

Circuit Design for Scalable and Fast Optical Circuit Switching

Erik Anderson

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2024-213

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2024/EECS-2024-213.html>

December 14, 2024



Copyright © 2024, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Circuit Design for Scalable and Fast Optical Circuit Switching

By

Erik Francis Anderson

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Vladimir Stojanović, Co-chair

Professor Ming C. Wu, Co-chair

Professor Martin White

Fall 2024

Circuit Design for Scalable and Fast Optical Circuit Switching

Copyright 2024
by
Erik Francis Anderson

Abstract

Circuit Design for Scalable and Fast Optical Circuit Switching

by

Erik Francis Anderson

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Science

University of California, Berkeley

Professor Vladimir Stojanović, Co-chair

Professor Ming C. Wu, Co-chair

Data centers and large-scale distributed computers have become limited by the latency, throughput, and inflexibility of traditional electronic packet switches (EPSs). As evidenced by the recent introduction of optical circuit switches (OCSs) into Google's datacenters and TPU clusters, OCSs provide a way to circumvent many of the limitations of EPS networks. Silicon-Photonic (SiPh) MEMS-based OCSs have been shown to offer a scalable and low-latency approach compared to other integrated and non-integrated OCSs. Yet to be realized, however, is the electrical control and digital interface required for integrating high-radix MEMS SiPh chips into an application environment. This work demonstrates two novel approaches to controlling SiPh MEMS OCSs in both a scalable and efficient manner.

To my parents.

Contents

Contents	ii
List of Figures	iii
List of Tables	iv
1 Introduction to Optical Circuit Switching	1
1.1 Circuit vs. Packet Switching	1
1.2 Current Applications	2
1.3 Future Applications	3
1.4 OCS Technologies	4
2 SuperSwitch 1: A Digital Approach	6
2.1 Problem Statement	6
2.2 CMOS Design & Simulation	9
2.3 3D SiPh-CMOS Packaging	21
2.4 PCB Design	23
2.5 Electro-optic Characterization Setup	24
2.6 Results	27
3 SuperSwitch 2: An Analog Alternative	31
3.1 Problem Statement	31
3.2 CMOS Design & Simulation	37
3.3 Wirebond-based SiPh-CMOS Packaging	52
3.4 Test Setup	53
3.5 Results	54
4 Conclusion	60
4.1 Final Results	60
4.2 Thesis Contribution	62
4.3 Future of SiPh MEMS OCS	62

List of Figures

1.1	(a) Optically-connected 4x4 electronic packet switch (EPS). (b) All-optical 4x4 optical circuit switch (OCS).	2
1.2	(a) Initial uniform ring topology. (b) Traffic matrix extracted from a specific application showing the relative amount of transmitted data. (c) New ring topology optimized for the extracted traffic matrix.	3
2.1	(a) Schematic of (a) OFF state and (b) ON state of vertical MEMS coupler - figures taken from [14] (c) SEM of MEMS switch cell - figure taken from [22] . .	7
2.2	A representative transfer function of a MEMS switch element.	8
2.3	GDS screenshots of (a) 128x128 SiPh MEMS crossbar switch (b) 4x4 CMOS high voltage driver chiplets flip-chip bonded to SiPh MEMS chip.	9
2.4	(a) Micrograph of SuperSwitch 1 high voltage driver chiplet (b) Cartoon floorplan of driver chiplet.	10
2.5	Simple schematic for activating 1 of 128 rows assuming a single CMOS chiplet. .	11
2.6	Schematic for controlling a 128x128 switch assuming a 4x4 array of CMOS chiplets.	12
2.7	(a) Logic for columns 0 and 1 for $N_c = 1$ (b) Logic for columns 0 and 1 for $N_c = 2$.	13
2.8	(a) Final schematic of SuperSwitch1 control chiplet scan architecture w/ loopback mux for debug. (b) Final parameters for SuperSwitch1 controller chiplet.	14
2.9	(a) Schematic of SuperSwitch1 high voltage driver circuit. (b) List of all supplies and their nominal values.	15
2.10	Shmoo plot for 10-100 fF bootstrap capacitance.	17
2.11	Basic initialization procedure for SuperSwitch1 high voltage driver circuit. . . .	18
2.12	Equivalent circuit of supply and ground connection for a single chiplet column in the flip-chip bonded CMOS-SiPh package.	19
2.13	(a) VSS resistance shmoo plot for all corners for HVDD = 70 V, HVSS = 65 V. (b) Same plot but for HVDD = 70 V, HVSS = 66 V.	20
2.14	Rise (a) and fall (b) times extracted from post-layout simulations.	20
2.15	(a) Transient simulation of capacitance of MEMS model (C_{MEMS}) and output driver voltage (V_{OUT}). (b) Parasitic output capacitance for all 1024 drivers (due to short routes and pads).	21
2.16	(a) Micrograph of bumped CMOS pads. (b) Cross section of Au UBM and Au micro-bumps. (c) Illustration of bonding process using differing thicknesses of UBM on SiPh chip to compensate for CMOS pad height differences.	22

2.17	(a) 4x4 CMOS flip-chip bonded to electrical interposer. (b) 1 CMOS flip-chip bonded to 32x32 subsection of 128x128 OCS.	23
2.18	(a) The chip board for wirebonding chips to. (b) The host board PCB used for testing chip boards.	24
2.19	Simplified diagram of the PC-to-chip interface with UART-controlled scan chain controllers.	25
2.20	Illustration of the SuperSwitch1 characterization test setup.	26
2.21	The physical test setup for fully automated switch characterization.	26
2.22	(a) Location of switch cell being tested. (b) Captured scope waveforms used for extracting rise and fall times. (c) Rise and fall times. vs applied voltage. (d) Optical transfer curve.	28
2.23	Measured on-chip loss (a/d), rise time (b/e), and fall time (c/e) for all $32^2 = 1024$ switch elements. Measurements were done at an HVDD of 60V and a wavelength $\lambda = 1310\text{nm}$	29
2.24	(a) Example 4x4 OCS w/ one driver "ON" per row/column. (b) Worst case functional pattern for power measurements.	29
2.25	(a) Measured worst case power for 32x32 OCS w/ HVDD=60V. (b) Power breakdown by supply for maximum reconfiguration frequency of $\approx 1.7\text{MHz}$	30
3.1	(a) Cross section of AIM Photonics base active PIC process - figure taken from [31]. (b) Simplified model of MEMS cantilever switch.	32
3.2	(a) Cartoon of example 4x4 crossbar switch. "On" refers to a high voltage being applied. (b) Screenshot of GDS for SuperSwitch 2 silicon photonic MEMS 8x8 OCS.	33
3.3	Circuit model for high voltage driver and MEMS device.	34
3.4	Optical transmission vs. displacement for a single MEMS switch element.	35
3.5	(a) Step response for 550 nm displacement (b) Dual step response for 550 nm displacement	36
3.6	(a) Step response peak displacement and time vs. applied voltage (b) Step response steady state displacement vs. time	37
3.7	SuperSwitch 2 - high voltage DAC array for analog MEMS control	38
3.8	(a) Dual step waveforms w/ voltage offset (b) Cantilever displacement w/ voltage offset (c) Static error (d) Settling time for +/- 10%	39
3.9	Simplified chip-level digital architecture showing control and select scan chains	40
3.10	Simplified schematic of a single cell controller	43
3.11	(a) Schematic of high voltage driver cell - M_4 and M_5 are high voltage transistors. (b) Sizing guide for driver given a particular I_{BIAS}	44
3.12	(a) Histogram of parasitic output cap from all 64 HV DAC outputs to their corresponding pads. (b) Capacitance of MEMS device as voltage is stepped.	45
3.13	(a) Schematic of 6-bit segmented current-steering DAC. (b) Schematic of single current-steering DAC segment.	47
3.14	Schematic for current reference to set LSB current of current-steering DAC.	48

3.15	Selected and deselected waveforms. The select scan enable signal is shown as well as the load signal that will trigger each cell controller's FSM to either start cycling through the programmed selected or deselected waveform.	49
3.16	(a) Simplified FSM diagram for stepping through selected and deselected waveforms. (b) Description of FSM state actions. (c) Logic to generate the load signal that triggers the FSM.	51
3.17	Global current distribution circuit.	51
3.18	(a) Output voltage of HV DAC going from minimum to maximum code driving a 1pF load. (b) Extracted rise times. (c) Extracted fall times.	52
3.19	Wirebond diagram for SuperSwitch 2 using a 3-chip package for controlling the 8x8 OCS.	53
3.20	CMOS only test package with (a) CMOS 1, the west side control chip from Fig. 3.19, (b) side-view of wirebonds, and (c) top-view of wirebonds.	53
3.21	The test setup for SuperSwitch 2.	54
3.22	Simulated and measured DNL and INL.	55
3.23	(a) Ideal voltage vs. measured DAC voltages. (b) Ideal displacement and displacement using measured DAC values.	57
3.24	(a) Estimated and measured power for an HVDD of 50V and 70V. (b) The estimated average power per high voltage DAC (HVDAC) vs. radix.	58

List of Tables

2.1	Global "address" scan chain bits listed in the order that they are connected. . .	14
3.1	Model parameters for the SuperSwitch 2 MEMS unit cell.	34
3.2	Mapping between column (0-7 for 8x8 switch) and HVDAC IDs.	40
3.3	Global "Select" scan chain bits listed in the order that they are connected. . . .	41
3.4	Global "Control" scan chain bits listed in the order that they are connected. . .	42
3.5	List of cell-controller control scan chain fields. Each cell controller has its own unique set of these fields.	42
3.6	List of all IDAC control scan chain registers. These bits are labelled as idac_cfg in table 3.5	47
3.7	List of all IREFP control scan chain registers. These bits are labelled as irefp_cfg in table 3.5	48
3.8	List of all FSM control scan chain registers. These bits are labelled as fsm_cfg in table 3.5	50
4.1	Comparison table for optical circuit switches.	61
4.2	Comparison table for HV driver circuits.	62

Acknowledgments

Having the opportunity to pursue a PhD is an incredible privilege. My time at Berkeley has been transformative and will continue to shape me throughout my life. I am deeply grateful for this invaluable experience.

The biggest thanks must go to my advisor, Vladimir Stojanović, for believing in me from day one. Your passionate and patient approach to mentorship has made me a more confident engineer and person. I will strive to mentor young engineers in the same manner throughout my career.

I'd also like to thank Professors Ming Wu, Sophia Shao, and Martin White for providing invaluable feedback on my research. Ming has been a wonderful secondary advisor, and I deeply appreciate all the time he spent guiding me. It has been an absolute pleasure working with you.

Thank you to all my friends and peers who made this experience so enjoyable, despite COVID's attempts to derail our social lives: Derek, Ryan, Sunjin, Lars, Avi, Kevin, Johannes, Sidney, Kramnik, Sarika, Hyeong, and everyone from the Integrated Systems Group.

Balancing school and work hasn't been easy, and I absolutely could not have done it without the support of my boss, Norman Chan. Thank you for accommodating my odd hours and last-minute schedule changes. I've thoroughly enjoyed working with you and the rest of the ASIC team at Ayar Labs.

BWRC is truly a special place and it would cease to operate without the tireless efforts of Candy and Mikaela. Thank you for everything that you've done for me and for the many students before me!

Lastly, thank you to my family for supporting me through all my successes and failures. I would not be where I am without you. A special thanks to my partner, Megan, who stood by me through all the anxiety and challenges that came with pursuing a PhD. Your unwavering support and understanding mean the world to me, and I am so fortunate to have you in my life.

This work was supported by the DARPA PIPES program, BETR Center, Berkeley Wireless Research Center, and the SRC and DARPA co-sponsored CUbiC Research Center. I would also like to extend special thanks to Muse Semiconductor, S&C Micro, T-Microtec, and Micross for their chip fabrication and packaging support.

Chapter 1

Introduction to Optical Circuit Switching

This thesis explores new methods of optical circuit switching using specifically designed CMOS circuits for fast and scalable control. The following sections introduce the concept of optical circuit switching and identify the target application and key technology for the designs presented in Chapters 2 and 3.

1.1 Circuit vs. Packet Switching

Circuit switching gets its name from the old switchboard telephone days. The circuit, the phone in this case, was physically connected to another circuit (phone) through the switchboard. The operator had to physically disconnect one connection before they could establish another. Packet switching, on the other hand, allows for data transfer between many devices at once without physically reconfiguring the switch paths. Contention is handled by processing data coming into the switch and buffering it until a viable output path is available. Optical circuit switches (OCS) perform a similar function to the old telephone switchboards except that their data is optical in nature. As datacenters and high performance computing (HPC) systems continue to grow, more and more of the data is carried optically to benefit from low-loss and distance-indifferent high-speed operation compared to state-of-the-art electrical interconnects.

Fig. 1.1 (a) shows an illustration of an optically-connected electronic packet switch (EPS). These types of switches are prevalent in all HPC and datacenter networks as they leverage the benefits of both optical links and packet switching. All-optical packet switching has been heavily researched [3], but has never gained traction due to the difficulty of buffering data in the optical domain. Optically-connected EPSs are thus required to perform optical-to-electrical (O2E) and electrical-to-optical (E2O) conversions such that the data can be easily buffered in the electrical domain. These conversions add a huge power and latency penalty compared to the optical circuit switch (OCS) shown in Fig. 1.1 (b). The direct connections

offered by the OCS reduce the power and latency of the switching process but require that the connections be scheduled in advance. Bursty and random network traffic makes it hard to utilize a circuit switch as the primary switching element. Despite this, the replacement of some optically-connected EPSs with OCSs in a network can eliminate the inefficiencies of EPSs while maintaining, and in some cases augmenting, the network performance gains afforded by pure packet switching.

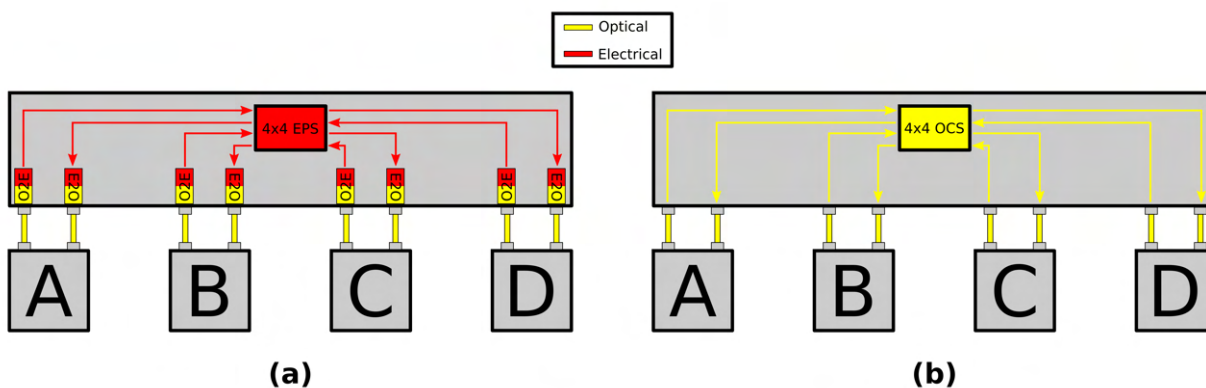


Figure 1.1: (a) Optically-connected 4x4 electronic packet switch (EPS). (b) All-optical 4x4 optical circuit switch (OCS).

1.2 Current Applications

Google published for the first time in 2022 on their use of OCSs within their datacenters [32]. They replaced the spine EPSs with an OCS connection layer to build a reconfigurable direct-connect mesh topology between aggregation blocks. Many of the stated advantages of using this OCS layer over the EPSs are related to the problem of deploying, upgrading, and maintaining large amounts of expensive equipment. Without the spine switches, the network can be more easily expanded and upgraded, as the OCS switching layer is trivially reconfigured via software. The OCSs are, in general, data rate agnostic and allow for incremental upgrades of the aggregation blocks without corresponding upgrades to the spine switching layer. All of these benefits save Google time, energy, and most importantly, money.

