

Scaling Limits of MEMS Beam-Steering Switches for Data Center Networks

William Maxwell Mellette and Joseph E. Ford, *Fellow, OSA*

Abstract—Transparent optical circuit switching can improve the aggregate bandwidth, scalability, and cost of data center networks, provided it can meet the performance requirements on switching speed, port count, and optical efficiency. Here, we examine the theoretical scaling limits of transparent nonblocking optical switches based on MEMS electrostatic tilt mirror devices. Using physical optics and electromechanics, we present a first principles analysis of how the response speeds of a set of canonical devices scale as a function of switch port count, crosstalk, and insertion loss. Our model indicates that the optimal actuator design (parallel plate versus vertically offset comb) and actuation method (digital versus analog) changes as a function of switch port count. It also suggests that conventional switch topologies do not allow a favorable tradeoff between switching speed and optical efficiency or crosstalk. However, high switching speeds can be achieved by multistage switch architectures such as the two examples we describe, a multiport wavelength switch and a wavelength-independent space switch.

Index Terms—Microelectromechanical systems, optical communication systems, optical design, optical fiber switches.

NOMENCLATURE

α	Mirror array fill factor (see Section II-D).
A	Area of electrode overlap (see Fig. 4).
β	Half divergence angle of optical beam (see Fig. 5).
η_{ac}	Angular confinement efficiency (see Section II-D).
η_f	Fiber coupling efficiency (see Section II-D).
η_{sc}	Spatial confinement efficiency (see Eq. (15)).
η_{switch}	Overall optical efficiency of switch (see Eq. (16)).
ε_0	Permittivity of free space (see Section II-B).
E	Young's modulus (see Eqs. (2), (12), and (13)).
f_0	Natural resonant frequency (Eq. (9)).
g	Finger to finger air gap (comb actuator) (see Fig. 4).
G	Shear modulus (see Eq. (11)).
h_m	Electrode-mirror air gap (plate actuator) (see Fig. 4).
i	Tilt axis, spanning x and y .
I	Moment of inertia (see Section II-C).
k	Rotational spring constant (see Eqs. (11) and (12)).
λ	Wavelength.
l_f	Comb finger length (see Fig. 4).
l_s	Spring length (see Section II-C).
M	Number of differentiable optical mirror states (see Eq. (14)).

Manuscript received December 23, 2014; revised April 7, 2015; accepted May 4, 2015. Date of publication May 7, 2015; date of current version June 20, 2015. This work was supported by the National Science Foundation under Grant #1314921.

The authors are with the University of California San Diego, La Jolla, CA 92093-0407 USA (e-mail: wmellett@ucsd.edu; jeford@ucsd.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JLT.2015.2431231

N	Number of ports in switch.
N_f	Number of comb fingers (see Eq. (4)).
ψ	Electrode ramp angle (plate actuator) (see Fig. 4).
r_m	Mirror radius.
R	Mirror reflectivity (see Section II-D).
S	Number of active switching stages (see Eqs. (1) and (2)).
$\theta, \theta_{\text{max}}$	Mechanical tilt angle, maximum tilt angle.
t_f	Comb finger thickness (see Fig. 4).
t_m	Mirror thickness (see Fig. 4).
t_s	Spring thickness (see Eq. (11)).
τ	Torque (see Section II-B).
V	Applied electrode voltage (see Section II-B).
w_0	Waist of Gaussian beam (see Section II-D).
w_s	Spring width (see Eq. (11)).
z_R	Rayleigh range of Gaussian beam.

I. INTRODUCTION

HIGH-SPEED optical circuit switches have recently been proposed for data center networks to provide reconfigurable high-bandwidth and data-rate agnostic channels without cascaded optical transceivers [1]. Optical switches are typically based on one of two prevalent technologies: planar guided-wave devices or three-dimensional (3-D) free-space devices. Guided-wave devices promise high speed and dense integration, but switching fabrics with acceptable optical performance have been limited to 16×16 ports [2]. Larger guided-wave switches have been demonstrated [3], but have higher insertion loss. Among free-space technologies, optical microelectromechanical systems (MEMS) based switches have an extensive publication record [4], [5], and have proven successful in telecommunications networks for bandwidth provisioning and fault protection, which requires large port counts and low loss, but relatively slow switching speeds of 10–100 ms. MEMS devices are not intrinsically slow; electrostatically actuated MEMS structures can have GHz resonant frequencies [6], but the optical requirements on beam-steering MEMS devices limits their response speed. Digital MEMS tilt mirrors are the fastest optical beam-steering devices, switching in 20 microseconds or less [7]. However, bistable operation has limited their use to small port-count switching. In data centers, both switching speed and port count are critical figures of merit, and sub-millisecond response times are essential to meet the network demands [8].

Here we re-examine canonical MEMS tilt mirror devices to quantify the tradeoffs between switching speed, port count, and optical performance. From a network-level perspective, our results can also be interpreted in terms of the number of reconfigurable ports achievable per second, by taking the product of




Fig. 1. (a) Schematic of a 1×8 switch constructed from three stages of 1×2 mirrors. (b) A 3×3 switch using two stages of 1×3 mirrors. The latter exhibits the topology of a conventional $N \times N$ 3-D-MEMS OXC.

the device resonant frequency and switch port count. A basic $1 \times N$ MEMS switch directs light from a single input fiber through free space to an electrostatically actuated mirror, which redirects the light to couple to one of N output fibers. $N \times N$ switches, with N inputs and N outputs, can be thought of as a collection of $1 \times M$ switches which use free space and relay optics to refocus light between a series of mirrors. The most scalable switching structure to date has been the MEMS-based optical cross connect (OXC), which uses two two-axis beam-steering micromirror arrays to interconnect a matrix of up to 1100 input and output fibers [9]. While the specific switch layouts can differ, we can still compare MEMS device performance based on the fundamental requirement that each micromirror discriminates between optical switch states.

Using fundamental physical mechanics, electrostatics, and free-space optics, we investigate how the response speeds of canonical one- and two-axis tilt mirror devices scale as a function of switch port count, crosstalk, and insertion loss. The electrostatic, mechanical, and optical properties of the MEMS devices as well as switch topologies are discussed in Section II. In Section III, we describe the numerical approach used to quantify device performance, analyze the results by considering specific design cases in more detail, and compare the modeling results to a commercial switch. The findings motivate us to explore new overall optical switching configurations, which are discussed in Section IV.

II. GENERALIZED MEMS BEAM-STEERING SWITCH MODEL

Beam-steering switches fall into two major categories, those which incorporate wavelength selectivity using spectral demultiplexing and $1 \times N$ port topologies, or those that use wavelength-independent $N \times N$ port topologies (see Fig. 1). A typical tilting micromirror device consists of a flat region to reflect a beam of light, a supporting structure to suspend the mirror and provide angular restoring force, and a set of nearby electrodes which apply electrostatic force to tilt the mirror. To explore a wide variety of switch configurations, we consider a switch as consisting of four modular parts: 1) an overall switch architecture, 2) a MEMS actuation structure, 3) a mirror structure, and 4) a set of resolvable optical beam paths which meet insertion loss and crosstalk requirements.

Our analysis technique needed to be general enough to cover the scope of MEMS tilt mirror actuators without sacrificing accuracy by excessive simplification. The theoretical basis of

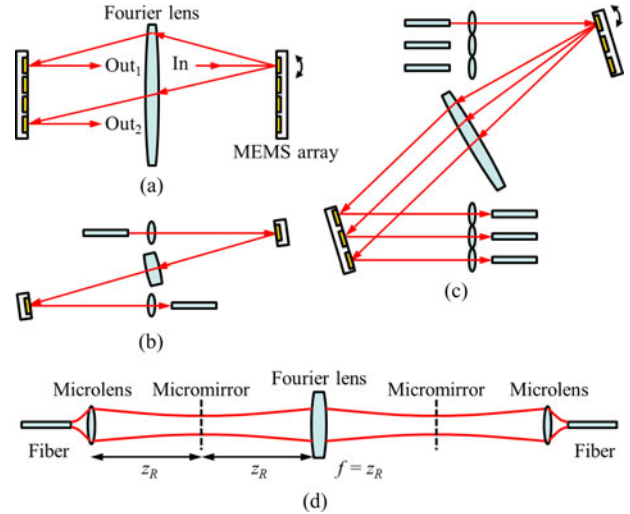


Fig. 2. (a) Top view of an $N \times N$ OXC using a Fourier lens to make full use of the tilt range of all micromirrors in the MEMS array. Side views show separation of beam paths in the switch using (b) one-axis and (c) two-axis micromirrors. (d) Gaussian beam profile through the unfolded system, showing relaying of the beam waist between micromirror planes when the focal length of the Fourier lens and distance between MEMS array and Fourier lens equal the Rayleigh range of the beam.

our numerical Matlab model used a straight-field approximation, Euler–Bernoulli beam theory, Hooke’s law, and Gaussian beam optics. This model is less accurate than a device-specific finite element simulation, but provides orders of magnitude faster execution. This allowed a search for optimal device designs over a large parameter space and the observation of scaling behavior over a large range of switch structures.

