially viable option with some of the best optical properties to date. However, MEMS micromirrors are typically dismissed for their slow switching speed and high cost per port. In this work we have demonstrated that a modern, simple to use, open loop controls technique can improve the switching time of an existing commercial mirror by a factor of 10. These controls in conjunction with PWM provides a fast, cheap method to control electrostatic micromirrors. We further discussed how a mirror with a 10 kHz natural frequency could achieve a 10 μs switching time. We believe that an appropriately designed 3D MEMS switch with modern controls can be competitive for the optical switching technology of the future.




Fig. 4. (a) Illustration of how the double step drive is converted to a PWM signal. (b) Illustration of a UOD 1x drive driven by PWM. (c) Responses from the PWM control versions of the double step and 1x unipolar overdrive. The analog (dashed lines) are the same curves presented in Fig. (2a) for comparison.

Funding

National Science Foundation (NSF) (EEC-0812056);

References

1. T. J. Chainer, M. D. Schultz, P. R. Parida, and M. A. Gaynes, "Improving Data Center Energy Efficiency With Advanced Thermal Management," *IEEE Transactions on Components, Packag. Manuf. Technol.* **7**, 1228–1239 (2017).
2. Q. Cheng, M. Bahadori, M. Glick, S. Rumley, and K. Bergman, "Recent advances in optical technologies for data centers: a review," *Optica* **5**, 1354 (2018).
3. P. J. Winzer, D. T. Neilson, and A. R. Chraplyvy, "Fiber-optic transmission and networking: the previous 20 and the next 20 years [Invited]," *Opt. Express* **26**, 24190–24239 (2018).
4. Q. Cheng, S. Rumley, M. Bahadori, and K. Bergman, "Photonic switching in high performance datacenters [Invited]," *Opt. Express* **26**, 16022 (2018).
5. D. Kilper, K. Guan, K. Hinton, and R. Ayre, "Energy challenges in current and future optical transmission networks," in *Proceedings of the IEEE*, vol. 100 (2012), pp. 1168–1187.
6. V. Alwayn, *Optical Network Design and Implementation* (Cisco Press, 2004).
7. C. Kachris and I. Tomkos, "A survey on optical interconnects for data centers," *IEEE Commun. Surv. Tutorials* **14**, 1021–1036 (2012).