Perhaps the most interesting use of these OCSs is highlighted by Google's dynamic reconfiguration of their network topologies. Google detailed their so-called "topology engineering" for both their datacenters [32] and TPU clusters [18] but much research has already been done (and continues to be done) in this vein [47] [46] [42]. Google was simply the first major service provider to implement this type of topology optimization at scale. Fig. 1.2 shows a simple example of this topology optimization. The initial topology, Fig. 1.2 (a), is a basic uniform ring topology with each node having equal bandwidth between itself and its two

neighboring nodes. This example allocates each node A-D 6 transmitters and 6 receivers (this number is chosen to make the example simple). This means that each row and column in the topology matrix shown in Fig. 1.2 (a) must sum to a maximum of 6. The traffic matrix for a given application is then extracted for a given period of time (Fig. 1.2 (b) shows the relative amount of transmitted data). Assuming that this extracted traffic matrix is representative of future traffic, the network topology can be reconfigured to better utilize the total available bandwidth (Fig. 1.2 (c)). Google was able to perform a more complicated topology optimization to improve specific network metrics relative to their baseline Clos topology.



Figure 1.2: (a) Initial uniform ring topology. (b) Traffic matrix extracted from a specific application showing the relative amount of transmitted data. (c) New ring topology optimized for the extracted traffic matrix.

1.3 Future Applications

Google has shown that OCS-enabled dynamic topology optimization can be done to increase network performance. The frequency at which this reconfiguration is done is ultimately limited by the physical speed of switching. Google's current datacenter workloads only require topology optimizations at timescales much larger than the slowest OCS switching speeds. Google reports optimizing their topology hourly [32], much longer than their quoted millisecond-scale switching speed [45]. Several papers, including those from Google, have highlighted the potential benefits of more frequent reconfiguration for different kinds of workloads, namely distributed AI/ML workloads [19] [24]. These workloads experience both repetitive and bursty network traffic, making them ideal candidates for pre-scheduled topology changes. Unfortunately, millisecond-scale switching will likely incur too high a latency penalty to make frequent topology changes beneficial. Future OCS technologies must investigate faster switching solutions if the full benefits of dynamic topology optimization are to be realized in distributed AI/ML applications.

1.4 OCS Technologies

The first technology to consider is the 3D MEMS mirror technology used by Google to create its 136x136 Palomar OCS [45]. The biggest advantage of these types of switches is their low insertion loss. Switches as big as 1100x1100 have been fabricated with only 4dB of maximum insertion loss [20]. Google's switch boasts a maximum insertion loss of only 2dB. The light itself only needs to bounce off two mirrors while making its way from input to output, reducing any loss associated with coupling into and out of waveguides as required by integrated photonic approaches. Low loss coupled with the 3D MEMS mirrors' inherent polarization insensitivity and broadband operation makes this technology a clear choice for the current OCS applications. Piezo-based switches, an alternative beam-steering technology commercialized by Polatis, have been shown to have loss as low or lower than 3D mems, with a commercial 576x576 switch reported to have a median insertion loss of 1.4dB and a maximum loss of 3dB [16] [17]. The switching speed for both technologies is quite slow, millisecond scale, and thus these switches may not provide a viable path forward for future applications as discussed in Section 1.3. Current applications, however, do not require fast switching and thus Piezo and 3D MEMS mirror based switches represent the current state of the art for optical circuit switches. Piezo-based switches, while also fabricated using standard MEMS processes, likely are more expensive to manufacture as the assembly process is more involved than 3D MEMS [10].

Significantly faster switching speeds, from microsecond [26] to even nanosecond scale [9], have been demonstrated by numerous silicon photonic (SiPh) OCSs. In addition, SiPh switches promise a lower potential cost than 3D MEMS as their manufacturing process is even closer to a standard CMOS process. Decently large radix switches (32x32) have been demonstrated using Mach-Zehnder interferometer (MZI) switches, but high per element losses and high crosstalk ultimately limit the scalability of such switches [33]. Benes style architectures can be used to minimize the total number of switching elements, and thus limit the per element losses, but this comes at the cost of increased cross talk compared to other MZI-based architectures [7]. Additionally, the rearrangeably non-blocking operation of these architectures is inherently bad for applications requiring frequent topology changes.

Micro-ring resonator (MRR) based SiPh switches generally experience the same fundamental scalability issue as MZI based switches as each additional ring adds significant loss in both the off and on states. Only switches up to 8x8 have been demonstrated thus far [13]. These types of switches also do not easily support wavelength division multiplexing (WDM) operation, an essential technology present in all modern datacenters [23]. The high thermo-optic coefficient in silicon also requires complex and power-hungry control for MRR switches. Recent work on multi-layer Si+SiN MRR OCSs solves many of these problems and shows promising scalability [25], but still does not provide pure WDM support with predefined wavelength channels limiting the type of transceivers that can be used.

There have been some promising results from indium phosphide (InP) based OCSs, as the presence of a direct bandgap material allows for the manufacture of semiconductor optical amplifiers (SOAs) on the same chip. This has allowed manufacture and testing of numerous

lossless (after amplification) switches. Unfortunately, III-V semiconductor processes such as InP are nowhere near as mature as CMOS processes, thus offering a more expensive and less economically scalable solution than MEMS or silicon photonics. In addition, the size of passive optical structures is much larger in InP processes meaning that the maximum port size of such switches is more limited than in other smaller photonic processes (such as SiPh) [5]. Hybrid architectures have been proposed to leverage the strengths of both SiPh and III-V, but port counts beyond 8x8 have yet to be demonstrated with larger switches only being proposed as a connection of multiple chips [49] [48] [43]. The SOAs themselves also consume a huge amount of power and degrade the signal-to-noise ratio through amplified spontaneous emission (ASE).

Crossbar style SiPh MEMS based OCSs, such as those demonstrated in [38] [22] [37], offer sub-microsecond switching speeds without the fundamental insertion loss and WDM issues of MZI and MRR based switches. This is due to the near zero off-state loss of the SiPh MEMS switch element and the fact that only 1 on-state switch element must be traversed per input. This means the total on-chip loss is limited only by the number of low-loss (0.015dB) waveguide crossings and the propagation loss in silicon ($<1.1\text{dB/cm}$) [37]. Recent work has also demonstrated a way to completely eliminate the waveguide crossings [12] such that the total loss is limited only by waveguide propagation loss. The fast speed and potential for low total insertion loss suggests that SiPh MEMS is an attractive solution for future OCS applications. Unfortunately, the crossbar style OCS requires k^2 MEMS devices to be individually controlled for a $k \times k$ switch. At large port counts, the total number of devices becomes prohibitively large to control via off-chip drivers. Row/column addressing schemes have been shown to alleviate this issue, but at the cost of greatly increased power consumption and switching time [22]. To make matters worse, each SiPh MEMS device requires actuation voltages in excess of 40V. These control challenges have been cited as one of the biggest drawbacks to this type of OCS [8]. The remainder of this thesis explores two novel approaches to solving this control challenge to enable fast, low-power, and scalable crossbar style SiPh MEMS OCSs.

Chapter 2

SuperSwitch 1: A Digital Approach

2.1 Problem Statement

The goal of SuperSwitch 1 was to build a high voltage CMOS controller designed specifically for digital SiPh MEMS switches such as those described in [22]. As discussed in the previous chapter, compared to other SiPh switches, these switches have low loss and fast reconfiguration times ($< 1\mu\text{s}$). The remainder of this chapter will highlight how the SuperSwitch 1 CMOS controller was designed to complement the SiPh switch such that its fast speed and low loss can be fully realized.

2-Layer SiPh MEMS Switch

The MEMS devices are based on a 2-layer adiabatic coupler architecture that has been researched in great detail in [14, 22]. Figures 2.1 (a) and (b) illustrate the OFF and ON state of the basic MEMS coupler device. The bottom waveguide layer is used as the bus layer. Fig 2.1 shows how this basic MEMS coupler can be turned into a crossbar switch element. The top coupler layer is pulled in towards the bus waveguide when a high voltage is applied across the device. Light is adiabatically coupled from the input bus waveguide (left side of Fig. 2.1 (c)) and is coupled up into the coupler layer. The coupler layer then bends 90 degrees before coupling light back into the bus waveguide layer, but now in the vertically oriented direction. This 90 degree bend is what allows this device to act as a crossbar switch element, coupling light from the horizontal input waveguides to the vertical output waveguides. When the MEMS device is in the OFF state the light remains in the horizontal input waveguide and continues on through the multi-mode interference (MMI) crossing. The MMIs are placed at all intersections of horizontal and vertical waveguides and help reduce the total on-chip loss.




Figure 2.1: (a) Schematic of (a) OFF state and (b) ON state of vertical MEMS coupler - figures taken from [14] (c) SEM of MEMS switch cell - figure taken from [22]

Digital High Voltage MEMS Control

A representative transfer function of the MEMS device is shown in Fig 2.2. There are two things to note about this transfer function: 1. the voltage required to operate the device is in excess of 40 V and 2. there exists a hysteresis loop of around 10 V. The hysteresis in the transfer function is due to the fact that the MEMS device makes physical contact with a mechanical stopper, highlighted in green in Fig 2.1, when the applied voltage is greater than around 40 V. Once contact is made, surface forces, such as Van der Waals force, act in addition to the electrostatic force. The device retracts, or pulls-out, only when the applied voltage is lowered enough such that the elastic force of the MEMS device can overcome the combined contact and electrostatic forces.

The transfer function gives us a first order understanding of the type of high voltage driver that needs to be designed for this switch. Obviously the driver needs to operate in excess of 40 V in order to fully pull-in the device. It can also be seen that the optical power remains relatively constant once the voltage applied to the device is above a certain value. This suggests that as long as the driver can apply a voltage higher than the pull-in voltage, and lower than the pull-out voltage, that it can operate in a digital manner. This is very important as a digital high-voltage driver is much easier to design and burns much less power than a high voltage driver capable of outputting analog values.

Scaling the OCS Radix

Given a standard fiber array unit pitch of 127 μm, a 128x128 radix crossbar SiPh MEMS switch fits comfortably within the standard reticle size of 26mmx33mm. While reticle stitching has been shown to allow for scaling SiPh MEMS switches to port counts of 240 [37] (comparable to 256-port state-of-the-art packet switches [6]), SuperSwitch 1's radix was limited to 128 for an initial demonstration of co-packaged CMOS-SiPh operation. This radix choice was further validated when Google published on their own OCS that has a similar radix of 136 [45].

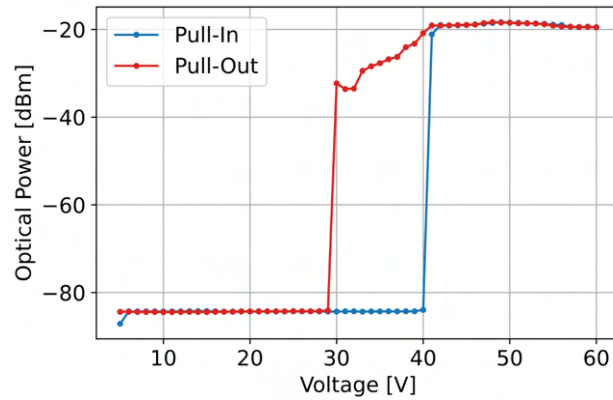


Figure 2.2: A representative transfer function of a MEMS switch element.

The 128-radix SiPh switch is shown in Fig. 2.3 (a). This switch was made in the same style as [22] but with individual connections to each MEMS switch element exposed throughout the center of the chip. This corresponds to 16,384 switch elements each w/ a positive and negative terminal that must be connected through the CMOS control chip. Fig. 2.3 (b) shows how a 4x4 grid of CMOS chiplets could be used to cover 16 separate 32x32 switch sections. This allows us to design the CMOS control chiplet through a multi-project wafer and use multiple chips to cover the nearly reticle-size SiPh chip.

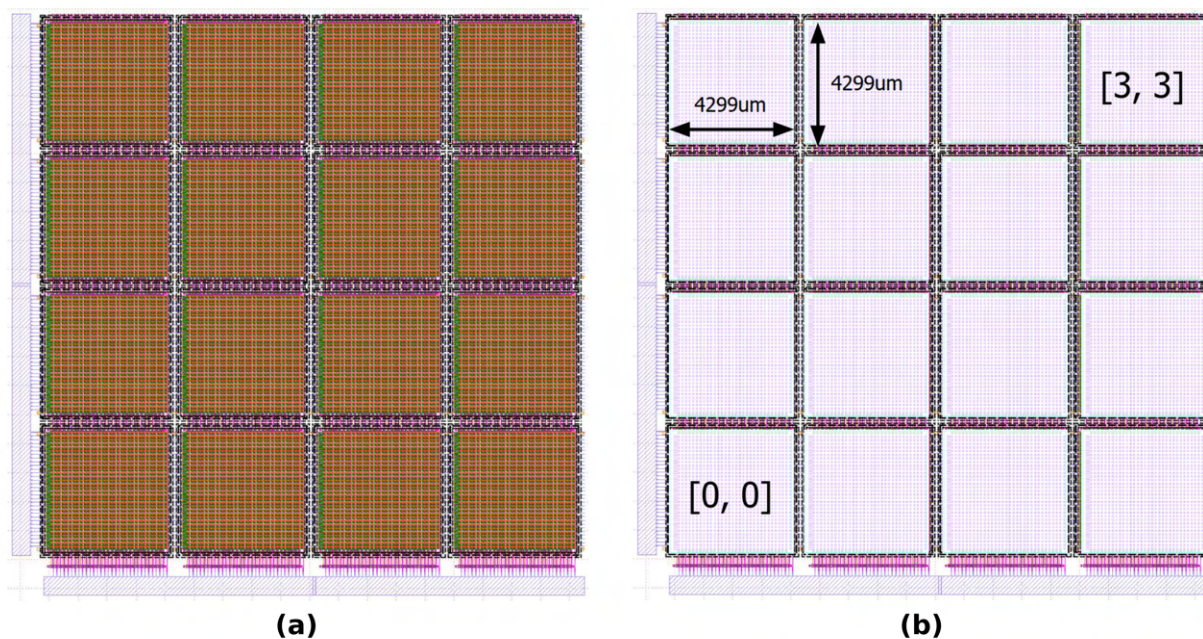


Figure 2.3: GDS screenshots of (a) 128x128 SiPh MEMS crossbar switch (b) 4x4 CMOS high voltage driver chiplets flip-chip bonded to SiPh MEMS chip.

2.2 CMOS Design & Simulation

Fig. 2.4 (a) shows a micrograph of the finished chip which was taped out in TSMC's 180nm High Voltage BCD Gen2 process. At its core, the SuperSwitch 1 control chiplet is a 32x32 array of digital high voltage level shifters. A simple floorplan of the chip is shown in Fig. 2.4 (b). The MEMS electrode and back-bias pads, shown as MEMS (+) and (-) in Fig. 2.4 (a), are pitch matched to those on the SiPh chip.

In order to connect the CMOS chiplets within each chiplet column, the power and digital control signals must be supplied from the north side of the chip and forwarded to the south side of the chip. For example, in Fig. 2.3 (b) the lower left CMOS chiplet is labeled as [0, 0], corresponding to the 0th chiplet row and 0th chiplet column. The northern most chiplet in the same column, chiplet [3, 0], will receive the digital signals and power and ground from the wirebonded connections on the north side of the SiPh chip. These signals will then be forwarded down all rows within the 0th column until they reach chiplet [0, 0]. The SiPh chip has a top routing layer that connects the south side of one CMOS chiplet and the north side of another. Each chiplet column has its own independent set of wirebond pads on the north side of the SiPh chip. The impact of forwarding the supplies from chip-to-chip is explored later in this chapter.