A. Overall Switch Architecture

The port count of a cross connect is determined by the switch architecture. $1 \times N$ switches typically use a tree topology with one or more switching stages, where one input node branches sequentially into N output nodes (or vice versa), with each micromirror acting as a branching node in the tree (see Fig. 1(a)). The number of stages in the tree, S , is related to the number of optical mirror states of each mirror, M , and the number of output ports, N , by

$$S = \log_M(N). \quad (1)$$

$N \times N$ switches typically use a folded multi-rooted tree topology with at least two switching stages (see Fig. 1(b)). In this topology,

$$S = 2\log_M(N). \quad (2)$$

Conventional free-space OXCs use $S = 2$ stages of N -state mirror elements ($M = N$), allowing N input and N output ports. Fig. 2 illustrates $N \times N$ OXC geometries using 1- and 2-axis micromirrors. 1-axis switches use a linear array of mirrors and 2-axis switches use a two-dimensional array of mirrors. Introducing passive optics to aim the beam paths toward the center of the second array makes full use of the micromirrors, scan range, increasing the port count of a 2-axis switch by $4 \times$ com-

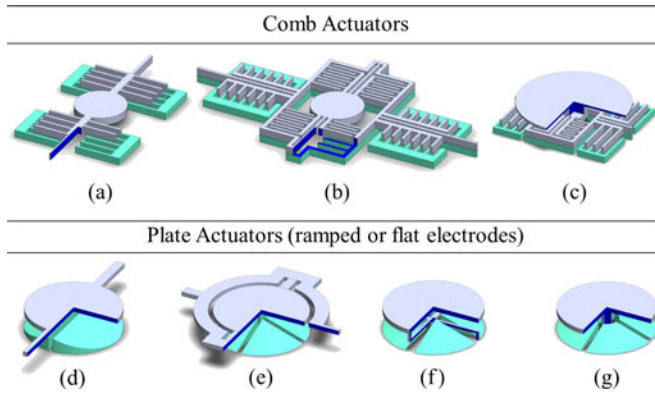


Fig. 3. Conceptual illustrations of MEMS tilt mirror actuators considered in the study: (a) one-axis comb, (b) two-axis in-plane comb, (c) two-axis hidden comb, (d) one-axis plate, (e) two-axis plate with gimbal, (f) two-axis plate with hidden “crossbar” springs, and (g) two-axis plate with hidden “post” spring. Plate actuators are shown with ramped electrodes, but were also analyzed with flat electrodes. Partial cross sections have been taken to reveal the structure.

pared to designs which do not incorporate passive beam aiming. There are a number of nearly equivalent techniques to aim the beams with passive optics, including field lenses at the collimator arrays, field lenses at the micromirror arrays, or a Fourier lens between micromirror arrays. Here we focus on the Fourier lens switch geometry for subsequent modeling and analysis [9]. Choosing the focal length of the Fourier lens to equal the Rayleigh range, z_R , of the optical beam and placing the lens one focal length from each MEMS array relays the beam waist between micromirrors. This reduces the aperture requirement on the mirrors, which lowers inertia and increases switching speed. Prior work has assessed the scaling of beam-steering cross connects which do not employ a Fourier lens or other means of passive beam aiming [10].

B. Tilt Mirror Electrostatic Actuation

Choosing from the large number of actuators which have been proposed in the literature, we analyzed a set of commonly employed one- and two-axis torsional actuators using gap-closing plates and vertically offset combs [11] (see Fig. 3). Gap closing actuators are typically fabricated with parallel (flat) plate electrodes [12], but we also considered a ramped electrode design, which has been shown to have improved voltage response [13]. In addition to in-plane gimbaled two-axis tilt mirrors, we considered variations with the support structures hidden under the mirror [14], [15]. Although more difficult to fabricate, designs with hidden springs reduce the rotational inertial and increase the density of mirrors in the array. There are alternative hidden actuator designs with different design constraints (e.g., [16]) which are not considered here.

The maximum optical beam-steering angle is determined by the mechanical tilt range of the mirror, which we found by balancing the restoring torque of the supporting springs with the electrostatic torque applied by the electrodes. The mechanical restoring torque is approximated by Hooke’s law,

$$\tau_i = k_i \theta_i, \quad (3)$$




Fig. 4. Cross sectional illustrations of (a) comb and (b) plate actuation mechanisms showing design variables.

where k_i is the rotational spring constant and θ_i is the mechanical tilt angle about the i th axis, where i spans x and y .

The driving torque of the comb actuator, calculated by differentiating the stored energy in the effective capacitor, is

$$\tau_{\text{comb},i} = N_{f,i} \frac{\epsilon_0 V_i^2}{2g} \frac{dA_i}{d\theta_i}, \quad (4)$$

where N_f is the number of comb fingers, ϵ_0 is the vacuum permittivity, V is the applied voltage, g is the air gap between comb finger electrodes, A is the area of electrode overlap, and θ is the tilt angle, all for the i th axis [17]. As shown in Fig. 4(a), the electrode overlap area A depends on θ , the finger thicknesses, finger length, and fixed finger offset. For the in-plane comb actuator, we assume the mirror, comb arm, comb fingers, and torsion spring are fabricated from the same device layer in a single etch step, and must therefore have the same thickness, t_m . The fixed finger thickness, t_f , is defined by a separate device layer, and can have a different thickness. We found that thicker fixed fingers ($t_f > t_m$) increased the performance of the one-axis comb actuator by allowing larger tilt angles, because large tilt angles are necessary to achieve large port counts in one-axis actuators. Two-axis actuators have an increased dimensionality of tilt, and do not require such large tilt angles along each axis. We found that when optimizing the two-axis comb actuators for speed, the optimal devices always tilted slightly less than the thickness of the moving comb finger. This can be explained because the rate of change in capacitance with angle begins to diminish when the top of the moving finger tilts below the top of the fixed finger, reducing the applied torque past this point. Operation in this regime allows larger tilt, but requires softer torsion springs (and lower resonant frequency) for the same electrode voltage. The tradeoff between spring stiffness and tilt angle favored stiffer springs for the two-axis devices, and did not require the fixed fingers to be thicker than the moving fingers, at least for the port counts considered here. Note that optimized

comb actuators still had larger scan angles than plate actuators, fulfilling the expected design advantages of comb drives.

Small structural asymmetries can excite lateral failure modes of the comb drive, imposing additional limits on the maximum tilt angle [18]. We consider these effects as inherent to comb drives and include them in our model. We set the gap between comb fingers, g , and the comb finger width, w_f , to be two micrometers to maximize comb density [18].

The torque generated by the plate actuator was calculated by integrating the forces exerted on the mirror by the electrode, neglecting fringing fields. For one-axis tilt, we integrated over the electrode in a radial direction, giving an applied torque of

$$\tau_{\text{plate},x} = \int_0^{r_m} x dF, \quad (5)$$

where r_m is the mirror radius, and x is the direction normal to the rotation axis. The incremental applied force dF is

$$dF = \frac{\varepsilon_0 V^2 \sqrt{r_m^2 - x^2} dx}{(h_m - r_m \psi - x(\theta - \psi))^2}, \quad (6)$$

where h_m is the nominal air gap between the mirror and a flat electrode and ψ is the electrode ramp angle, defined in Fig. 4(b) [19]. Note that $\psi = 0$ for a flat electrode. In the two-axis plate actuator, the ramped electrode is conical in shape and we integrate in two dimensions, giving an applied torque of

$$\tau_{\text{plate},i} = \int_{\rho_1}^{r_m} \int_{\varphi_1}^{\varphi_2} \rho \cos \varphi dF, \quad (7)$$

where ρ is the radial coordinate and φ is the azimuthal coordinate of a cylindrical coordinate system and

$$dF = \frac{\varepsilon_0 V^2}{(h_m - r_m \psi - \rho(\theta \cos \varphi - \psi))^2} \rho d\varphi d\rho. \quad (8)$$

The mechanical tilt range of plate actuators considered here was less than $\pm 13^\circ$, so (6) and (8) remain reasonably valid.