8. A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano, A. Kanagala, J. Provost, J. Simmons, E. Tanda, J. Wanderer, U. Hölzle, S. Stuart, and A. Vahdat, "Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network," *Proc. 2015 ACM Conf. on Special Interest Group on Data Commun.* pp. 183–197 (2015).
9. C. Lam, H. Liu, B. Koley, X. Zhao, V. Kamalov, and V. Gill, "Fiber optic communication technologies: What's needed for datacenter network operations," *IEEE Commun. Mag.* **48**, 32–39 (2010).
10. G. Lee, *Cloud Data Center Networking Topologies* (2014).
11. DARPA, "HR001119S0004," (2018).
12. T. Mori, S. Naito, R. Kawahara, and S. Goto, "On the characteristics of Internet traffic variability: Spikes and elephants," in *2004 International Symposium on Applications and the Internet*, (IEEE, 2004), pp. 99–106.
13. L. Guo and I. Matta, "The war between mice and elephants," in *Proceedings Ninth International Conference on Network Protocols*, (2001), pp. 180–188.
14. V. A. Aksyuk, F. Pardo, D. Carr, D. Greywall, H. B. Chan, M. E. Simon, A. Gasparyan, H. Shea, V. Lifton, C. Bolle, S. Arney, R. Frahm, M. Paczkowski, M. Haueis, R. Ryf, D. T. Neilson, S. Member, J. Kim, C. R. Giles, and D. Bishop, "Beam-Steering Micromirrors for Large Optical," *J. Light. Technol.* **21**, 634–642 (2003).
15. V. A. Aksyuk, S. Arney, N. R. Basavanahally, D. J. Bishop, C. A. Bolle, C. C. Chang, R. Frahm, A. Gasparyan, J. V. Gates, R. George, C. R. Giles, J. Kim, P. R. Kolodner, T. M. Lee, D. T. Neilson, C. Nijander, C. J. Nuzman, M. Paczkowski, A. R. Papazian, F. Pardo, D. A. Ramsey, R. Ryf, R. E. Scotti, H. Shea, and M. E. Simon, "238 Å Å 238 Micromechanical Optical Cross Connect," *IEEE Photonics Technol. Lett.* **15**, 587–589 (2003).
16. M. Stepanovsky, "A Comparative Review of MEMS-based Optical Cross-Connects for All-Optical Networks from the Past to the Present Day," *IEEE Commun. Surv. & Tutorials* **PP**, 1–1 (2019).
17. J. Kim, C. J. Nuzman, B. Kumar, D. F. Lieuwen, J. S. Kraus, A. Weiss, C. Lichtenwalner, A. R. Papazian, R. E. Frahm, N. R. Basavanahally, D. A. Ramsey, V. A. Aksyuk, F. Pardo, M. E. Simon, V. Lifton, H. B. Chan, M. Haueis, A. Gasparyan, H. R. Shea, S. Arney, C. A. Bolle, P. R. Kolodner, R. Ryf, D. Neilson, and J. V. Gates, "1100 x 1100 port MEMS-based optical crossconnect with 4-dB maximum loss," *IEEE Phot. Tech. Lett.* **15**, 1537–1539 (2003).
18. W. Mellette and J. E. Ford, "Scaling limits of MEMS beam-steering switches for data center networks," *J. Light. Technol.* **33**, 3308–3318 (2015).
19. P. De Dobbelaere, K. Falta, L. Fan, S. Gloeckner, and S. Patra, "Digital MEMS for optical switching," *IEEE Commun. Mag.* **40**, 88–95 (2002).
20. X. Ma and G.-S. Kuo, "Optical switching technology comparison: Optical MEMS vs other technologies," *IEEE Opt. Commun.* **41**, S16–S23 (2003).
21. Calient, "S320 PHOTONIC SWITCH," (2013).
22. M. Imboden, J. Chang, C. Pollock, E. Lowell, M. Akbulut, J. Morrison, T. Stark, T. G. Bifano, and D. J. Bishop, "High-speed control of electromechanical transduction," *IEEE Control. Syst. Mag.* **36**, 48–76 (2016).
23. D. T. Neilson, V. A. Aksyuk, S. Arney, N. R. Basavanahally, K. S. Bhalla, D. J. Bishop, T. W. V. Blarcum, L. Zhang, C. R. Giles, B. Labs, L. Technologies, M. Avenue, and M. Hill, "Fully Provisioned 112x112 Micro-Mechanical Optical Crossconnect With 35.8 Tb/s Demonstrated Capacity," in *Optical Fiber Communication Conference*, (2000), pp. 202–204.
24. C. Pollock, L. K. Barrett, P. G. del Corro, A. Stange, T. G. Bifano, and D. J. Bishop, "PWM as a Low Cost Method for the Analog Control of MEMS Devices," *J. Microelectromechanical Syst.* **28**, 1–9 (2019).
25. C. Pollock, M. Imboden, A. Stange, J. Javor, K. Mahapatra, L. Chiles, and D. J. Bishop, "Engineered PWM drives for achieving rapid step and settle times for MEMS actuation," *J. Microelectromech. Syst.* **27**, 513–520 (2018).
26. F. L. Lewis, *Optimal Control* (Wiley, 1986), 1st ed.
27. D. N. Burghes and A. Graham, *Introduction to control theory, including optimal control* (Horwood Ltd., 1980), 1st ed.
28. S. T. S. Holmstrom, U. Baran, and H. Urey, "MEMS laser scanners: a review," *J. Microelectromech. Syst.* **23**, 259–275 (2014).
29. J. B. Hopkins, R. M. Panas, Y. Song, and C. D. White, "A High-Speed Large-Range Tip-Tilt-Piston Micromirror Array," *J. Microelectromechanical Syst.* **26**, 1–10 (2017).