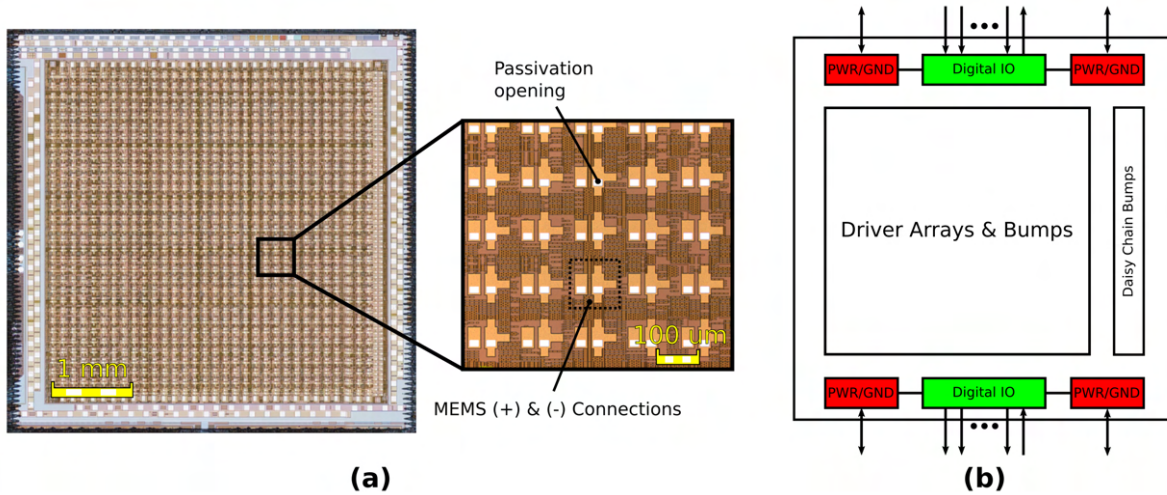


Figure 2.4: (a) Micrograph of SuperSwitch 1 high voltage driver chiplet (b) Cartoon floorplan of driver chiplet.

Digital Architecture

The crossbar switch requires that only one switching element per column be actuated. Given this, the simplest digital control scheme would be to allocate a 7-bit scan chain per column to select 1 of 128 rows. Fig. 2.5 illustrates how this simple scheme would work for any given column X . The scan chain contains a 7-bit "instruction" that the decoder would use to produce a 1-hot 128-bit value. This 128-bit value then serves as the inputs to the 128 high voltage drivers (labelled as HVDRIVER in Fig. 2.5) within that column. The drivers are connected to the pads corresponding to the positive terminals of the MEMS switching elements.

Unfortunately, the scheme presented in Fig. 2.5 does not work as we have to control a 128-radix switch using 32-radix CMOS chiplets. Each column still has a 7-bit instruction scan chain but only 5 bits are decoded as there are only $2^5 = 32$ rows and columns within each chiplet (this is highlighted in red in Fig. 2.6). The remaining 2 bits of the instruction are checked against 2 address bits loaded into a separate "address" scan chain (shown on the right side of the image highlighted in blue in Fig. 2.6). The decoder will produce a 1-hot output only if the instruction address is equal to the loaded address, otherwise it will output all 0s. This allows us to program the chiplet "addresses" at power-up such that every chiplet within a chiplet column has a unique address. Then, to actuate a specific row, the 2-bit chiplet address would be included in the instruction along with the 5-bits to denote the row within that selected chiplet.

One important thing to note about Fig. 2.6, is that the address scan chain is connected serially between chiplets in a single chiplet column. The instruction scan chains are connected in parallel across all chiplets within a column. This is done by forwarding the instruction

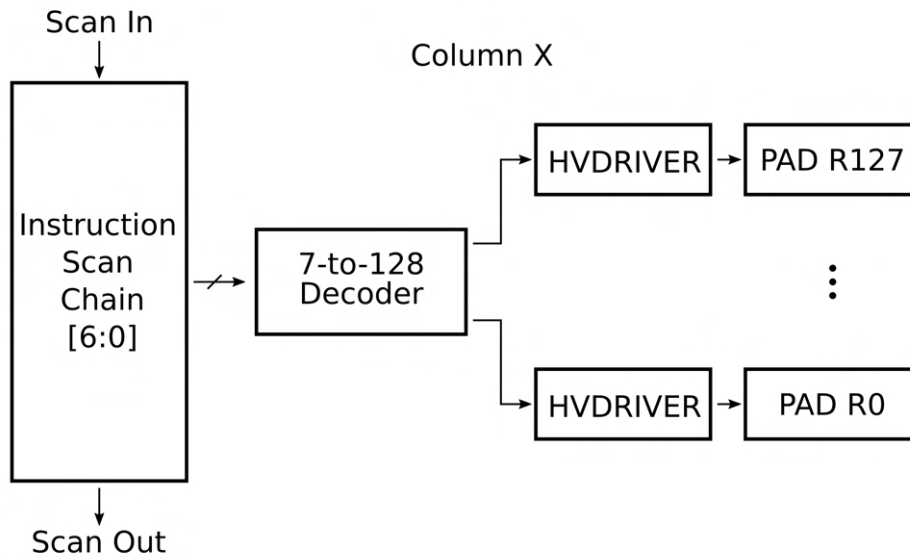


Figure 2.5: Simple schematic for activating 1 of 128 rows assuming a single CMOS chiplet.

scan chain input via the south side of the CMOS chiplet. In contrast, the address scan chain *output* is forwarded via the south side of the CMOS such that each address scan chain of each chiplet can be programmed independently.

Column Folding

The digital scan-in time directly adds to the switch reconfiguration time. If the scan-in time is large, then the fast reconfiguration time of the MEMS devices no longer affords us any benefit. In order to minimize the total switch reconfiguration time, all column instructions should ideally be scanned-in in parallel. This would require a scan-chain input signal for every column in the chip. For the 32x32 chiplet, this would mean dedicating 32 IOs just for the instruction scan chain inputs (not to mention the other signals required for operating the scan chains). To decrease the total number of scan chains, the column instructions can be "folded" such that multiple column instructions are contained w/in a single instruction scan chain. This "folding factor", or number of columns per scan chain, is denoted by the symbol N_c . Fig. 2.7 shows what the scan chains would look like for (a) $N_c = 1$ and (b) $N_c = 2$. Naturally, the scan chain increases in length by 7, i.e. $\log_2 128$, for each additional column contained within the scan chain.

Eq. 2.1 is used to calculate the total number of instruction scan chain pads, N_p , given the total switch radix, k , and the number of columns per instruction scan chain, N_c . The factor of 4 in the denominator comes from the fact that 128 columns are split across 4 chiplet columns.

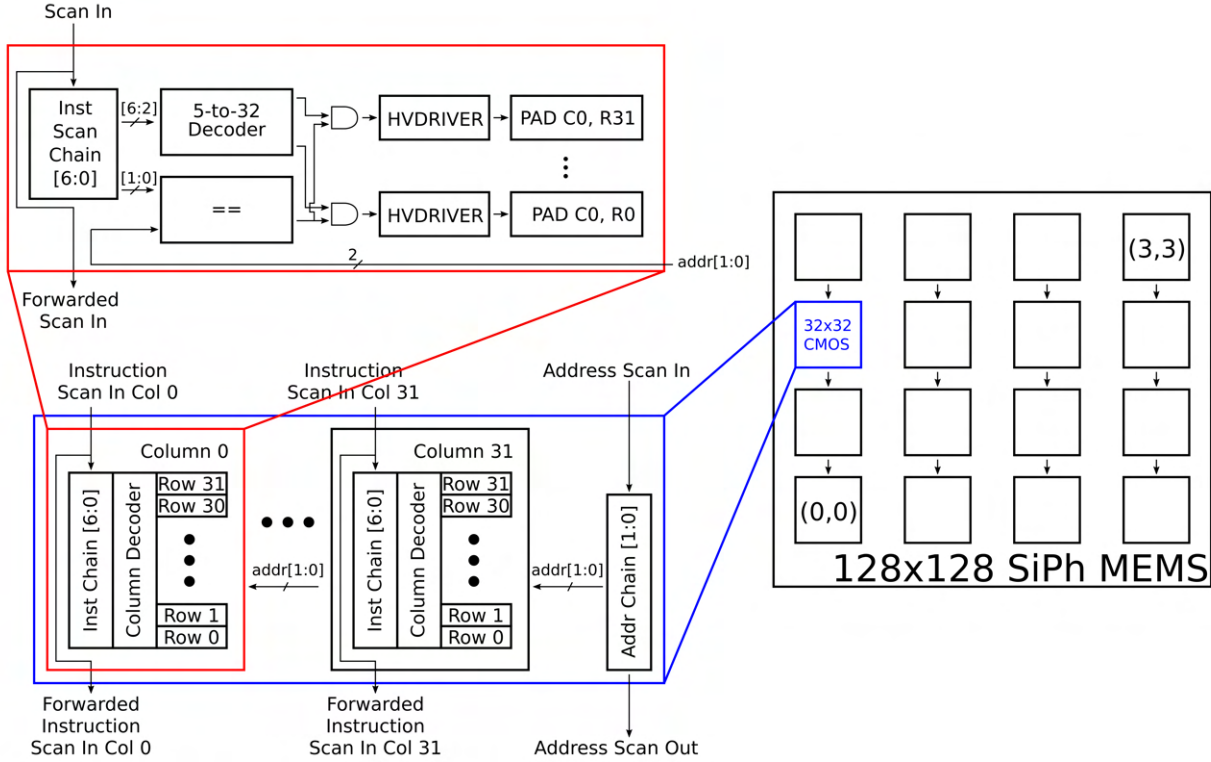


Figure 2.6: Schematic for controlling a 128x128 switch assuming a 4x4 array of CMOS chiplets.

$$N_p = \frac{k}{4N_c} \quad (2.1)$$

Eq. 2.2 is used to calculate the additional reconfiguration time due to the scan chains, T_s , given the number of columns per scan chain, N_c , the radix of the switch, k , and the clock frequency of the scan chains, f_{clk} .

$$T_s = \frac{N_c \log_2 k}{f_{clk}} \quad (2.2)$$

Final Scan-Chain Architecture

Unfortunately, 32 instruction chain inputs, i.e. $N_c = 1$ and $N_p = 32$, for a chiplet w/ north side edge of approximately 4.3mm is prohibitively high. This is doubly true since the SiPh chip must route every signal from one chiplet to another over a single routing layer. For this reason, an N_c of 4 was chosen to reduce N_p to 8. The effects of this choice on the scan chain architecture are shown in Fig. 2.8 (a). The CMOS chiplet has a total of 8 instruction scan chains each with 28 bits for controlling 4 separate columns. The address scan chain is

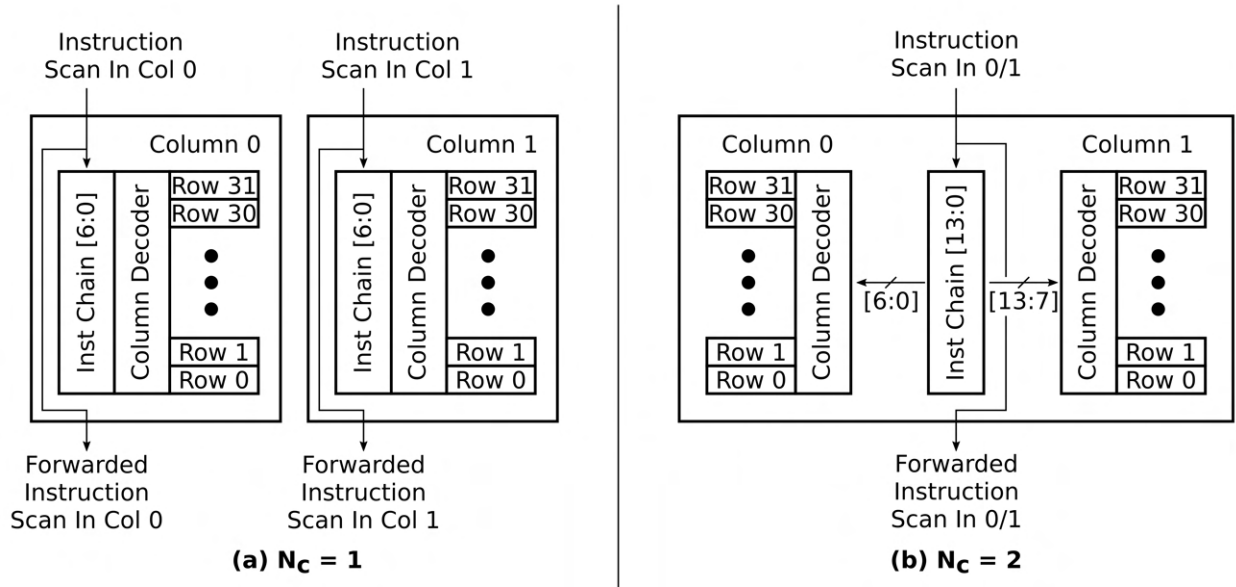


Figure 2.7: (a) Logic for columns 0 and 1 for $N_c = 1$ (b) Logic for columns 0 and 1 for $N_c = 2$.

comprised of 7 bits with only 2 of those bits actually dedicated for the chiplet address. The other bits include a global enable signal for the high voltage driver circuits, discussed later in this chapter, as well as a loopback-select field that controls the mux shown in Fig. 2.8 (a) (these fields are described in detail in Table 2.1). The loopback select allows us to test the functionality of the instruction scan chains. The mux can be configured to loopback any of the instruction scan chain output OR the loopback signal from the adjacent southern CMOS chiplet. This allows us to test any instruction scan chain within any given CMOS chiplet.

At an f_{clk} of 30MHz, this gives us a T_s of less than 1 μ s. Seeing as how similar style MEMS switches have observed switching times less than 1 μ s ([14, 22, 37]), this seems an appropriate target for T_s . A list of all final design parameters can be found in Fig. 2.8 (b).

HV Driver Architecture

The schematic of the high voltage driver is shown in Fig. 2.9 (a). The various supplies and their nominal values are listed in Fig. 2.9 (b). This circuit is a modified version of a capacitive level shifter first introduced in [41] and then further detailed in [50]. The two digital inputs, V_{in} and V_{en} , are generated by the logic shown in Figures 2.6 and 2.8. V_{in} controls the output at V_{out} such that a 0 to 1.8 V swing at the input causes a 0 to 70 V swing at the output. The input signal is first level shifted to 5 V, via transistors M_1 - M_4 , before being used to push/pull the voltage at the bottom plates of capacitors C_1 and C_2 . Via capacitive coupling, a changing voltage at the bottom plates of these capacitors will cause a corresponding change in voltage

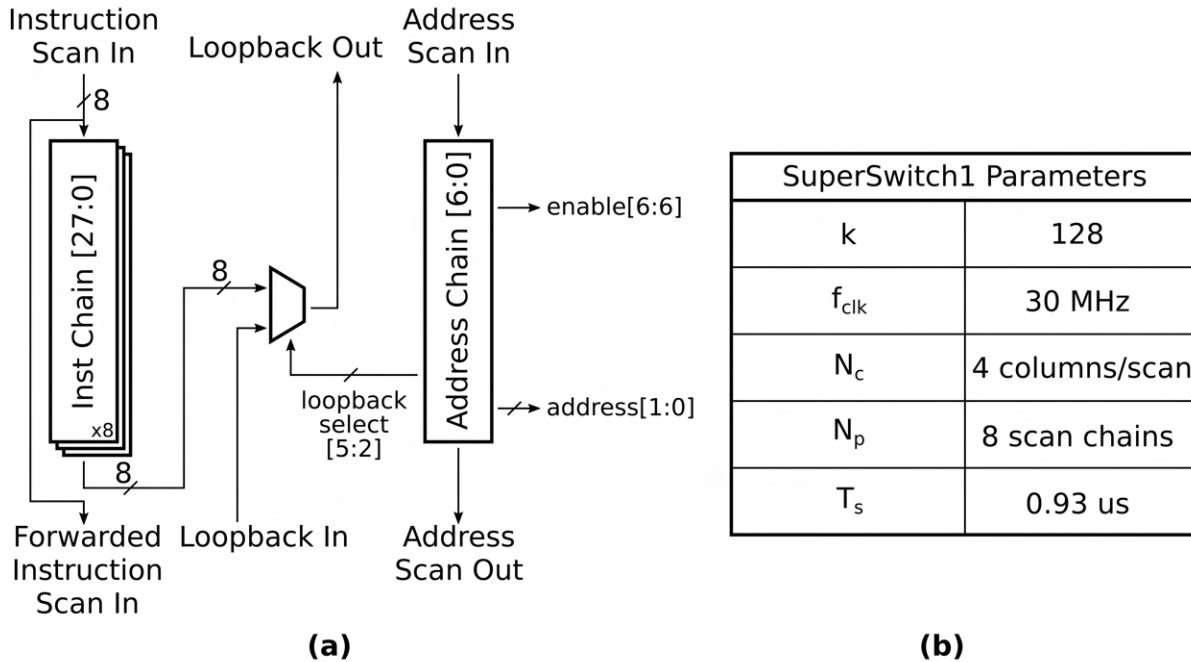


Figure 2.8: (a) Final schematic of SuperSwitch1 control chiplet scan architecture w/ loopback mux for debug. (b) Final parameters for SuperSwitch1 controller chiplet.