We assumed the use of four quadrant (90°) electrodes as in Fig. 4(b), such that when tilting the mirror in a direction centered on a quadrant (45° from a quadrant boundary line), higher torque is applied by activating three electrodes rather than a single electrode. The gimbal design allows $\rho_1 = 0$ in (7), but for the post design ρ_1 must be greater than the post radius. In the hidden crossbar design, the width of the springs cuts into the area of the electrodes, and we must modify the integration limits in (7) accordingly. We found that maximal electrode ramp angles produced the highest performing plate actuator devices, except in the case of the hidden post spring, where the center cut-out in the electrodes to allow for the post negated the benefit of the ramped electrodes. Complex electrode designs can improve device performance [20], [21], but to maintain the large scope of our study, we focused on the most common designs, illustrated in Fig. 4.

We limited the applied electrode voltage to 275 V to avoid electrostatic breakdown [22]. Care was taken in calculating the maximum tilt angle for the plate actuator; past some tilt angle (typically 44% of the maximal angle), the nonlinearity in the torque exerted by the electrode overcomes the linear restoring

torque of the spring and the mirror is snapped down to the substrate. This is the well-known ‘‘pull-in’’ phenomenon [12].

The plate actuator can be purposely operated in the pull-in regime, allowing the mirror to be snapped to a discrete number of mechanical states in a ‘‘digital’’ fashion [7], where the mirror structure accelerates until it reaches contact with a mechanical stop. This mode of operation allows switching on microsecond time scales, but the small number of mechanical and optical states limits the port count of the switch. Alternatively, the mirror can be operated with continuous ‘‘analog’’ positioning over a smaller angular range, allowing more optical states but with a slower reconfiguration rate. Digital versus analog actuation is a critical switch design choice.

C. Tilt Mirror Dynamics

The maximum device switching speed is primarily limited by the resonant frequency at which the mirror structure oscillates. While driving the mirror faster than its natural resonant frequency is possible, this requires sophisticated high-voltage closed-loop control which is likely to be impractical to implement at high switching speeds due to the necessarily high device driver currents. The natural resonant frequency, f_0 , of the device is given as

$$f_{0,i} = \frac{1}{2\pi} \sqrt{\frac{k_i}{I_i}}, \quad (9)$$

where I_i is the moment of inertia about the i th axis. Mass located farther from the axis of rotation has a larger contribution to the rotational inertia. There can be multiple resonances of the structure [23], and the actuator design must ensure the desired torsional mode has the lowest resonant frequency in order to suppress unwanted motion in parasitic modes. The operational resonant frequency is proportional to the natural resonant frequency, but depends on both damping, Γ , and driving torque, τ_{drive} , and can be found by solving the full equation of motion given by

$$I \frac{d^2 \theta}{dt^2} + \Gamma \frac{d\theta}{dt} + k\theta = \tau_{\text{drive}}(\theta). \quad (10)$$

The transient solution can be found by detailed calculation of the damping term [24]. The damping can be fine-tuned by changing the ambient gas pressure or shape of the cavity beneath the mirror, or by etching small holes in the mirror [25]. For our analysis, the driven resonant frequency in the absence of damping is a sufficient metric for comparing the response speeds of different devices because damping effects establish a proportionality between driven resonant frequency and response time, and because that proportionality factor is tunable, it can be made similar in all devices considered.

Comb actuators have a nearly linear response because the driving torque is nearly constant as a function of θ (up to the angle at which the comb fingers are fully interdigitated). In the absence of damping, then, the driven resonant frequency of the comb actuator is its natural frequency. The driving torque of the plate actuator, on the other hand, is highly nonlinear in θ . The driven resonant frequency drops as a function of tilt angle, approaching zero at the pull-in angle. Closed-loop control can

extend the analog tilt range of the mirror [26], but providing sufficient voltage and current for closed-loop control becomes extremely challenging with fast switching devices, so we assumed open-loop control. This simplifies drive electronics and allows a direct comparison to the comb actuator, which does not exhibit the same vertical pull-in effect. This means some angular margin must be maintained between the maximum operational tilt angle and the pull-in angle in order to drive the plate actuator in an analog fashion at high speeds.

The physical geometry of the suspension structures determines their stiffness. The rotational spring constant of the torsional elements is given by

$$k_{\text{torsion}} = \frac{2G}{l_s} ab^3 \left(\frac{1}{3} - 0.21 \frac{b}{a} \left(1 - \frac{b^4}{12a^4} \right) \right), \quad (11)$$

where G is the shear modulus of the material (polysilicon), l_s is the length of the spring, and a and b are the longer and shorter dimensions of the beam cross section, respectively. Depending on the design, either a or b can assume the spring width, w_s , or the spring thickness, t_s . For in-plane devices, the spring thickness was set equal to the mirror thickness so both structures could be fabricated from the same device layer in a single etch step. The hidden actuator devices decouple mirror thickness from spring thickness. We approximated the flexure structure in the hidden post design as having a rotational spring constant given by

$$k_{\text{flexure}} = \frac{\pi E r_{\text{post}}^4}{l_s}, \quad (12)$$

where E is Young's modulus of the material (polysilicon), r_{post} is the post radius, and l_s is the spring length. Nonlinear springs are commonly used in MEMS structures and have been shown to extend the tilt range of micromirrors [27]. However, to maintain the scope of the study, we used linear springs in our model because they do not require case-by-case optimization.

In practice, the finite translational stiffness of the springs means that an applied electrostatic force will contribute to moving the mirror vertically (in a piston mode), and will slightly reduce the torsional deflection. We found the torsional spring constant was at least an order of magnitude weaker than the flexure spring constant for the high-aspect ratio springs considered here, so we approximated that all applied force contributed to the torsional mode. Case-by-case spring optimization could further suppress the piston mode.

Because of its finite stiffness, the micromirror bends under static and dynamic actuation. We constrained the mirror to maintain a flatness of 1/8th the wavelength to satisfy the Rayleigh criterion. Because the plate actuator applies force directly onto the mirror, we required the mirror thickness to increase with mirror radius and spring constant to maintain flatness under static deflection:

$$t_m = \left(\frac{16r_m^k \tan\theta}{E \lambda} \right)^{1/3}. \quad (13)$$

The comb actuated mirror does not experience direct electrostatic force, and can typically be thinner for the same radius. In

this case, the limiting thickness is determined by the dynamic deformation of the mirror [28].

D. Tilt Mirror Optical Response

The micromirror must be able to discriminate between a discrete number of optical switch states without excessive optical loss or crosstalk. For the tilt angles used in MEMS beam-steering switches, the number of optical states resolved along each rotational axis can be approximated as

$$M_i = \frac{2\theta_{\text{max},i}}{\beta} + 1, \quad (14)$$

where $\pm\theta_{\text{max},i}$ is the maximum mechanical tilt angle along the i th axis, and β is the half divergence angle of the optical beam.

In $N \times N$ switches, as shown in Fig. 2, the beam propagates to a second array of mirrors. The physical size of the spring structures, gimbal, and comb fingers surrounding the mirror all contribute to the footprint of a single device, and limit how close adjacent mirrors can be positioned in the array. We assumed plate actuators to be separated by at least twice the mirror height, h_m , to prevent electrical crosstalk. We define the linear mirror fill factor, α_i , as the ratio of the mirror diameter to the mirror pitch along the i th dimension. This value changes with the physical structure of the actuator, and was calculated on a case-by-case basis for each actuator design. Mathematically, α_i scales the first term on the right-hand side of (14), so that lower fill factors reduce the number of addressable optical states. We did not include the potential reduction in fill factor from electrical routing because it is highly design dependent. Multilayer electrical routing can increase device density significantly compared to planar routing [29]. We did not include the skew angle of the MEMS array or path length variability in our optimization model. The impacts on tilt angle and insertion loss are a second-order correction to the model, and these impacts are quantified for example design cases in Section III-B using physical optics modeling in Zemax. We also note that skew can be completely removed by choice of switch geometry, while still using a Fourier lens configuration [30].

The number of optical states for devices with two rotation axes can be approximated by determining the number of states along each principle axis of rotation. The two-axis comb actuator has two independent axes, each with its own torsion springs and comb fingers. The maximum tilt angle of the comb actuator traces a rectangle in angular space (for small angles) and the device can resolve $M_{\text{comb}} = M_x M_y$ beams, where M_x and M_y are the number of 1-D resolvable beams along each axis, given by (14). The two axes of the plate actuator have independent torsion springs, but are coupled by a common electrode. The maximum tilt angle of the plate actuator traces an ellipse in angular space, resolving $M_{\text{plate}} = \pi M_x M_y / 4$ beams. The distributions of mass and spring constants differ between the two axes, so in general both the resonant frequencies and maximum tilt angles differ for each axis (i.e. $f_{0,x} \neq f_{0,y}$ and $M_x \neq M_y$). Because the maximum response rate of the device is limited by its slowest axis, devices optimized for speed tend to have comparable resonant frequencies (and different angular ranges)

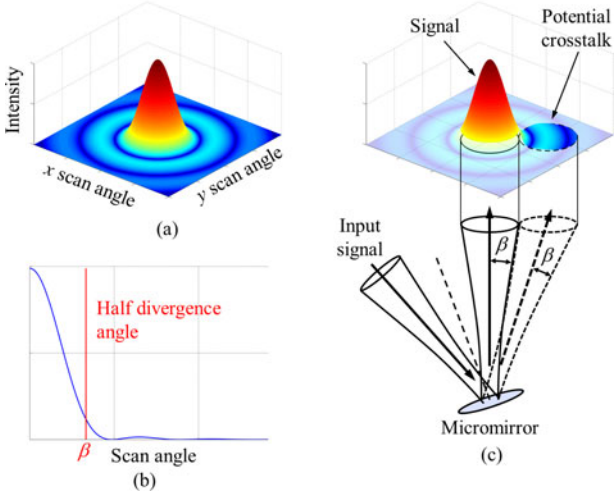


Fig. 5. (a) Far field angular intensity diffracted from a micromirror. (b) Cross section of diffraction pattern and choice of half divergence angle, β . (c) An input beam diffracts light into signal and crosstalk beams, whose distinguishability is defined in angular space by β .

along both axes. This leads to an inequality in the number of resolvable optical states between axes, which skews the shape of the mirror array to rectangular instead of square.

We modeled the light emitted by the input single mode fiber as a Gaussian beam parameterized by a nominal wavelength, λ , of 1550 nm. The Gaussian approximation of a fiber mode is less accurate far from the optical axis, and more detailed analysis may be necessary for systems where very high extinction ratios are required [31]. Using the Fourier lens OXC geometry (see Fig. 2), we place the waist of the beam at the micromirror. The Gaussian beam is infinite in spatial extent and is clipped at each mirror, resulting in a spatial confinement efficiency, η_{sc} , at the mirror given by

$$\eta_{sc} = 1 - \exp\left(-\frac{2r_m^2}{w_0^2}\right), \quad (15)$$

where w_0 is the beam waist. The angular distribution of light reflected from the mirror is altered as a result of diffraction from the edges of the mirror, and is no longer a pure Gaussian beam. We calculated the far-field intensity distribution of light reflected from a mirror by convolving the Fourier transforms of the Gaussian field and the mirror aperture. Using the far-field diffraction pattern (see Fig. 5(a)), we defined a nominal angular subtense, β , to distinguish the signal portion of the beam from the surrounding potential crosstalk (see Fig. 5(b) and (c)). The fraction of power encircled within the signal portion defines the angular confinement efficiency, η_{ac} , while the nearest neighbor crosstalk is found by integrating the appropriate region of the surrounding power. Thus, the nominal divergence angle of the beam is related to the mirror radius, beam waist, angular confinement, and crosstalk.

Our approximation of crosstalk using encircled energy was necessary to limit computation time during optimization. This method gives an upper bound on the crosstalk, and becomes more accurate in the limit of high insertion loss, which is where

crosstalk becomes a significant concern. A more accurate assessment of crosstalk requires a mode overlap calculation between the fiber and the potential crosstalk signal after it is focused by the corresponding microlens. We perform this detailed analysis for design examples in Section III-B.

For a given confinement efficiency, the overall insertion loss of the switch is driven by the number of stages. We modeled the overall switch throughput efficiency, η_{switch} , as a series of lumped element efficiencies at each micromirror:

$$\eta_{\text{switch}} = \eta_f (R\eta_{sc}\eta_{ac})^S, \quad (16)$$

where η_f is the fiber coupling efficiency at the output (assumed to be 90% based on experimental demonstrations in large port count OXCs [9]), R is the mirror reflectivity (assumed to be 97% for gold at 1550 nm), and S is the number of micromirror stages in the switch (see (1) and (2)).

III. SWITCH SCALING STUDY

A. Optimization and Theoretical Scaling Limits

We implemented the model discussed in Section II numerically in MATLAB. The initial goal was to determine how the resonant frequency of each MEMS device scales with switch port count and optical performance levels in a conventional OXC. The geometrical form of the device, including the mirror radius and thickness, spring width and length, electrode shape and air gap, as well as the optical beam parameters all constitute a design space which determines the resonant frequency and optical performance of a device.

For a given device, switch port count, insertion loss, and crosstalk, the problem of determining the optimal values of all free design variables is underdetermined. Consequently, we implemented a global search over the design space, with the geometrical form and optical beam parameters as inputs to the algorithm. Although computationally slower than other optimization methods, such a brute-force search is immune to local maxima and does not require assumptions about the optimization space other than its value limits. We bounded the search algorithm on the bottom end by assuming a minimum feature size of 1 micrometer and a minimum beam waist of $3 \mu\text{m}$ (twice the wavelength at 1550 nm). The upper end was bounded by the size of the mirror, and corresponding mechanical structures, necessary to achieve the maximum port count we considered ($N = 4096$ ports). We calculated that mirrors larger than 4 mm in diameter had excessive optical performance to meet the maximum port count, and would be unnecessarily slow due to increased inertia. We checked the solutions to ensure that the imposed boundaries did not arbitrarily constrain the design space. We discretized each design variable linearly or logarithmically with sufficiently fine sampling that we saw convergence in the solution.

The electrostatics, mechanics, and optics coupled many of the design variables. To save computation time, we separated the algorithm into an electromechanics component and an optics component which were coupled through the mirror radius. In the electromechanics code, we used the geometrical form parameters of the mirror, springs, and electrodes to compute the resonant frequency using (9), (11), and (12) and the maximum

mechanical tilt angle using (2)–(8) for every realization of each device within the design space. Dual axis devices took into account that the resonant frequency of the outer axis depends on the parameters of the inner axis, and that different spring constants are needed to achieve the same resonant frequency along both axes. The number of unique electromechanical realizations of a single device ranged from 10^5 to 10^7 , depending on the number of design variables. The optical portion of the code used the mirror radius, beam waist, and beam divergence angle as inputs to compute the spatial confinement, angular confinement, and crosstalk. We considered roughly 10^5 unique optical configurations.

For a specified switch topology and number of ports, we used (16) to calculate the insertion loss of the switch for each optical realization. We then eliminated any optical realizations which did not satisfy specified levels of insertion loss and crosstalk. Next, we used (14) to eliminate a portion of the electromechanical realizations based on the number of resolvable optical states required. Finally, we sorted the remaining device realizations by driven resonant frequency to determine the fastest device capable of meeting the specified switch parameters. This process was repeated for different devices, port counts, and optical performance parameters.