Table 2.1: Global "address" scan chain bits listed in the order that they are connected.

Field	Width	Description
driver_enable	1	Global (at a chiplet level) high voltage driver enable signal.
loopback_select	4	Selects the loopback signal for debug purposes. Values 0-7 select the output of the 1st through 8th instruction scan chain. Any value above 7 selects the loopback input from the connected chiplet to the south.
chiplet_address	2	Sets the address of the chiplet for decoding instructions. This field should be set to a unique value for each chiplet within a chiplet column.

at nodes X and Y. Assuming X and Y are pushed above/below the switching threshold of the latch made up by transistors M_9 - M_{12} , then positive feedback from the latch will continue pushing/pulling X and Y to HVDD and HVSS respectively. Node Y can then be used to drive the gate of the high voltage PMOS in the output stage of the high voltage driver. M_{13} , the high voltage NMOS in the output stage, is driven by the logical AND of the inverted and level shifted V_{in} and the level shifted V_{en} . When the driver is enabled, i.e. $V_{en} = 1.8V$,

M_{13} is driven high when $V_{in} = 0V$ such that only one transistor in the output stage is ever turned on. On power-up, V_{en} is set to 0 which forces M_{13} to be off. This is desired as we do not know the initial state of the latch immediately after power up. As discussed in a following section, a specific initialization routine must be completed to place the latch in a known state before V_{en} can be set to 1.8V. Otherwise, there is a chance that both M_{13} and M_{14} could be turned on after power-up.

70 V corresponds to the highest voltage transistor offered in TSMC's 180nm HV BCD Gen2 process. These transistors were chosen for M_{13} and M_{14} as their area was not prohibitively large and the 70 V limit gives ample margin for error given the required 40 V range estimated by Fig. 2.2.

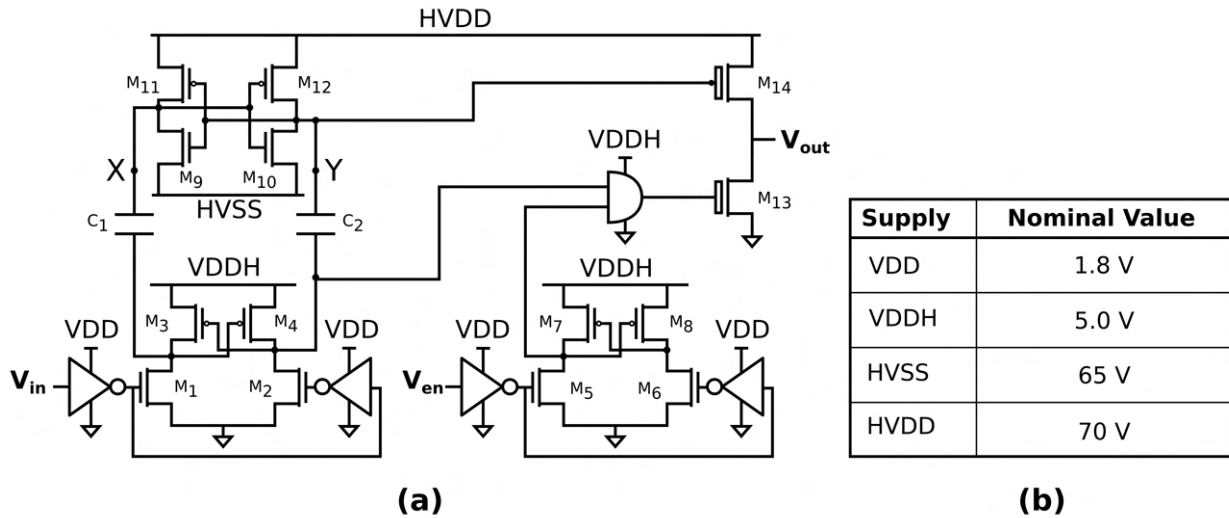


Figure 2.9: (a) Schematic of SuperSwitch1 high voltage driver circuit. (b) List of all supplies and their nominal values.

Optimizing Bootstrap Capacitance

Perhaps the most important parameter when sizing the circuit in Fig. 2.9 is the size of the bootstrap capacitors C_1 and C_2 . The larger we make these capacitors the better capacitive coupling we will have from the drains of M_1 and M_2 to nodes X and Y. The primary failure mechanism of this circuit, discussed in detail in [50], is the failure to flip the latch when transitioning V_{in} . This can occur if the capacitances, C_1 and C_2 , or the pullup/pulldown resistances of the latch itself, become too small. When either of these values become too small then any change in voltage on nodes X and Y will be quickly neutralized through the pullup/pulldown paths in the latch. The capacitances and resistances must be large enough such that the voltages on nodes X and Y can increase/decrease beyond the switching

threshold of the latch. This is the only way that we can get the latch to change states when toggling V_{in} .

Given a standard 2:1 sizing of PMOS-to-NMOS widths (w/ minimum sized NMOS), the key question becomes rather simple: how large should we make the bootstrap capacitors? Fig. 2.10 shows the results of transient simulations as the bootstrap capacitance is swept from 10fF to 100fF. The goal is to find a capacitance value that is small enough that it does not dominate the area of the driver but large enough that switching is guaranteed across all corners. The corners, shown on the y-axis, denote the process corner (NMOS strength followed by PMOS strength) and the temperature at which the simulation is run (m40c referring to *minus* 40 degrees celsius). Failure occurs when the driver is not able to flip the state of the latch and thus the output cannot be driven between HVDD (70 V) and VSS (0 V).

As shown very clearly by Fig 2.10, only the mismatched corners, $fs27c$ and $sf27c$, seem to affect the performance of the circuit relative to the nominal corner $tt27c$. For this particular circuit, the worst performance occurs when the NMOS process corner is slow and the PMOS process corner is fast. This is due to the operational mechanism of the input level shifter (transistors M_1 - M_4). As discussed in [50], the rate of change in voltage on the bottom plates of C_1 and C_2 is directly related to the change in voltage seen on nodes X and Y. The rate of change of voltage is governed by the speed at which the input level shifter shifts. Because the inputs of the level shifter, transistors M_1 and M_2 , are NMOS, the speed of shifting depends on the strength of NMOS relative to the PMOS. For stronger NMOS, i.e. corner $fs27c$, the outputs of the input level shifter will change sharply as the NMOS quickly overpowers the PMOS. For the other corner, $sf27c$, the NMOS transistor will more slowly overcome the PMOS transistor and thus give us a smaller change in voltage on nodes X and Y. This is why $sf27c$ gives the worst circuit performance and $fs27c$ gives the best. Despite this, Fig. 2.10 tells us that a capacitance value of 30fF or above should allow us to switch under any process corner.

The capacitors themselves are made from standard MOM caps. This type of capacitor is preferred as it can withstand voltages up to 70 V without experiencing dielectric breakdown. Other more area efficient capacitors, such as MIM caps, cannot withstand high enough voltages to be used in this circuit. Using 3 metal layers, 90fF capacitors were designed for C_1 and C_2 . Fig. 2.10 shows that this size gives us ample margin across all corners. The size of each capacitor is approximately $240\mu\text{m}^2$. The total driver size is $\approx 9160\mu\text{m}^2$ which means the capacitors account for less than 6% of the total area. The high voltage transistors, on the other hand, account for greater than 50% of the total area of the high voltage driver.

Initialization Procedure

As previously stated, the latch shown in Fig. 2.9, formed by transistors M_9 - M_{12} , will initially be in an unknown state after power up. The latch must be initialized before the driver is enabled in order to ensure that a short never occurs from HVDD to VSS. The following initialization procedure must be followed to properly enable the high voltage driver circuits.

Bootstrap Capacitance vs. Corners

fs27c	Fail	Fail	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass
ss125c	Fail	Fail	Fail	Pass	Pass	Pass	Pass	Pass	Pass	Pass
tt27c	Fail	Fail	Fail	Pass	Pass	Pass	Pass	Pass	Pass	Pass
ffm40c	Fail	Fail	Fail	Pass	Pass	Pass	Pass	Pass	Pass	Pass
sf27c	Fail	Fail	Fail	Fail	Fail	Pass	Pass	Pass	Pass	Pass
	10fF	20fF	30fF	40fF	50fF	60fF	70fF	80fF	90fF	100fF

Figure 2.10: Shmoo plot for 10-100 fF bootstrap capacitance.

1. Ramp the digital voltages VDD and VDDH to their nominal values (1.8 and 5 V). All input and enable signals for all drivers should remain at 0 V.
2. Ramp both HVDD and HVSS to HVSS's nominal value (65 V).
3. Ramp HVDD to its final and nominal value (70 V).
4. Toggle V_{in} of all drivers from 0 V to 1.8 V and back to 0 V. This will place all latches into a known state.
5. Set the global enable signal to 1.8 V such that the high voltage NMOS of each driver circuit is turned on. The previous step ensured that the HVPMOS was turned off so there is no danger in turning the HVNMOS off. All drivers are now enabled.

Fig. 2.11 shows a transient simulation of the above initialization procedure. For this particular simulation, node Y of Fig. 2.9 settles to 65V after all supplies are ramped. If V_{en} signal is set to 1.8V immediately after power up then both the HVNMOS (on when $V_{in} = 0V$) and HVPMOS (on after power-up) would be on causing a short between HVDD and VSS. Fig. 2.11 shows the keeper being placed into a known state around $3.5\mu s$ once V_{in} is set to 0 V after being set to 1.8V. V_{en} is then set to 1.8V around $3.75\mu s$. V_{in} is then toggled

one more time to show that the output correctly rises to HVDD (70V) when $V_{in} = 1.8V$ and correctly falls to VSS (0V) when $V_{in} = 0V$.




Figure 2.11: Basic initialization procedure for SuperSwitch1 high voltage driver circuit.

Simultaneous Switching

The circuit model for a single column of flip-chip bonded chiplets is shown in Fig. 2.12. The supply and ground resistances between chiplets are due to the flip-chip bonds between the SiPh and CMOS chips as well as the short amount of interconnect on the SiPh chip. One concern with this design is that given a large enough supply and/or ground resistance, the chip may cease to function properly. The ground or VSS resistance is the most sensitive as all current from all supplies will return via the same ground path. For this reason, a transient simulation experiment was done to understand the maximum allowable ground resistance. The results of this experiment are shown in Fig. 2.13. During the simulation 32 high voltage drivers simultaneously turn on while 32 separate high voltage drivers turn off. This gives us the worst-case/maximum current draw that any chiplet column should ever experience. This is because of the crossbar nature of the switch and the fact that only 1 device/column ever needs to be actuated. The simulation assumes that all the current from all supplies is forced through the same resistance. This would occur only when all 32 drivers turning off and on are located in chip 0 of Fig. 2.12.

Alarmingly, Fig. 2.13 (a) indicates that only 10 ohms of ground resistance is needed to cause circuit failure at the $sf27c$ corner. This is the same worst case corner discussed in Section 2.2. In contrast to the experiment of Fig. 2.10, there also seems to be some




Figure 2.12: Equivalent circuit of supply and ground connection for a single chiplet column in the flip-chip bonded CMOS-SiPh package.

performance degradation associated with the *tt27c* and *ffm40c* corners relative to the *ss125c* and *fs27c* corners. This is likely due to the fact that the stronger the transistors are the more current that will be drawn and thus the more voltage drop across the ground resistance. Thankfully, there is a relatively simple remedy to this problem. Fig. 2.13 (b) shows the same experiment but with HVSS set to 66V instead of 65V. This slightly degrades the rise time of the high voltage driver (as the HVP MOS can now only be turned on w/ a $|V_{gs-max}|$ of 4V instead of 5V) while affording significantly improved resilience to ground resistance. This improved resilience is due to the increased HVSS reducing the switching threshold of the latch circuit inside the high voltage driver. The increased HVSS will also decrease the peak current drawn by the HVP MOS and thus reduce the voltage drop across the ground resistance. HVSS can be increased further to accomodate even larger ground resistance.

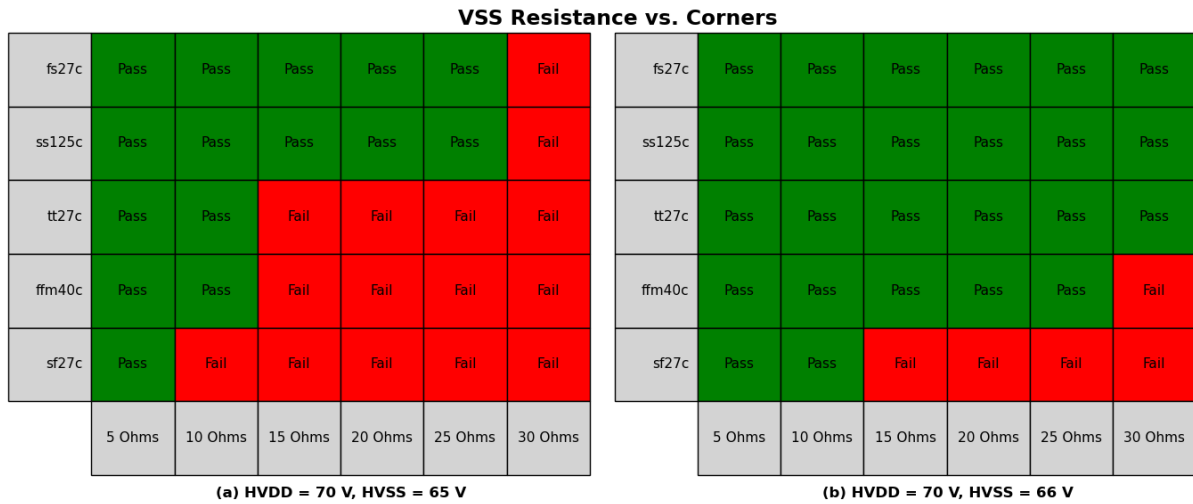


Figure 2.13: (a) VSS resistance shmoo plot for all corners for HVDD = 70 V, HVSS = 65 V. (b) Same plot but for HVDD = 70 V, HVSS = 66 V.

Transient Simulations

The simulated rise and fall times of the high voltage driver circuit, shown in Figures 2.14 (a) and (b) respectively, are well below $1\mu\text{s}$ even for 70V operation and an output load of 1pF . These simulations are done on the extracted post layout netlist with an ideal capacitor as the output load.




Figure 2.14: Rise (a) and fall (b) times extracted from post-layout simulations.