Fig. 6 shows how each device's resonant frequency scales with switch port count for a conventional $N \times N$ 3-D-MEMS OXC (see Fig. 2). The optical performance was constrained to have 3 dB insertion loss using (16) and less than -20 dB crosstalk. We found that insertion loss imposed the stronger constraint, and that all optimal designs had approximately 3 dB loss and much less than -20 dB crosstalk. There is a clear tradeoff between switching speed and port count, which can be understood through two functional relationships. First, the tilt range of a mirror is inversely proportional to resonant frequency through the spring constant in (9) and directly proportional to port count in (14). Second, the mirror radius is inversely proportional to resonant frequency through rotational inertia and directly proportional to port count through diffraction and beam divergence angle.

From Fig. 6, ramped plate actuators always outperform parallel (flat) plate actuators, and digitally operated plate actuators always outperform their analog counterparts. The latter can be understood by considering the additional optical steering range gained by allowing the mirror to snap down to the substrate, and the independence of resonant frequency and tilt angle when driving past pull-in. The drawback of digital devices is that they do not scale beyond a few ports in a conventional OXC. One-axis devices are faster than two axis devices in the small port count regime, where lower inertia makes up for the reduced dimensionality of tilt.

One interesting result seen in Fig. 6 is that within the single- and dual-axis subgroups, the fastest actuator changes as a function of port count. Focusing on single-axis devices, the plate actuator operated digitally has $10\times$ the resonant frequency of the next fastest device. The comb actuator is faster than the plate actuator for larger port counts, when, as a consequence of diffraction, the mirror radius has become sufficiently large to overcome the inertial impact of the comb fingers. The one-axis

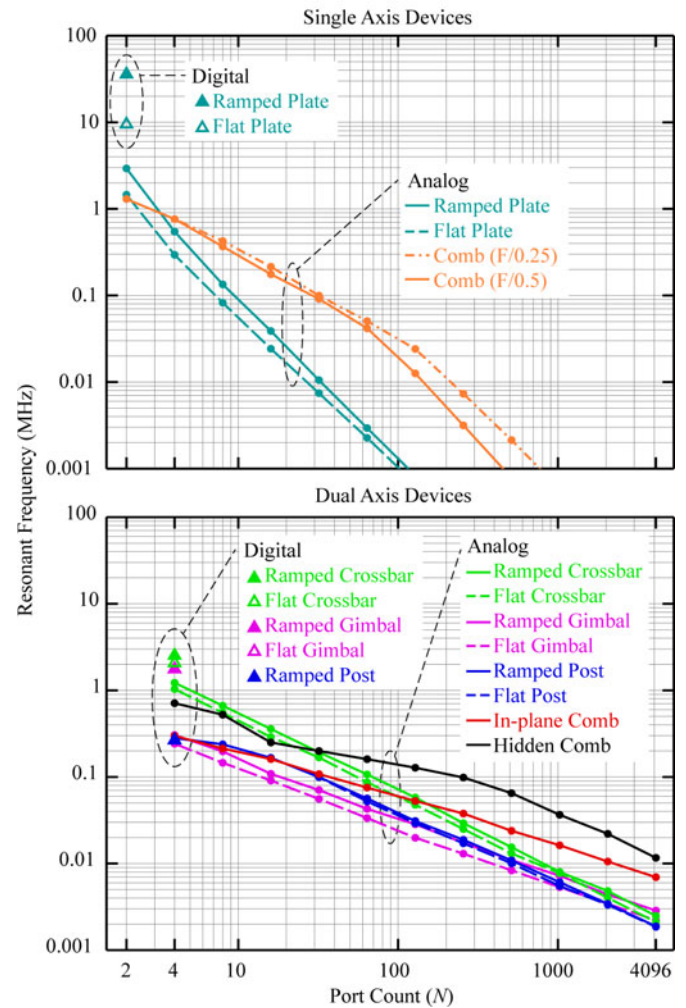


Fig. 6. Log-log plot of resonant frequency versus port count for single-axis (upper) and dual-axis (lower) devices (see Fig. 3) arranged in a conventional $N \times N$ free-space OXC (see Fig. 2), with insertion loss better than 3 dB and crosstalk better than -20 dB. The tilt angle of the one-axis comb drive must be limited for compatibility with a reasonable $F/\#$ Fourier lens.

comb drive naturally optimizes to large tilt angles ($>20^\circ$), so we imposed restrictions on the tilt angle to maintain compatibility with the f -number ($F/\# = \text{focal length divided by full aperture}$) of the Fourier lens. $F/0.25$ may be impractical to achieve due to lens aberrations, but shows the theoretically allowed scaling limit. $F/0.5$ may be achievable with an aspheric curved mirror. Examining the dual-axis devices, we see that more complex designs (hidden comb and crossbar) have better performance than simpler designs. The scaling trends of resonant frequency with port count are largely explained by those of the mirror aperture. Fig. 7 shows the corresponding mirror apertures for each device necessary to achieve the performance shown in Fig. 6. We see that single axis devices require larger mirrors than dual axis devices to reach high port count, accounting for the different scaling trends in resonant frequency.

We used our model to investigate the tradeoff between optical performance parameters (insertion loss and crosstalk) and switching speed in the $N \times N$ OXC topology to determine the

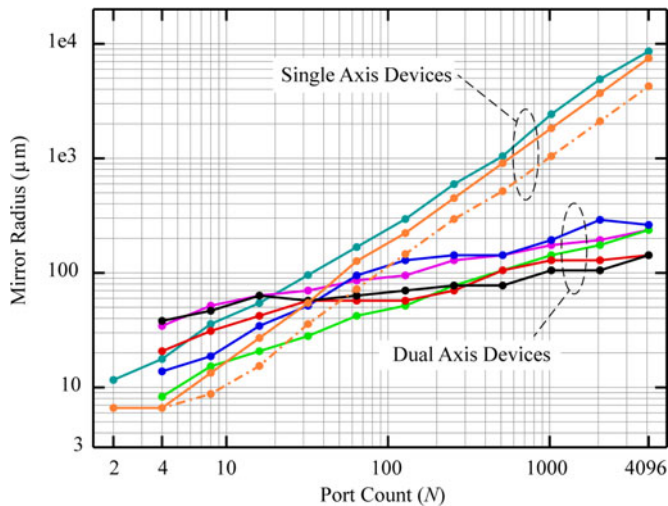


Fig. 7. Log-log plot of mirror radius versus port count for the devices in Fig. 6.




Fig. 8. Log-log plot of resonant frequency versus insertion loss for a 256×256 port OXC. Crosstalk was constrained to be better than -10 dB. Similar trends were seen for all port counts considered.

speed increase that can be gained by surrendering optical performance. Physically, the rotational inertia can be reduced, and the resonant frequency increased, by shrinking the size of the mirror. The smaller mirror, however, spatially clips more of the optical beam and diffracts more light into adjacent ports. Our model showed that a significant increase in switching speed cannot be achieved by a reasonable sacrifice in optical performance. Fig. 8 shows that in the best case, while maintaining a crosstalk of better than -10 dB in a 256 port switch, the resonant frequency (switching speed) of a device can only be improved by about $3\times$, and requires 27 dB excess insertion loss. Similar scaling trends hold for all port counts considered.

Finally, because the finite slew rate of the MEMS driver can limit the speed of the device, we investigated how the resonant frequency varies as a function of electrode voltage. Fig. 9 shows the scaling trends for two of the most promising devices in a 256 port OXC.

B. Detailed Analysis and Comparison to Commercial Switches

To assess the accuracy of our optimization model and the practicality of the optimal systems, we extracted the actuator and switch parameters from our model for the designs shown in Fig. 6. We used Zemax to construct 3-D switch models to ac-




Fig. 9. Log-linear plot of resonant frequency versus electrode voltage for a 256×256 port OXC. Insertion loss was constrained to be ≤ 3 dB and crosstalk better than -20 dB. Similar trends were seen for all port counts considered.




Fig. 10. (a) Optimal two-axis in-plane comb actuator for a 132 port switch, arranged in an array accounting for the asymmetric fill factors and tilt angles. (b) Zemax model of the corresponding system including skew angle, microlenses, and Fourier lens.

count for the skew of the MEMS arrays, the variability of optical path length, and the aberrations associated with the microlenses and Fourier lens. Fig. 10 shows the micromirror and switch system for a 132 port OXC based on the in-plane two-axis comb drive. We tiled the mirrors into an array, accounting for the asymmetric fill factors and tilt angles along each dimension, then used physical optics propagation to model the single mode fiber coupling and crosstalk for both an ideal (paraxial) and biconvex silicon Fourier lens. The results for this and a few other selected designs are summarized in Table I. These devices all operate at 275 V. Based on the Zemax results, we found that our model accurately predicted optical performance for moderate port counts, and still maintained reasonable accuracy at extreme port counts. Our approximation of crosstalk in Section II was conservative, and did not impose an unintended constraint during optimization. We used a reflective Fourier mirror to achieve the $F/0.5$ requirement for the 256 port switch using the one-axis comb drive. Corrections to tilt angle to account for array skew and the non-paraxial Fourier lens were less than 10% of the model output value in all cases. Using an optimized triplet Fourier lens instead of a simple biconvex singlet could further improve performance, but is beyond the scope of this paper.

TABLE I
MODELED OPTIMAL ACTUATOR AND SWITCH PARAMETERS FOR SELECTED CASES

Type	Actuator Parameters														Switch Parameters					
	f (kHz)	θ_{mech} (\pm $^\circ$)	r_m (μm)	t_m (μm)	h_m (μm)	Fill factor	k (Nm/rad)	l_s (μm)	w_s (μm)	t_s (μm)	N_f (μm)	l_f (μm)	t_f (μm)	Ports [layout]	Skew ($^\circ$)	z_R (mm)	Insertion loss (dB)		Crosstalk (dB)	
2-axis In-plane comb	52.8	7.9 (5.8)	57	6.4	-	0.69 (0.56)	1.5×10^{-8} (6.6×10^{-8})	24 (41)	1.1 (1.7)	6.4 (13)	6 (10)	46 (63)	6.4	132 [11 \times 12]	25	4.9	2.4	3.3	-39	-39
2-axis Hidden comb	11.6	10.9 (11.6)	143	1	55	0.95 (0.95)	6.7×10^{-9} (1.2×10^{-8})	123 (123)	1.6 (1.5)	5.2 (10)	30 (30)	18 (22)	5.2	4154 [62 \times 67]	35	23.4	3.0	4.2	-29	-27
2-axis Hidden crossbar	2.5	10.2 (10.2)	237	2.6	105	0.70 (0.70)	6.7×10^{-9} (6.7×10^{-9})	237 (237)	1.1 (1.1)	26 (26)	-	-	-	4096 [64 \times 64]	23	85.6	2.9	3.5	-34	-33
1-axis Comb (F/0.5)	3.2	12.8	449	22	-	1	6.7×10^{-7}	12	2	22	236	126	22	256 [1 \times 256]	0.5	240	3.6	5.2	-31	-26

Table showing design parameters extracted from the optimization model, at a few interesting points shown in Fig. 6. The insertion loss and crosstalk are calculated using physical optics propagation in Zemax for a paraxial Fourier lens and a biconvex singlet Fourier lens. For two-axis actuators, the value for the inner axis is indicated without parentheses and the value for the outer axis is indicated with parentheses. The layout of the micromirror array is indicated by brackets.

The idealized switch structures in this model lead to a theoretically achievable performance which is not necessarily compatible with a practical switch constrained by the many design factors required for a manufacturable and cost-effective product. However, it is useful to discuss these factors to understand how a theoretical design can be translated into reality. Consider the optimized 132 port switch shown in Fig. 10. The predicted 3 dB insertion loss requires perfect alignment of the collimator and fiber arrays. To account for misalignments between the fiber array, microlens array, and MEMS array while keeping insertion loss low, we must redesign for low theoretical insertion loss and let the mirror aperture grow. Redesigning for 0.2 dB nominal loss requires the mirror aperture to increase by $2\times$, and the resonant frequency is reduced by $2\times$. To reduce cost, the Fourier lens may be omitted, which requires the tilt range to increase by $2\times$ in each dimension. Because the tilt range is already large, we can instead double the distance between mirror arrays and double the mirror aperture. Also, because the beam is no longer focused onto the mirror, the mirror aperture must increase by $1.5\times$ and the mirror pitch must increase by $2\times$, again requiring larger tilt, or a larger mirror. The omission of the Fourier lens cumulatively increases mirror aperture by $4\times$ and reduces resonant frequency by $4\times$. Next, to maximize the reliability of drive electronics, the electrode voltage might be reduced to 150 V. This requires weaker springs to maintain the same tilt range, and reduces the resonant frequency by $2\times$. To account for imperfect fiber and MEMS array yield, we can add redundant elements to the arrays. Assuming an 80% yield for fiber and MEMS arrays, we must increase the number of elements in the array by 50%. Finally, to avoid complicated multilayer electrical routing, we might use planar routing and increase the pitch between mirror elements. To account for yield and planar routing, we may let the mirror radius increase by $1.5\times$ and reduce the spring constant further, reducing the resonant frequency by $2\times$. Accounting for all these factors, the resonant frequency is reduced to 1.6 kHz and the mirror radius is increased to 680 μm . Commercial switches are often operated at some fraction (1/10th) of the mirror's resonant frequency to allow mirror ringing to subside. This gives a response time of 6.2 ms, which is more comparable to that of commercial switches with ~ 100 ports. Besides the changes in switching speed and physical scale, the switch still looks very

similar to the one shown in Fig. 10 (ignoring the Fourier lens), but the cost of manufacture is significantly reduced. Supposing that similar relationships hold for all port counts considered in our study, the relative changes in the theoretical limits calculated here (see Figs. 6 and 8) should be reflected in real switches.

IV. MULTISTAGE SWITCH ARCHITECTURES

Our scaling study indicates that new overall optical switch topologies will be needed to achieve microsecond-scale switching with the large port count necessary for data center networks, as opposed to simply modifying device structures within existing telecommunications switches. Here we describe two switch architectures to illustrate how multistage topologies can allow better scaling properties. We used the skew ray representation of Gaussian beams [32] to design the switches and physical optics propagation in Zemax to model single mode fiber coupling.

A. Multiport Wavelength Selective Switch

Telecom multiport wavelength switches use digital beam steering and aperture division [33]. Introducing an array of relay lenses located near a digital micromirror array can extend the port count of $1 \times N$ wavelength selective switches while retaining the microsecond-scale reconfiguration rate of two-state micromirrors.

Fig. 11 illustrates this switch structure. The input signal is spectrally demultiplexed by a reflective dispersive element in the Fourier plane, with the fiber and MEMS arrays located at the focal plane of the Fourier lens. Each wavelength channel is spatially separated in y at the MEMS array and is independently relayed laterally in x by the tilting mirrors. By purposely adding spherical power to the micromirrors (possible by greyscale lithography [34] or stress induced bending [35]) and using an array of reflective micro optics to form a 4-f relay, the beam parameters can be relayed between mirrors with a minimal number of reflections (see Fig. 11(b)). A refractive microlens at the input of the relay focuses the spectral components of a single wavelength channel onto the micromirror to maintain a wide passband (see Fig. 11(c)). Because the 4-f relay spatially inverts the spectral components of the passband with each pass through the relay, an output port can only be located at every second micromirror.

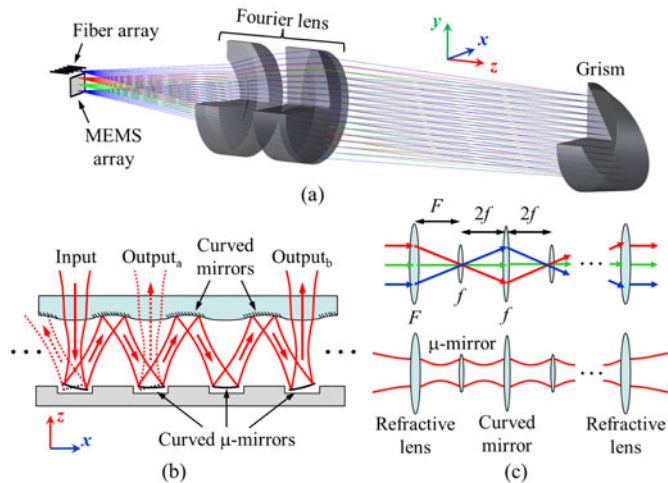


Fig. 11. (a) Zemax model showing dispersion along y at the MEMS array. Ray color indicates different wavelength channels. (b) Schematic x - z cross section at the MEMS array showing a single wavelength channel being relayed to the right or left in x using an array of curved mirrors. (c) Unfolded schematic showing the simultaneous relaying of three spectral components of a single channel and the Gaussian beam parameters when $z_R = f$ and $F = 2f$.