The actual load of the driver consists of both the MEMS device and the parasitic output capacitance of the driver. The MEMS device acts as a variable capacitor, having a high capacitance when pulled in and a lower capacitance when pulled out. In order to understand the effects of the variable capacitance on the driver, a Verilog-A model was developed similar to the model detailed in Chapter 2 of [40]. Fig. 2.15 (a) shows the results of a transient simulation using the MEMS model as the output load. The output voltage drops as the MEMS device pulls in and the capacitance suddenly increases from around 10fF to around 100fF . The sudden drop in voltage at the output could potentially affect the state of the latch controlling the HVPMOS (through C_{gd} of M_{14} in Fig. 2.9 (a)). Thankfully, even the minimum sized HVPMOS in this process was strong enough to limit the maximum voltage droop to less than 10%, resulting in a minimal change in voltage on Node Y of Fig. 2.9 (a). The effects of this changing capacitance are further reduced by any additional output capacitance. Fig. 2.15 (b) shows the extracted parasitic output capacitance for all 1024 drivers. Each driver has a slightly different parasitic output cap because of the way the pads are arrayed in relation to the driver array. The average parasitic capacitance is only $\approx 23\text{fF}$, suggesting the MEMS capacitor will be the majority of the capacitive load.

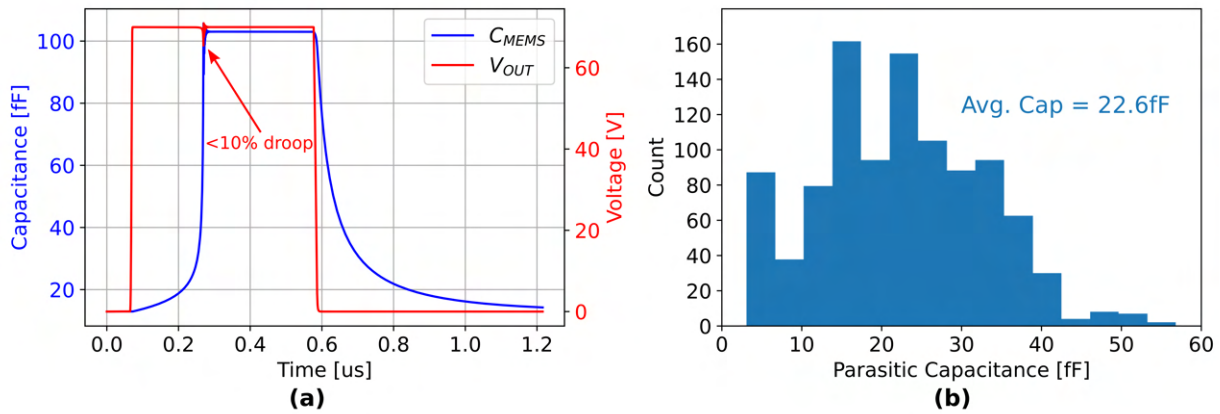


Figure 2.15: (a) Transient simulation of capacitance of MEMS model (C_{MEMS}) and output driver voltage (V_{OUT}). (b) Parasitic output capacitance for all 1024 drivers (due to short routes and pads).

2.3 3D SiPh-CMOS Packaging

The SiPh and CMOS chips were bonded together using a process that relies on deformable Au micro-bumps patterned using nano-particle deposition. The process bonds the micro-bumps together as result of mechanical pressure and does not require reflow (max temperature is 200C). This reduces the thermal cycling stress as each package must be temperature cycled 16 times, once for each CMOS chip. It also completely eliminates the need for flux. This is essential as flux or other similar substances could infiltrate the released MEMS structures rendering the switch inoperable post bonding. More details on the specifics of this bonding process are given in [21].

The post Au micro-bumped CMOS chips are shown in Fig. 2.16 (a). Two types of pads are present on the CMOS: a small $51\mu\text{m} \times 31\mu\text{m}$ pad for the switch connections and a larger $77\mu\text{m} \times 70\mu\text{m}$ pad for the north and south IO connections. Fig. 2.16 (b) shows an illustration of the cross section of a switch pad. A small passivation opening is made over the CMOS Al pad so that the deposited Au UBM can make contact. The UBM covers only the area above the CMOS pad such that there exists a flat surface for bumping and bonding (right side of Fig. 2.16 (b)). After several failed attempts at bonding, a surface height difference of nearly $1\mu\text{m}$ was discovered between the switch and IO pads on the CMOS chip. Fig. 2.16 shows how this difference was compensated for by adjusting the thickness of the SiPh chip's UBM for the two pad types.

A 4x4 test package is shown in Fig. 2.17 (a). 16 CMOS chips were bonded to an electrical interposer as an early first attempt at bonding. Fig. 2.17 (b) shows a single CMOS chip bonded to a 32x32 subsection of the 128x128 SiPh OCS. This was the highest yielding package and was used for generating all results in Section 2.6. A fully yielding 128x128

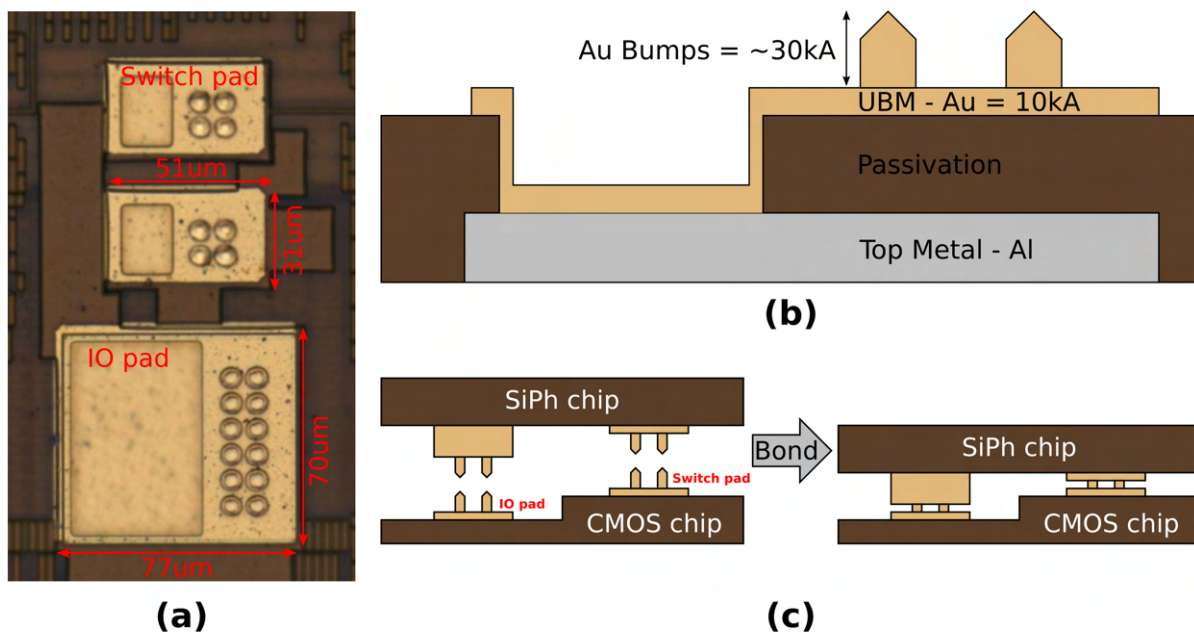


Figure 2.16: (a) Micrograph of bumped CMOS pads. (b) Cross section of Au UBM and Au micro-bumps. (c) Illustration of bonding process using differing thicknesses of UBM on SiPh chip to compensate for CMOS pad height differences.

package has not yet been accomplished and is part of the future work for this project. The resistance of all supplies through the 32x32 switch was measured to be below 10Ω with the VSS resistance under 2Ω , well under the maximum resistance for all corners as shown in Fig. 2.13 (b). [21] measures an average resistance of $0.25\Omega/\text{bump}$ suggesting that the bumps are likely not the limiting factor.




Figure 2.17: (a) 4x4 CMOS flip-chip bonded to electrical interposer. (b) 1 CMOS flip-chip bonded to 32x32 subsection of 128x128 OCS.

2.4 PCB Design

The PCB design was split into two different boards. The first board, shown in Fig. 2.18 (a), serves as the wirebond package for the flip-chip bonded SiPh-CMOS assembly. Wirebonds are made from the north side of the SiPh chip directly to this "chip board" PCB. The chip board then fans out the signals and connects them to a set of pins that fit into a standard 17x17 pin-grid-array (PGA) socket. A dust cover is also fit on top of the chip board to protect the released MEMS structures from getting contaminated. Fig. 2.18 (a) shows a picture of a chip board w/ only a single CMOS chiplet bonded to the upper lefthand corner of the SiPh chip.

The second board, shown in Fig. 2.18 (b), or "host board", contains the PGA socket that the chip board plugs in to. The host board then contains all ICs required to test the chip. These ICs include power regulators, power monitors, level shifters, buffers, and other miscellaneous ICs. The separation between the chip and the host board allows reuse of the host board across all assembled chip boards.

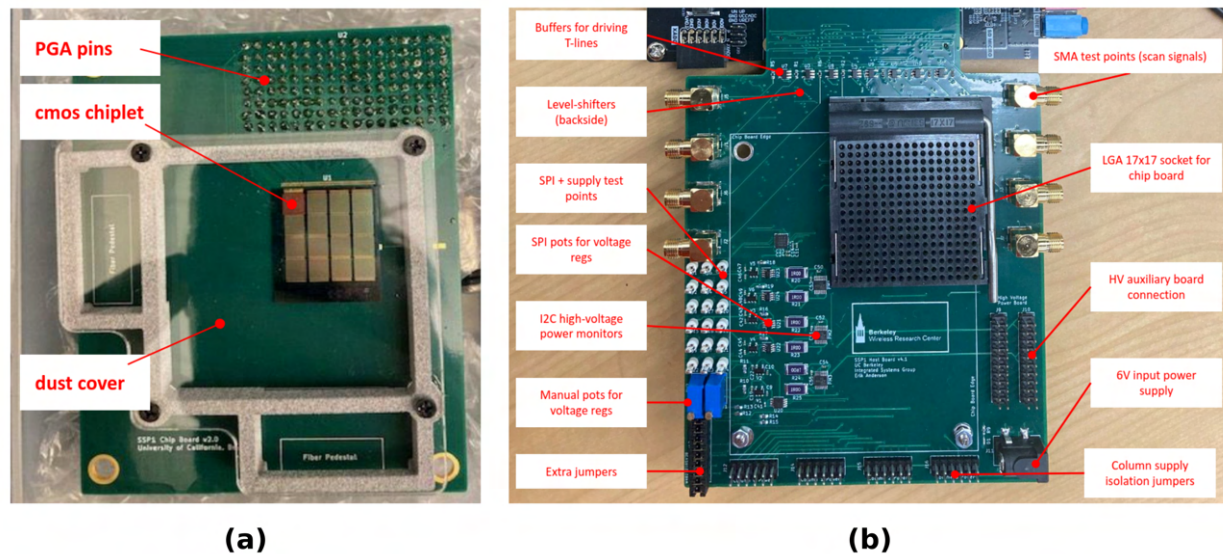


Figure 2.18: (a) The chip board for wirebonding chips to. (b) The host board PCB used for testing chip boards.

2.5 Electro-optic Characterization Setup

The first step to setting up a full electro-optic characterization setup is establishing electrical communication with the chip itself. As the CMOS chips use a custom scan chain architecture (detailed in Section 2.2), an FPGA provides a simple and reliable way to generate compatible scan controllers. Fig. 2.19 shows how the host PC communicates to the chip through the scan controllers instantiated on the FPGA. First, the PC sends UART commands (via pySerial) to read and write registers connected to a custom UART implementation. These registers set and get values inside a finite-state-machine that controls N scan chain controllers. N can be more or less depending on the number of CMOS chiplet columns bonded to the SiPh chip. For example, Fig. 2.18 (a) shows a single chiplet row in a single chiplet column. An accompanying FPGA image would have $N = 9$ scan chain controllers to control the 8 instruction scan chains and the 1 additional address scan chain. Each additional chiplet column would require an additional 9 scan chain controllers.

The FSM can directly control the scan controllers through the UART interface OR it can be programmed to continuously scan-in bits written into the block-RAM of the FPGA. The latter mode is required in order to measure the worst-case power consumption of the device. Without this continuous mode, the scan chains become limited by the speed of the UART interface which can be quite slow. This will be discussed in more detail in a later section.

Once electrical communication between the OCS and the host PC has been established, the rest of the characterization setup can be constructed. A cartoon illustration of the final characterization setup can be seen in Fig. 2.20. The host PC is not shown in this figure. In




Figure 2.19: Simplified diagram of the PC-to-chip interface with UART-controlled scan chain controllers.

reality, the host PC is connected to every device in the figure besides the few passive optical components such as the 99:1 splitter and the 50:50 splitter. The input light is provided by a tunable O-band laser. This light is then fed into a motorized polarization controller before being sent to the switch. This is required as the switch uses vertical grating couplers optimized for TE polarization of light. The polarization controller helps optimize the input polarization such that the minimum loss can be measured. After the polarization controller, the light is split and 1% is sent to a digital power meter and the other 99% is sent to a 1x64 OCS. This OCS is used to switch the light into any of 64 input ports of a fiber array unit (FAU). The FAU is responsible for shooting the light into the SiPh OCS, shown in yellow in Fig. 2.20. The reading from the digital power meter at the end of the 1% tap can be used to back calculate the input power to the switch. The FAU is initially aligned by hand but is continuously fine-aligned by the motorized positioners during the characterization routine. This compensates for any drift experienced by the FAU as well as allowing for optimized coupling for each individual input and output. While the FAU's fibers are nominally pitch matched to the SiPh OCS at $127\mu\text{m}$, there is some nonuniformity in the fiber spacing. By fine-aligning the FAUs for each input to the SiPh OCS, the minimum loss through each input can be observed. A similar coarse and fine alignment is done on the output FAU as well. There is also another 64x1 OCS to select a specific output for measurement. The light from the selected output is then fed into a 50:50 splitter which sends half the light to another digital power meter and the other half to an oscilloscope with an optical input. The power meter gives us a low noise-floor measurement of the output power. This measurement, combined with the input power measurement, allows us to calculate the loss through any given switch element in the SiPh OCS. The oscilloscope is used to measure the rise and fall time of the optical signal.

An annotated picture of the actual electro-optic characterization setup is shown in Fig. 2.21. For full 128x128 characterization, a total of 16,384 devices must be characterized. This is a massive number of devices and requires a reliable and repeatable measurement setup. Almost everything in the measurement setup is fully automated including the oscilloscope waveform capture, the polarization optimization algorithm, and the fine-alignment of the FAUs. The external 1x64 OCSs limit the automation to only $64^2 = 4096$ devices. After characterizing 4096 devices the FAUs either need to be realigned to a new set of 64

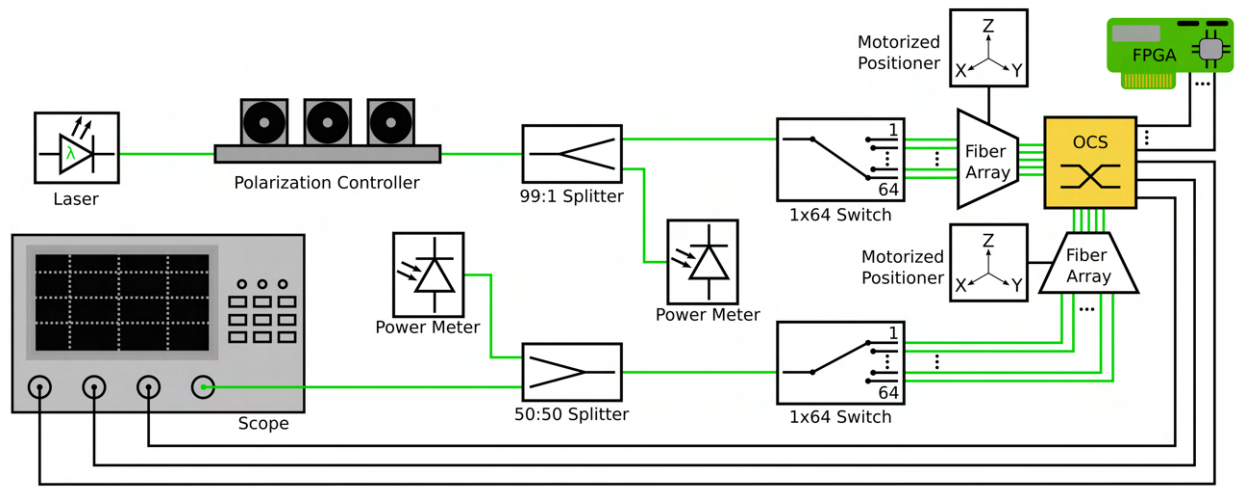


Figure 2.20: Illustration of the SuperSwitch1 characterization test setup.

input/output ports or the patch panels (shown in lower left and lower right of Fig.2.21) must be rewired. A more complete setup would make use of 1x128 OCSs such that coarse alignment must only be done once at the beginning of the characterization routine.