The size of the relay lenses imposes a spectral separation in y , and to keep the overall system length to a minimum we used a grism (grating-prism combination) with stronger dispersion than a standard grating. We modeled a 1×12 port switch supporting 128 50 GHz-spaced C-band channels in Zemax (see Fig. 11(a)). We assumed all surfaces had a high reflectivity coating, such as Newport DM.8, which is 99.5% reflective up to 45° at 1550 nm. Performing Gaussian beam propagation in Zemax, and assuming a Grism efficiency of 80% for a custom blaze angle, we found the worst case insertion loss was 3.1 dB, with a 25 GHz passband with 0.5 dB excess loss.

B. Multistage OXC

In Section III and Fig. 6, we observed that the switching speed of an $N \times N$ OXC can be increased by reducing the port count. The second illustrative switch geometry uses free-space optics to interconnect many small port count OXC “sub-switches” in a multistage network to form an $N \times N$ switch which retains the faster reconfiguration rate of the small sub-switches.

Fig. 12 shows an illustration of the switch structure. The drawing seems to show a cascade of three fully interconnected $N \times N$ switches. In fact, the tilt range of every micromirror in each OXC switch has been reduced, sacrificing full connectivity within a single OXC, but allowing a faster reconfiguration rate through the inverse relationship between tilt range and resonant frequency. Full, non-blocking connectivity between all ports is regained by interconnecting three active switching stages in a Clos network [36], provided the stages are interconnected with a suitable port-mapping structure. This could be done with fiber cabling, but would triple the insertion loss. Instead, the necessary port mapping between stages can be accomplished with an optical transpose interconnection implemented by relay imaging [37], reducing optical loss and complexity compared to fiber optic connections. Fig. 12(b) shows the passive optics required, using two prism arrays and two lens arrays to redirect and relay light output from the first switch to the input of the next switch.

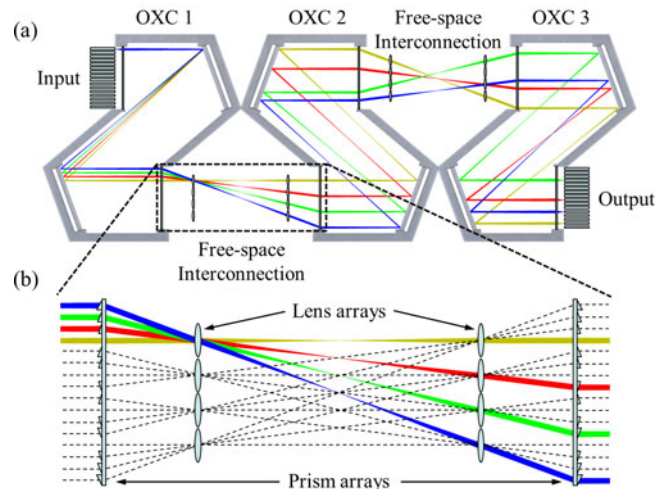


Fig. 12. (a) Cross section of three 256 port OXCs, each with $4\times$ reduced scan angle, interconnected with free-space optics. Ray color indicates different optical paths through the system, depending on mirror states. (b) Detail of the passive free-space interconnection with possible connections.

For the two-axis hidden crossbar device, our model predicts a $4\times$ reduction in tilt angle in a 256 port switch increases the switching speed by $3\times$. The reduction in tilt allows a higher resonant frequency due to stiffer springs, but the gain in speed is limited because stiffer springs require a thicker mirror to prevent bending of the mirror under actuation force. We performed Gaussian beam analysis in Zemax and found the worst-case insertion loss after three stages was 7.7 dB accounting for the accrued path length differences.

A more substantial increase in speed can be achieved by reducing the mirror aperture, as this increases resonant frequency by decreasing inertia and does not require an increase in mirror thickness. Our model indicates a $4\times$ reduction in mirror radius can increase switching speed by $10\times$ for a 256 port switch. However, the smaller micromirror aperture requires a larger beam divergence to maintain high spatial confinement efficiency at the mirrors. The increased beam divergence can be accommodated by adding relay optics between collimators and MEMS arrays and between MEMS arrays in the system shown in Fig. 12, providing a more substantial increase in switching speed at the cost of increased optical complexity.

V. SUMMARY

The requirements on switching speed, port count, and optical efficiency in the data center are strongly dependent on network topology and application. Without making assumptions about specific network topologies or applications, we have quantified the relationship between speed, port count, and optical efficiency for free-space MEMS beam-steering switches to help determine how this technology fits within the data center environment. Based on our prediction that conventional telecom switches may not scale far beyond their current performance without significant increase in optical loss, we suggested two multistage switch architectures that help extend the capability of MEMS tilt mirror technology into the data center regime. Future work will demonstrate multistage switch capabilities and investigate their use in data center network topologies.