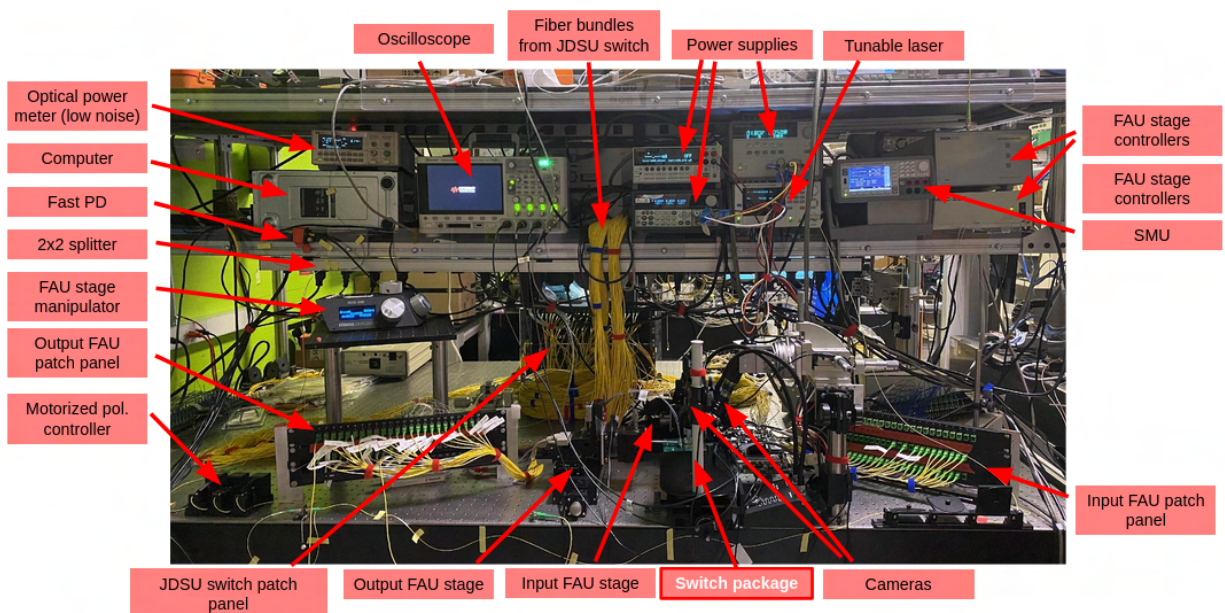


Figure 2.21: The physical test setup for fully automated switch characterization.

2.6 Results

As of writing this thesis, the only complete characterization that has been done was on the 32x32 OCS pictured in Fig. 2.17 (b). This figure also shows the package with both the input and output FAUs aligned to the switch’s vertical grating couplers. Only 1 CMOS chiplet is bonded to the 128x128 SiPh switch and thus the radix is reduced to 32. Looking back at Fig. 2.3 (b) the bonded CMOS chiplet has index [3, 0] (chiplet column 0 row 3).

Single Cell Characterization

The switching speed, optical loss, and optical transfer curve were measured for each of the 1024 MEMS switch elements. Fig. 2.22 shows these characterization results for the lower left switch element (row 96, column 0) of the 32x32 OCS. The switch rise and fall times are extracted from captured oscilloscope waveforms such as the one shown in Fig. 2.22 (b). The top 3 signals are the instruction scan chain control signals for the CMOS chip (scan clock signals are not shown). The scan enable signal enables shifting of bits specified via the scan-in signal. The scan update signal latches the new shifted-in value of the scan chain into the design. This means the high voltage driver begins driving the switch element as soon as the scan update signal is asserted. The bottom signal shown in Fig. 2.22 (b) is the optical output from the switch (translated into a voltage through a PD+TIA). The rise time is calculated from the instant scan update is asserted to the time at which the optical signal has reached 90% of its final value. Fig. 2.22 (c) shows the rise and fall times for different values of HVDD. At 60V, both rise and fall are less than 1 μ s, 0.83 μ s and 0.46 μ s respectively. The rise time shows a strong dependence on HVDD as larger HVDD increases the maximum electrostatic force that pulls the device in, thus pulling it in quicker. The fall time shows the opposite trend as increased HVDD simply means the driver takes longer to discharge the MEMS capacitor to the point at which the combined surface and electrostatic forces are less than the elastic force of the MEMS. Once the surface forces are no longer acting on the device, the elastic force of the MEMS structure quickly overpowers the relatively weak electrostatic force. Fig. 2.22 (d) shows the optical transfer curve for this device.

32x32 OCS Results

The full characterization results for the 32x32 OCS are shown in Fig. 2.23. Figures 2.23 (a)-(c) show the on-chip loss, rise time, and fall time for each switch element at an HVDD of 60V and a wavelength of 1310nm. Only 3 switch elements, shown in white, were found to be non-operational. This corresponds to a yield of 99.7%. The on-chip loss distribution is shown in Fig. 2.23 (d). This switch includes excess loss of \approx 3.8dB as the output ports of the 32x32 OCS must propagate through an additional 96 rows of unused switch elements before reaching the output vertical grating couplers. This could be avoided by designing (or dicing) the SiPh OCS to only have a radix of 32. As shown by Fig. 2.23 (e), the rise and fall times for all elements are below 1 μ s for an HVDD of 60V.

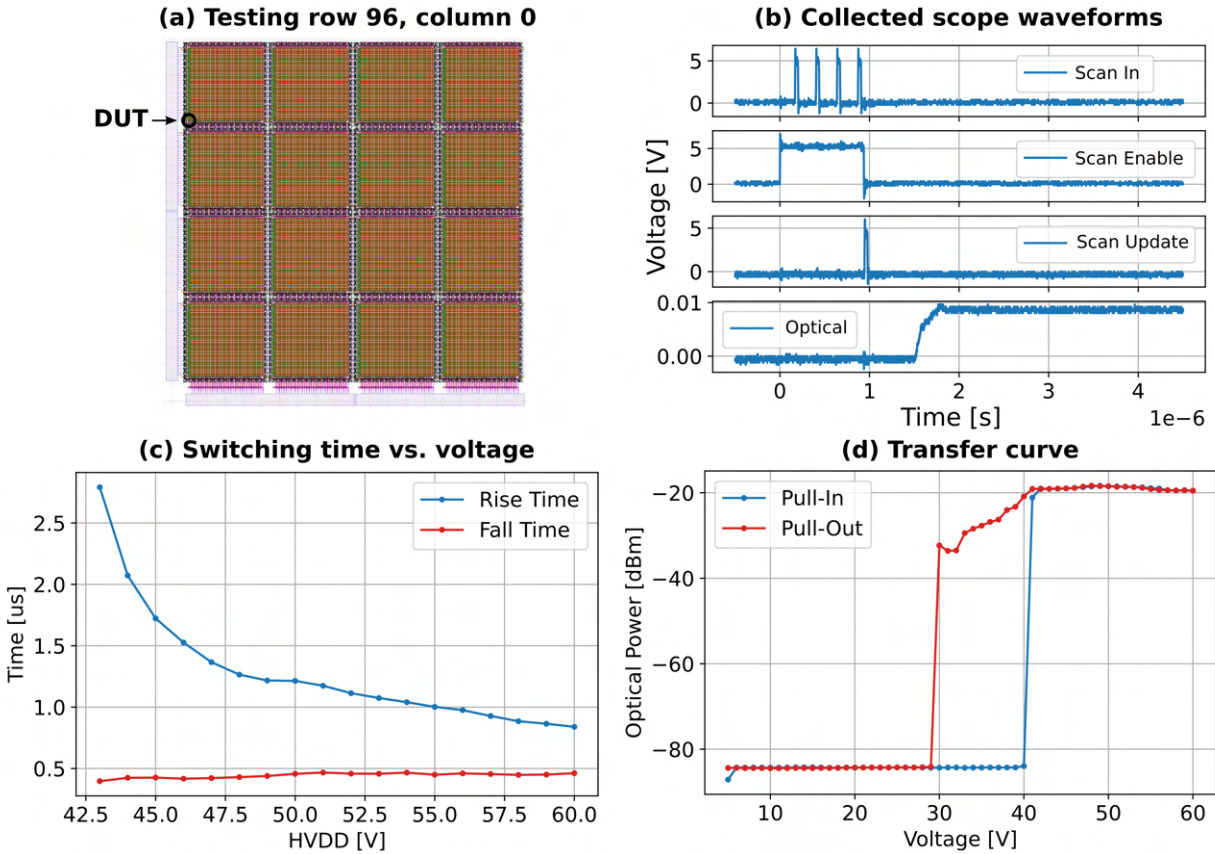


Figure 2.22: (a) Location of switch cell being tested. (b) Captured scope waveforms used for extracting rise and fall times. (c) Rise and fall times. vs applied voltage. (d) Optical transfer curve.

Power Consumption

The digital nature of the driver circuit shown in Fig. 2.9 (a) results in nearly 0 static power consumption. This leaves only the dynamic power consumption which can be measured by periodically turning off and on switch elements. The worst case dynamic power consumption is illustrated in Fig. 2.24. Fig. 2.24 (a) shows a cartoon representation of a 4x4 switch. The crossbar switch requires only 1 switch element to be ON, shown in green, in each row/column. This means that for any valid switch configuration a total of k switch elements are ON, k being the radix of the switch. The worst case dynamic power would then be observed when k new elements are turned ON and the previous k elements are turned OFF. Fig. 2.24 (b) shows the worst case pattern for a 4x4 switch.

A similar pattern to the one shown in Fig. 2.24 (b) is used to measure the power for the 32x32 OCS. In contrast to Fig. 2.24 (b), each frame, or individual switch configuration, will have 32 switch elements actuated. The speed at which the pattern iterates through its frames

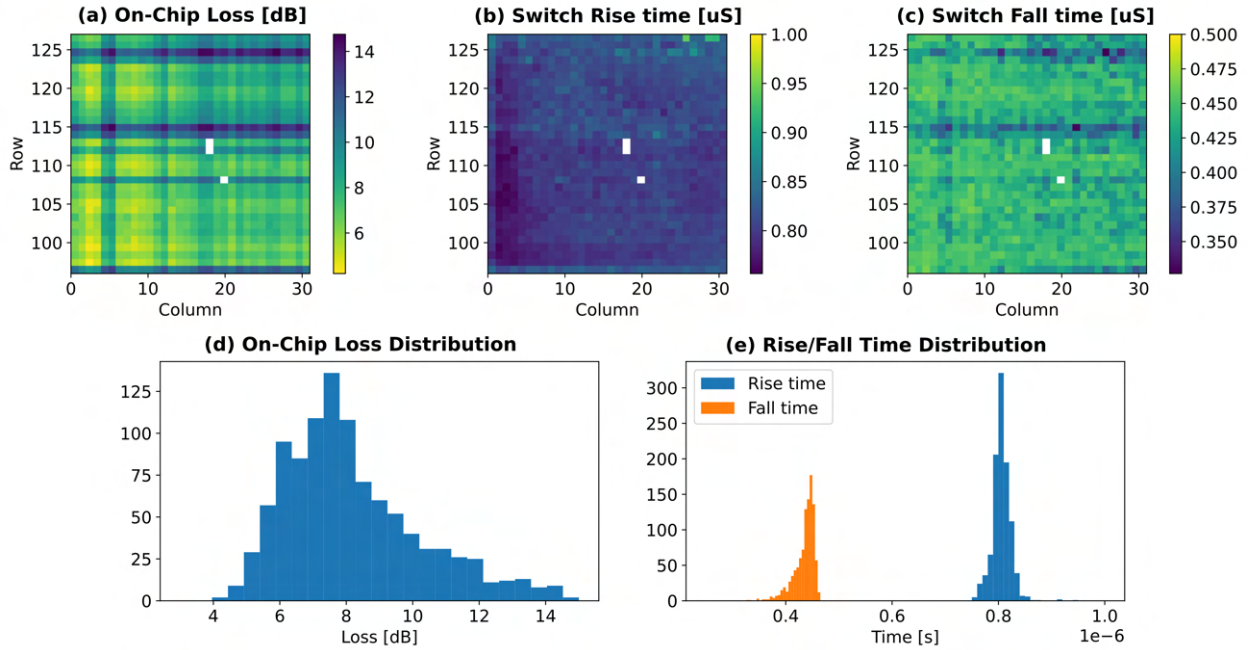


Figure 2.23: Measured on-chip loss (a/d), rise time (b/e), and fall time (c/e) for all $32^2 = 1024$ switch elements. Measurements were done at an HVDD of 60V and a wavelength $\lambda = 1310\text{nm}$.

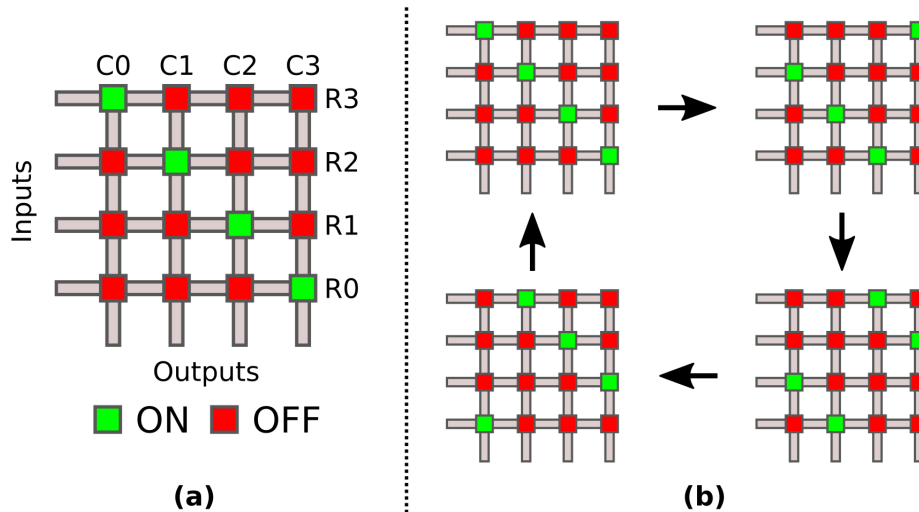


Figure 2.24: (a) Example 4x4 OCS w/ one driver "ON" per row/column. (b) Worst case functional pattern for power measurements.

is called the switch reconfiguration frequency. The total power, including both the dynamic and static contributions, is shown as a function of switch reconfiguration frequency in Fig. 2.25 (a). As previously mentioned, the static power is quite small at only 3.61mW. The maximum observed power is 72mW which was measured at the maximum reconfiguration frequency of 1.7MHz. This frequency corresponds to the time required to scan in 28 instruction bits at the maximum digital clock frequency, 50MHz: $f_{reconfig-max} = 1/\frac{28+1}{50MHz} \approx 1.7MHz$. 72mW represents the absolute maximum power for this 32x32 OCS. Typical power consumption will likely be much less as at least some consecutive frames will leave a few elements unchanged. The reconfiguration frequency will also likely not remain constant at the maximum frequency during typical operation. Fig. 2.25 (b) shows the power breakdown by supply for the maximum reconfiguration frequency. HVDD is set to 60V while all other supplies are set to their nominal values as shown in Fig. 2.9 (b). 60V is chosen as it is the minimum voltage required to ensure all rise and fall times are less than 1 μ s. $VDDPST$, not shown in Fig. 2.9 (b), is the IO supply voltage and is set to 5V. This supply consumes over 50% of all power at the highest reconfiguration frequency and is highly dependent on the PCB layout. A more optimized PCB layout than what is shown in Fig. 2.18 could be used to further reduce the power consumption. A lower IO supply voltage could also be used to reduce power (down to 1.8V for this chip) but would also reduce the maximum attainable reconfiguration frequency.