REFERENCES

- [1] N. Farrington, G. Porter, S. Radhakrishnan, H. Bazzaz, V. Subramanya, Y. Fainman, G. Papan, and A. Vahdat, "Helios: A hybrid electrical/optical switch architecture for modular data centers," in *Proc. ACM SIGCOMM Conf.*, 2010, pp. 339–350.
- [2] T. Goh, M. Yasu, K. Hattori, A. Himeno, M. Okuno, and Y. Ohmori, "Low-loss and high-extinction-ratio silica-based strictly nonblocking 16×16 thermo-optic matrix switch," *Photon. Technol. Lett.*, vol. 10, no. 6, pp. 810–812, 1998.
- [3] S. Han, T. Seok, N. Quack, B.-W. Yoo, and M. Wu, "Monolithic 50×50 MEMS silicon photonic switches with microsecond response time," in *Proc. Opt. Fiber Commun. Conf.*, 2014, vol. 1, no. 3, pp. 9–13.
- [4] M. Wu, O. Solgaard, and J. Ford, "Optical MEMS for lightwave communication," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4433–4454, Dec. 2006.
- [5] P. Hagelin, U. Krishnamoorthy, J. Heritage, and O. Solgaard, "Scalable optical cross-connect switch using micromachined mirrors," *Photon. Technol. Lett.*, vol. 12, no. 7, pp. 882–884, 2000.
- [6] Y. Xie, S.-S. Li, Y.-W. Lin, Z. Ren, and C. Nguyen, "1.52-GHz micromechanical extensional wine-glass mode ring resonators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 55, no. 4, pp. 890–907, Apr. 2008.
- [7] R. Knipe, "Challenges of a digital micromirror device: Modeling and design," *Proc. SPIE, Micro-Opt. Technol. Meas., Sens., Microsyst.*, vol. 2783, pp. 135–145, 1996.
- [8] G. Porter, R. Strong, N. Farrington, A. Forencich, P. Chen-sun, T. Rosing, Y. Fainman, G. Papan, and A. Vahdat, "Integrating microsecond circuit switching into the data center," in *Proc. ACM SIGCOMM Conf.*, 2013, pp. 447–458.
- [9] J. Kim, C. Nuzman, B. Kumar, D. Lieuwen, J. Kraus, A. Weiss, C. Lichtenwalner, A. Papazian, R. Frahm, N. Basavanthally, D. Ramsey, V. Aksyuk, F. Pardo, M. Simon, V. Lifton, H. Chan, M. Haueis, A. Gasparyan, H. Shea, S. Arney, C. Bolle, P. Kolodner, R. Ryf, D. Neilson, and J. Gates, "1100 \times 1100 Port MEMS-based optical crossconnect with 4-dB maximum loss," *Photon. Technol. Lett.*, vol. 15, no. 11, pp. 1537–1539, 2003.
- [10] R. Syms, "Scaling laws for MEMS mirror-rotation optical cross connect switches," *J. Lightw. Technol.*, vol. 20, no. 7, pp. 1084–1094, Jul. 2002.
- [11] U. Krishnamoorthy, D. Lee, and O. Solgaard, "Self-aligned vertical electrostatic combdrives for micromirror actuation," *J. Microelectromech. Syst.*, vol. 12, no. 4, pp. 458–464, 2003.
- [12] G. Nielson and G. Barbastathis, "Dynamic pull-in of parallel-plate and torsional electrostatic MEMS actuators," *J. Microelectromech. Syst.*, vol. 15, no. 4, pp. 811–821, 2006.
- [13] H. Ishii, M. Urano, Y. Tanabe, T. Shimamura, J. Yamaguchi, T. Kamei, K. Kudou, M. Yano, Y. Uenishi, and K. Machida, "Fabrication of optical microelectromechanical-system switches having multilevel mirror-drive electrodes," *Jpn. J. Appl. Phys.*, vol. 43, no. 9A, pp. 6468–6472, 2004.
- [14] J. Tsai and M. Wu, "Gimbal-less MEMS two-axis optical scanner array with high fill-factor," *J. Microelectromech. Syst.*, vol. 14, no. 6, pp. 1323–1328, 2005.
- [15] I. Jung, U. Krishnamoorthy, and O. Solgaard, "High fill-factor two-axis gimbaled tip-tilt-piston micromirror array actuated by self-aligned vertical electrostatic combdrives," *J. Microelectromech. Syst.*, vol. 15, no. 3, pp. 563–571, 2006.
- [16] F. Pardo, R. Cirelli, E. Ferry, W. Lai, F. Klemens, J. Miner, C. Pai, J. Bower, W. Mansfield, A. Kornblit, T. Sorsch, J. Taylor, M. Baker, R. Fullowan, M. Simon, V. Aksyuk, R. Ryf, H. Dyson, and S. Arney, "Flexible fabrication of large pixel count piston-tip-tilt mirror arrays for fast spatial light modulators," *Microelectron. Eng.*, vol. 84, pp. 1157–1161, 2007.
- [17] D. Hah, H. Toshiyoshi, and M. Wu, "Design of electrostatic actuators for MOEMS applications," *Proc. Symp. Des., Test, Integr. Packag. MEMS/MOEMS*, May 2002, pp. 200–207.
- [18] D. Lee and O. Solgaard, "Pull-in analysis of torsional scanners actuated by electrostatic vertical combdrives," *J. Microelectromech. Syst.*, vol. 17, no. 5, pp. 1228–1238, 2008.
- [19] R. Sattler, F. Plotz, G. Fattinger, and G. Wachutka, "Modeling of an electrostatic torsional actuator: Demonstrated with an RF MEMS switch," *Sens. Actuators A*, vols. 97/98, pp. 337–346, 2002.
- [20] C. Pu, S. Park, P. Chu, S. Lee, M. Tsai, D. Peale, N. Bonadeo, and I. Brener, "Electrostatic actuation of three-dimensional MEMS mirrors using sidewall electrodes," *J. Sel. Topics Quantum Electron.*, vol. 10, no. 3, pp. 472–477, 2004.
- [21] U. Krishnamoorthy, K. Li, K. Yu, D. Lee, J. Heritage, and O. Solgaard, "Dual-mode micromirrors for optical phased array applications," *Sens. Actuators A*, vols. 97/98, pp. 21–26, 2002.
- [22] H. Shea, A. Gasparyan, H. Chan, S. Arney, R. Frahm, D. Lopez, S. Jin, and R. McConnell, "Effects of electrical leakage currents on MEMS reliability and performance," *IEEE Trans. Device Mater. Rel.*, vol. 4, no. 2, pp. 198–207, Jun. 2004.
- [23] H. Urey, C. Kan, and W. Davis, "Vibration mode frequency formulae for micromechanical scanners," *J. Micromech. Microeng.*, vol. 15, pp. 1713–1721, 2005.
- [24] M. Bao and H. Yang, "Squeeze film air damping in MEMS," *Sens. Actuators A, Phys.*, vol. 136, no. 1, pp. 3–27, 2007.
- [25] D. Greywall, P. Busch, and J. Walker, "Phenomenological model for gas-damping of micromechanical structures," *Sens. Actuators A, Phys.*, vol. 72, no. 1, pp. 49–70, 1999.
- [26] P. Chu, I. Brener, C. Pu, S.-S. Lee, J. Dadap, S. Park, K. Bergman, N. Bonadeo, T. Chau, M. Chou, R. Doran, R. Gibson, R. Harel, J. Johnson, C. Lee, D. Peale, B. Tang, D. Tong, M. Tsai, Q. Wu, W. Zhong, E. Goldstein, L. Lin, and J. Walker, "Design and nonlinear servo control of MEMS mirrors and their performance in a large port-count optical switch," *J. Microelectromech. Syst.*, vol. 14, no. 2, pp. 261–273, 2005.
- [27] D. Burns and V. Bright, "Nonlinear flexures for stable deflection of an electrostatically actuated micromirror," *Proc. SPIE Microelectron. Struct. MEMS Opt. Process. III*, vol. 3226, pp. 125–136, 1997.
- [28] P. Brosens, "Dynamic mirror distortions in optical scanning," *Appl. Opt.*, vol. 11, no. 12, pp. 2987–2989, 1972.
- [29] Y. Low, R. Scotti, D. Ramsey, C. Bolle, S. O'Neill, and K. Nguyen, "Packaging of Optical MEMS Devices," *J. Electron. Packag.*, vol. 125, pp. 325–328, 2003.
- [30] M. Kozhevnikov, R. Ryf, D. Neilson, P. Kolodner, D. Bolle, A. Papazian, J. Kim, and J. Gates, "Micromechanical optical crossconnect with 4-F relay imaging optics," *Photon. Technol. Lett.*, vol. 16, no. 1, pp. 275–277, 2004.
- [31] A. Ankiewicz and G. Peng, "Generalized Gaussian approximation for single-mode fibers," *J. Lightw. Technol.*, vol. 10, no. 1, pp. 22–27, Jan. 1992.
- [32] P. Colbourne, "Generally astigmatic Gaussian beam representation and optimization using skew rays," *Proc. SPIE*, vol. 9293, IODC, 92931S, pp. 1–11, 2014.
- [33] J. Ford, V. Aksyuk, D. Bishop, and J. Walker, "Wavelength add/drop switching using tilting micromirrors," *J. Lightw. Technol.*, vol. 17, no. 5, pp. 904–911, May 1999.
- [34] C. Chen, S. Tzeng, and S. Gwo, "Silicon microlens structures fabricated by scanning-probe gray-scale oxidation," *Opt. Lett.*, vol. 30, no. 6, pp. 652–654, 2005.
- [35] P. Townsend and D. Barnett, "Elastic relationships in layered composite media with approximation for the case of thin films on a thick substrate," *J. Appl. Phys.*, vol. 62, no. 11, pp. 4438–4444, 1987.
- [36] W. Dally and B. Towels, "Non-blocking networks," in *Principles and Practices of Interconnection Networks*. San Francisco, CA, USA: Elsevier, 2004, pp. 111–144.
- [37] G. Marsden, P. Marchand, P. Harvey, and S. Esener, "Optical transverse interconnection system architectures," *Opt. Lett.*, vol. 18, no. 13, pp. 1083–1085, 1993.

William Maxwell Mellette received the B.S. degree in engineering physics and the M.S. degree in photonics from the University of California San Diego (UCSD), La Jolla, CA, USA, in 2012 and 2014, respectively, where he is currently working toward the Ph.D. degree in photonics. He held internships with General Atomics Aeronautical Systems' Electromagnetics and Photonics groups and has been with the Photonic Systems Integration Lab, UCSD, since 2012. His current research interests include optical communications, free-space optics, and both imaging and nonimaging optics.

Joseph E. Ford is currently a Professor at the Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA, USA. He currently leads UCSD's Photonics Systems Integration Lab, La Jolla, a research group doing free-space optical system design, prototyping, and characterization for a wide range of applications in communications, energy, and sensing. At AT&T Bell Labs from 1994 to 2000, his research interests include MEMS attenuator, spectral equalizer, and wavelength add/drop switch. He is the Coauthor on 49 United States patents and more than 200 journal articles and conference proceedings. He was the General Chair of the first *IEEE Conference on Optical MEMS in 2000*, and General Chair for the 2008 OSA Optical Fiber Communications Conference. He is a Fellow of the Optical Society of America.