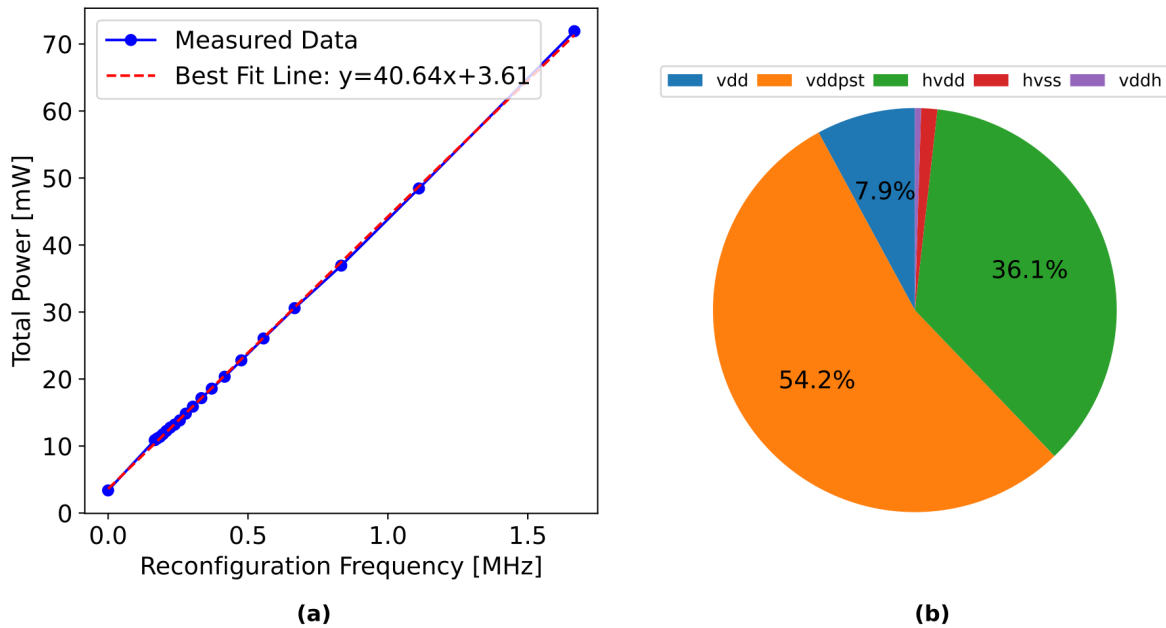


Figure 2.25: (a) Measured worst case power for 32x32 OCS w/ HVDD=60V. (b) Power breakdown by supply for maximum reconfiguration frequency of $\approx 1.7MHz$.

Chapter 3

SuperSwitch 2: An Analog Alternative

3.1 Problem Statement

Most foundry processes that specialize in one type of technology, either MEMS or SiPh, have never been used or are not optimized for the other type. For this reason, it's generally desirable to create a custom in-house process to develop these OCSs (such as the switch presented in [38]). Unfortunately, custom processes make it hard to scale the research into a real product. This is the reason why SuperSwitch 2 uses a photonics optimized process from the AIM Photonics foundry [31]. This means the photonics devices themselves (e.g. waveguides, couplers, etc...) should be of high quality and, most importantly, the potential exists for the OCS to be manufactured at high volumes. The MEMS release process still needs to be handled in-house, but the research done through SuperSwitch 2 provides incentive for commercial foundries to offer a photonics optimized MEMS process.

AIM Photonics 8x8 SiPh MEMS OCS

A cross section of the "Base Active PIC" process from AIM Photonics is shown in Fig. 3.1 (a). This cross section and other information on this process can be found on AIM's website [31]. This process unfortunately does not contain a compatible polysilicon layer to build a coupler waveguide layer similar to the one used in SuperSwitch 1 (see Fig. 2.1 (a)/(b)). This necessitates the use of lateral adiabatic coupling, such as that shown in [11], to create the basic OCS switch element.

Fig. 3.1 (b) shows a simplified cantilever model of the released waveguide layer. Light is guided through the SOI layer and into the page. A fixed SOI waveguide, not shown in Fig. 3.1, is located to the right of the MEMS-actuated SOI waveguide. This fixed waveguide provides the same 90-degree turn that was previously provided by the coupler waveguide layer (see Fig. 2.1 (c)). The released SOI waveguide will deflect upwards when a voltage is applied between the Metal 2 and SOI layers. This will force the light to continue through

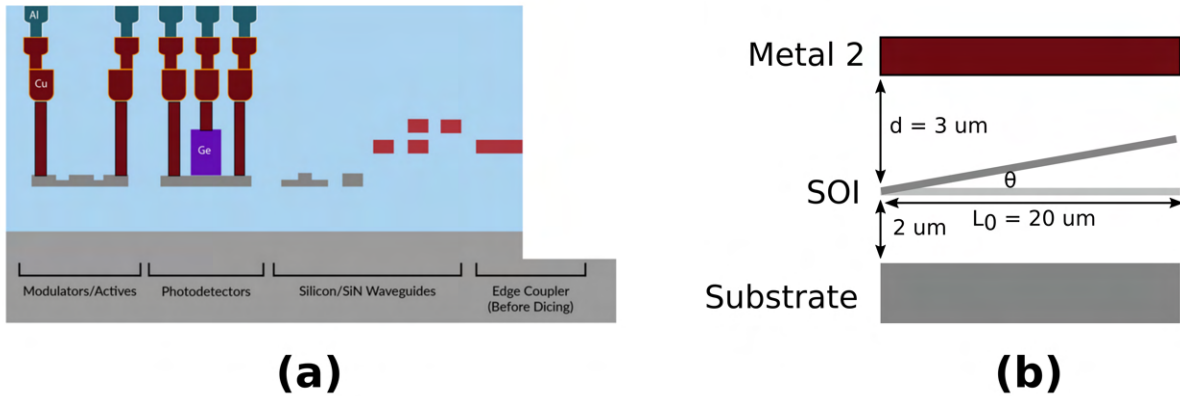


Figure 3.1: (a) Cross section of AIM Photonics base active PIC process - figure taken from [31]. (b) Simplified model of MEMS cantilever switch.

the released waveguide instead of coupling into the fixed waveguide. This creates the basic switch element which can then be used to build a full crossbar OCS similar to SuperSwitch 1. In contrast to SuperSwitch 1, light is coupled from input to output when no voltage is applied. Only when a high voltage is applied does the light pass through the switch element and into the adjacent column's input. This difference is illustrated in Fig. 3.2 (a). A 4x4 OCS is used as an example. The input waveguides are pictured as gray horizontal lines while the output waveguides are gray vertical lines. The guided light is shown as arrows of varying colors. Each switch element is depicted as either a red or green square. A green square, labelled as "ON", represents a switch element with a high voltage applied between its Metal 2 and SOI layer. The crossbar nature of the switch works the same as SuperSwitch 1 with only 1 switch element per column needing to couple light from input to output. This means that a majority of the switch elements, $k^2 - k$ for a $k \times k$ switch, will have a high voltage applied to them in the steady state. This will be an important detail when it comes to designing the high voltage driver.

The taped out SuperSwitch 2 MEMS SiPh OCS is shown in Fig. 3.2 (b). This switch is an 8x8 OCS with all positive and negative switch element connections made through wirebond pads on the east and west side of the chip. A smaller radix is chosen for the first proof-of-concept SuperSwitch 2. Wirebonding is chosen over flip-chip bonding to reduce the initial cost and difficulty of packaging for these new switch devices. Higher radix designs, that may require flip-chip bonding, will be explored once enough devices have been characterized in this process.

Cantilever MEMS Model

Both [30] and [36] offer detailed derivations of the capacitance, C , and applied electrostatic force, F , of the SuperSwitch 2 cantilever MEMS device (assuming Fig. 3.1 (b) can be mod-

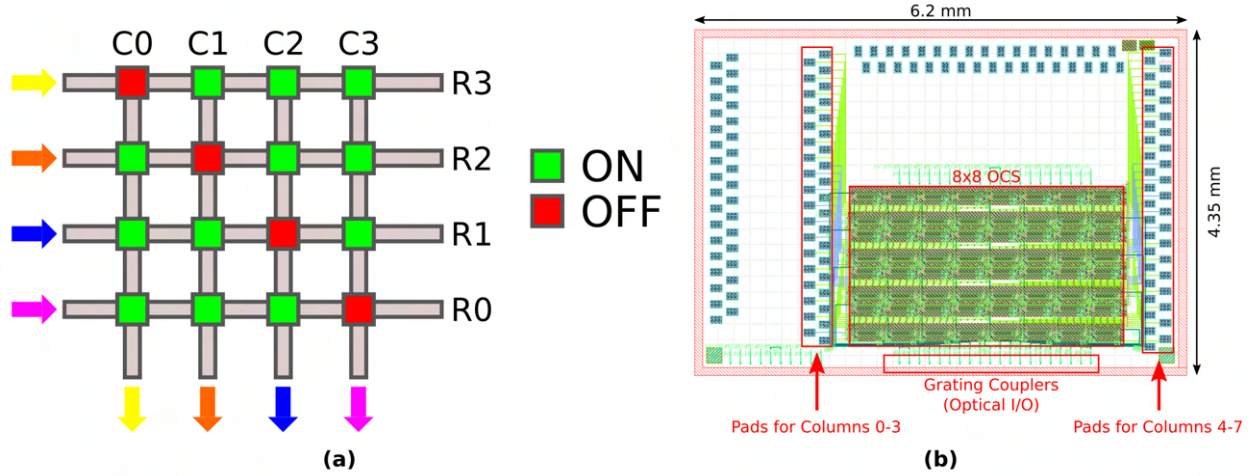


Figure 3.2: (a) Cartoon of example 4x4 crossbar switch. "On" refers to a high voltage being applied. (b) Screenshot of GDS for SuperSwitch 2 silicon photonic MEMS 8x8 OCS.

elled as a rigid beam with a spring at the tip). Equations 3.1 and 3.2 define approximations for the derived capacitance and force for a vertical displacement z and an applied voltage V . The electrostatic force is nonlinear with voltage making analytical modelling difficult.

$$C \approx \frac{\epsilon W L_0}{d} \left[\frac{L}{L_0} + \frac{z}{2d} \left(1 - \left(1 - \frac{L}{L_0} \right)^2 \right) + \frac{z^2}{3d^2} \left(1 - \left(1 - \frac{L}{L_0} \right)^3 \right) \right] \quad (3.1)$$

$$F \approx \frac{\epsilon W L_0 V^2}{d} \left[\frac{1 - \left(1 - \frac{L}{L_0} \right)^2}{2d} + \frac{2z \left(1 - \left(1 - \frac{L}{L_0} \right)^3 \right)}{3d^2} \right] \quad (3.2)$$

The parameters for the MEMS model are shown in Table 3.1. An additional fitting parameter of 0.5 is applied to equation 3.2 in order to better match the dynamics seen in the 3D COMSOL model of the device. Fig. 3.3 shows the equivalent circuit containing the MEMS capacitor, C_{MEMS} , the parasitic capacitance, C_P , and the high voltage driver cell with equivalent source resistance R_S . R_S is a design parameter that is solved for in Section 3.2. The same section will also evaluate C_P and C_{MEMS} and determine the total effective capacitance $C_{eq} = C_P + C_{MEMS}$.

Using the equivalent circuit in Fig. 3.3, the equations 3.1 and 3.2, and the parameters defined in Table 3.1, a state space representation can be made of the MEMS cantilever model:

Table 3.1: Model parameters for the SuperSwitch 2 MEMS unit cell.

Symbol	Meaning	Value	Unit
H	Thickness of cantilever	110	nm
L_0	Length of cantilever	20	μm
W	Width of cantilever	300	μm
L	Overlap length of capacitor	6	μm
f_{res}	Resonant frequency	250	kHz
k	Spring constant = $\frac{EWH^3}{L^3}$	2	N m^{-1}
E	Youngs Modulus	160	GPa
ζ	Normalized damping coefficient	0.1	
M_0	Real mass of cantilever = $\rho W L H$	1.5378e-12	kg
M	Effective mass = $\frac{k}{(2\pi f_{res})^2}$	8e-13	kg
ρ	Density of silicon	2330	kg m^{-3}
b	Unnormalized damping = $2\zeta\sqrt{kM}$	2.54e-7	N s m^{-1}
d	Initial gap of capacitor	3	μm




Figure 3.3: Circuit model for high voltage driver and MEMS device.

$$\dot{x}_1 = \frac{1}{R} \left(V_{in} - \frac{x_1}{C_{eq}} \right) \quad (3.3a)$$

$$\dot{x}_2 = x_3 \quad (3.3b)$$

$$\dot{x}_3 = \frac{1}{m} (F - kx_2 - bx_3) \quad (3.3c)$$

The state variable x_1 represents electrical charge while x_2 and x_3 represent the displacement and velocity of the tip of the cantilever beam respectively. Going forward, the above equations are numerically solved to find the displacement of the MEMS device as a function of time given an applied voltage waveform. The optical transmission as a function of displacement has been extracted from COMSOL and is shown in Fig. 3.4.

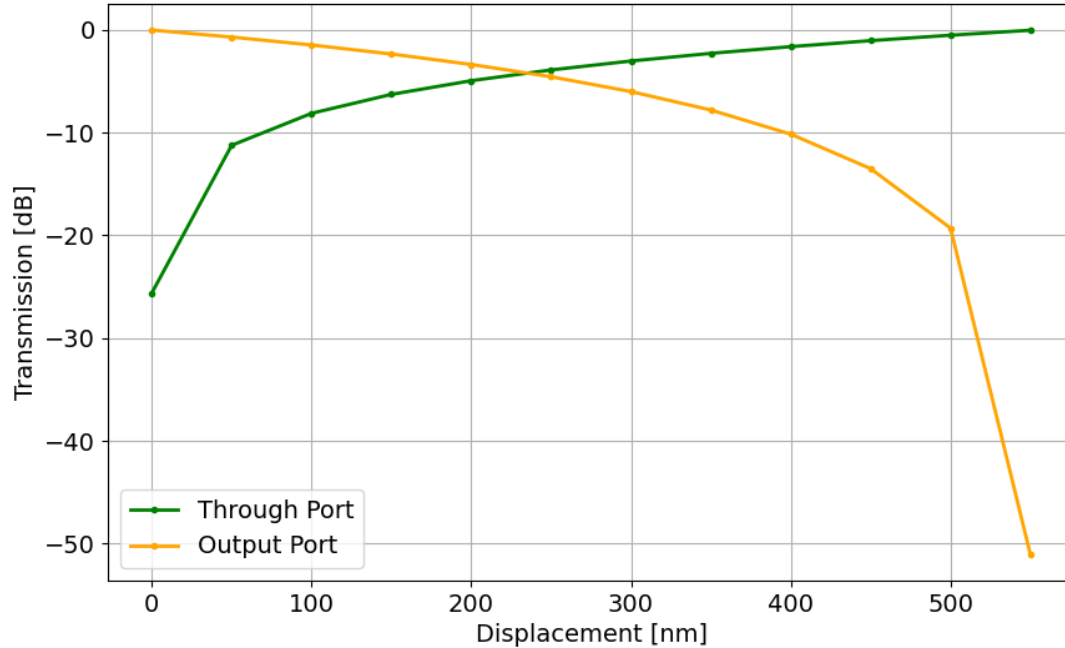


Figure 3.4: Optical transmission vs. displacement for a single MEMS switch element.

Analog High Voltage MEMS Control

Because these devices have yet to be tested, a similar transfer function to Fig. 2.2 cannot be used to determine the type of high voltage driver required. In lieu of experimental results, the model detailed in the previous section, as well as a more detailed COMSOL model, can be used to guide the design process. The SuperSwitch 1 driver was designed to force the switching element in and out of the pull-in condition. This made it easy to design the driver as the voltage only ever needed to swing between a voltage higher than the pull-in voltage and a voltage lower than the pull-out voltage. SuperSwitch 2 cannot operate in such a fashion because pull-in would cause the two MEMS capacitor plates, the released SOI waveguide layer and the Metal 2 layer, to come into contact. This would cause a short between the output of the high voltage driver and whatever is connected to the SOI waveguide layer (VSS typically). SuperSwitch 1 avoided this problem by adding mechanical stoppers, highlighted in green in Fig. 2.1, to stop the coupler waveguide at a defined distance away from the bus waveguide. Any oscillations that the SuperSwitch 1 device would experience are eliminated by these mechanical stoppers. Without these stoppers, the driver must carefully control the MEMS device such that it never experiences pull-in. As explained in [30], pull-in approximately occurs at 1/3 of the initial displacement between a beam and its corresponding electrode. Fig. 3.1 shows that the initial displacement is 3 μ m. This means the released waveguide can

displace by $\approx 1\mu\text{m}$ before it pulls-in and shorts out.

Fig. 3.5 (a) shows the step response of the model developed in Section 3.1 for a target steady-state displacement of 550nm. Going forward, 550nm will be used as the "ideal" target displacement as it is sufficiently far from the pull-in condition and provides sufficient optical extinction ($> 50\text{dB}$ as shown in Fig. 3.4. The peak displacement is around $0.9\mu\text{m}$ which is very close to the approximate pull-in condition of $1\mu\text{m}$. Manufacturing defects and other non-idealities could easily cause pull-in at displacements less than $1\mu\text{m}$. This highlights the importance of minimizing oscillations to avoid any potential pull-in events. In addition, these oscillations greatly increase settling time, which directly impacts the switch reconfiguration time, one of our key figures of merit. One common method of minimizing oscillations for MEMS devices, and similar dynamical systems, is to employ a dual-step response such as that shown in Fig. 3.5 (b). The applied voltage, shown in red, starts at 0 and steps twice. The displacement, shown in blue, ramps steadily to the target displacement before stopping. This dual-step response is described analytically in [15], but is fairly simple to understand intuitively. The first voltage should be chosen such that the initial peak displacement is equal to the target displacement. The second voltage should be applied at the instant at which the initial peak displacement occurs and should correspond to a steady state displacement equal to the target displacement. When the cantilever is at its initial peak displacement the velocity is zero but the acceleration is nonzero. If the forces can be cancelled at exactly this moment, by applying the voltage that corresponds to the target steady state voltage, then the acceleration also becomes zero and the cantilever comes to rest.




Figure 3.5: (a) Step response for 550 nm displacement (b) Dual step response for 550 nm displacement

The model can then be used to generate a lookup table of peak and steady state displacements as well as the time at which each peak displacement occurs. Both plots in Fig. 3.6 are made by numerically solving the step response for different voltages given the force equation 3.2 and the state space equations 3.3a, 3.3b, and 3.3c. Fig. 3.6 (a) shows the peak displacement and steady state displacement versus voltage. These plots determine the

voltages that should be used for the first and second steps. Fig. 3.5 (b) shows the time at which the peak displacement occurs which determines the time at which to apply the second voltage.




Figure 3.6: (a) Step response peak displacement and time vs. applied voltage (b) Step response steady state displacement vs. time

3.2 CMOS Design & Simulation

Assuming the simple dual-step control strategy from Section 3.1 is used, the SuperSwitch 2 high voltage driver must be capable of applying at least 2 arbitrary high voltages in sequence. To accomplish this, the final CMOS chip, shown in Fig. 3.7, contains 64 high voltage DACs to independently control all switching elements of the 8x8 SiPh MEMS OCS. This chip was taped out in the same process as the SuperSwitch 1 CMOS chiplet, TSMC’s 180nm HV BCD Gen2 process. Each high voltage DAC is contained within a ”cell controller” block that contains additional circuitry required to implement the MEMS control strategy. All 64 cell controllers are labelled in red in the interior of Fig. 3.7 chip. The leftmost cell controllers, 0-31, have their high voltage DACs connected to the wirebond pads on the west side of the chip. The other 32 are connected to the wirebond pads on the east side of the chip. All HV DAC wirebond pads are pitch matched to the east/west pads of the SiPh chip pictured in Fig. 3.2 (b). This makes the wirebond packaging much easier to implement, as will be discussed later in Section 3.3. The power supplies and digital IOs are connected to the south side wirebond pads.

DAC Resolution Simulations

The first question related to the design of the HV DAC is the number of bits of resolution to be used. Fig. 3.8 shows a number of transient simulations using the model previously

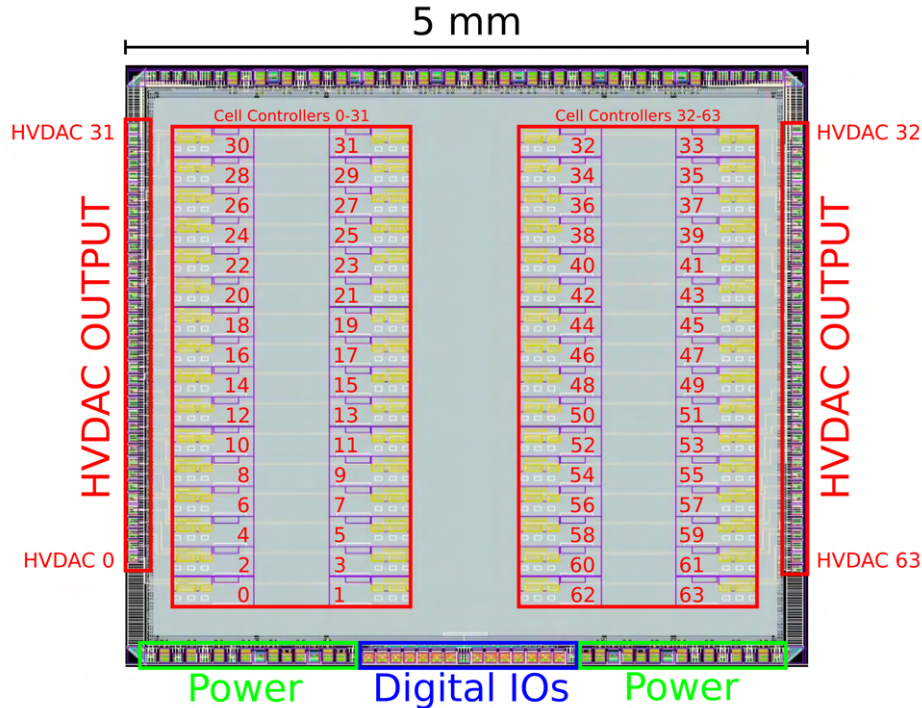


Figure 3.7: SuperSwitch 2 - high voltage DAC array for analog MEMS control

described in 3.1. These simulations have been done to understand the effects of voltage offset on the dual-step control strategy. The voltage values for the first and second voltage steps are offset from $-5V$ to $+5V$ from the ideal values found using Fig. 3.6. Figures 3.8 (a) and (b) show the applied voltages and the calculated responses respectively. Figures 3.8 (c) and (d) show the static and dynamic error for each transient simulation.

The maximum quantization error for a DAC is $1/2$ of an LSB, assuming that the target voltage is within the full scale range. For a $70V$ full scale range $1 \text{ LSB} = \frac{70V}{2^N}$ where N is the number of bits of resolution. Vertical lines are drawn on Figures 3.8 (c) and (d) to show the maximum quantization error for $N = 4, 5,$ and 6 . To ensure that the static error is significantly less than 10% , a resolution of at least 6 bits is required. Anything beyond 6 bits will further reduce our static error but also increase the area and power of the HV DAC, something of key concern when talking about building an array of HV DACs. **For this reason, a resolution of $N = 6$ is chosen for the SuperSwitch 2 high voltage DAC.** The dynamic error, or settling time, is less important as the difference between 4 and 6 bits is only $10s$ of nanoseconds. The one interesting thing about Fig. 3.8 is that the minimum settling time does not occur at an offset voltage of $0V$. This is expected as the dual-step waveform does not optimize for settling time. Other feedforward techniques, one of which will be discussed later, can be used to further reduce the settling time.

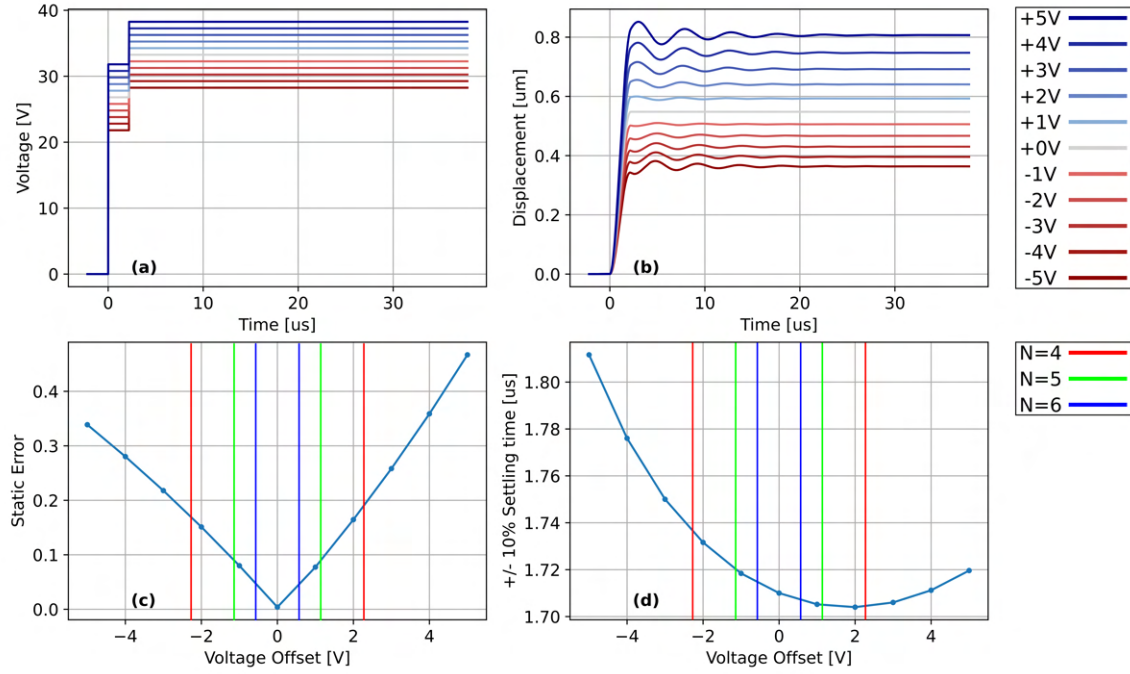


Figure 3.8: (a) Dual step waveforms w/ voltage offset (b) Cantilever displacement w/ voltage offset (c) Static error (d) Settling time for +/- 10%

Digital Architecture

The crossbar nature of the switch necessitates a similar high-level digital architecture to what was done in SuperSwitch 1. Each switching element is assigned a specific row and column and the CMOS driver is responsible for selecting 1 element per column (or per row) to switch light through. Because the wirebond pads are not physically arranged in an 8x8 grid, in contrast to the 32x32 grid of flip-chip pads for SuperSwitch 1, the row/column mapping from each cell controller becomes a bit complicated. The mapping from the switch column number to the HVDAC IDs is shown in Table 3.2. The row for each HVDAC is programmable such that the SuperSwitch 2 CMOS driver chip can be more easily reused in future OCS designs.

Fig. 3.9 shows an illustration of the full digital architecture of the switch. This figure is meant to complement the annotated layout shown in Fig. 3.7. The cell controllers, which contain the HV DACs, are numbered from 0-63. These numbers correspond to the HV DAC IDs in Table 3.2. There are 3 total scan chain interfaces shown in Fig. 3.9. 2 of these interfaces are labeled as "select 1" and "select 2" in the diagram. They are responsible for selecting a row within each column. The final interface, labelled as "control", is responsible

Table 3.2: Mapping between column (0-7 for 8x8 switch) and HVDAC IDs.

Column	HVDAC IDs
0	2, 5, 12, 16, 20, 23, 24, 31
1	3, 4, 9, 15, 21, 22, 25, 30
2	0, 7, 10, 13, 18, 19, 26, 29
3	1, 6, 8, 11, 14, 17, 27, 28
4	60, 59, 55, 51, 44, 43, 36, 35
5	61, 58, 54, 50, 45, 42, 37, 34
6	62, 57, 53, 49, 46, 41, 38, 33
7	63, 56, 52, 48, 47, 40, 39, 32

for setting various configuration bits across the chip and within each cell controller. All 3 of these interfaces will be discussed in more detail in the following subsections.

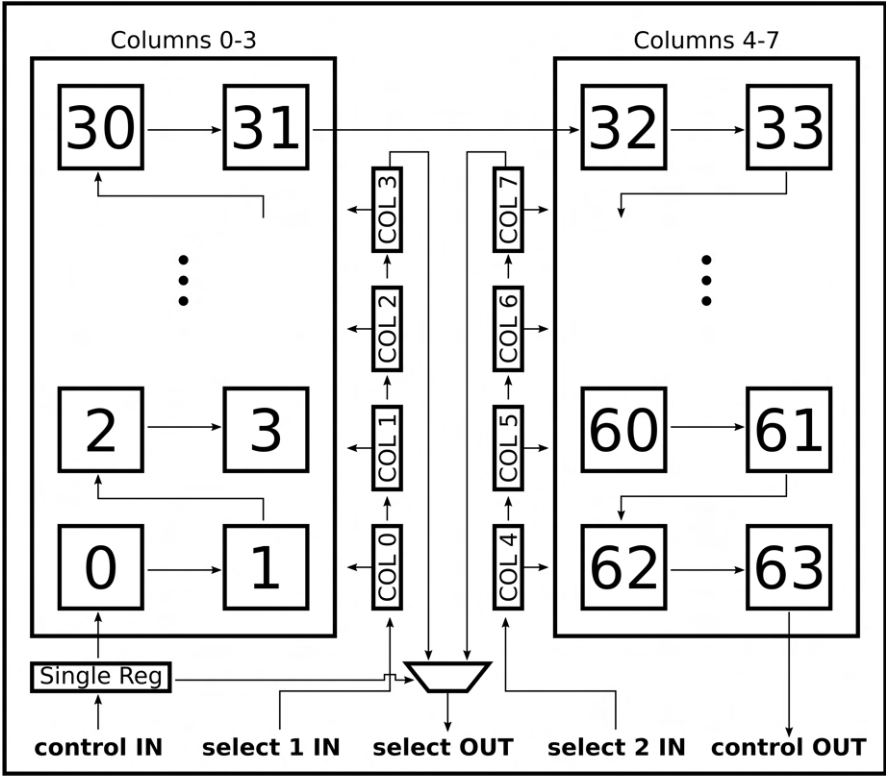


Figure 3.9: Simplified chip-level digital architecture showing control and select scan chains

Select Interface

The select interfaces provide 3 bits ($\log_2 8$) per column to select a row. The "select 1" interface is responsible for columns 0-3 and "select 2" is responsible for columns 4-7. Both interfaces are scan chains that contain a total of 12 bits each, 3 bits for each of the 4 columns that they are responsible for. Table 3.3 details each of the scan chain fields contained within the 2 select interfaces. To minimize the number of IOs required to support both select interfaces, a single select interface output is used. As shown by the mux at the bottom middle of Fig. 3.9, this single output is the scan chain output of either select interface. A single register, which is part of the control interface, determines which of the two select interfaces should be output from the chip. This does not affect the functional behaviour of the chip as the scan chain output is only used for debug purposes.

Table 3.3: Global "Select" scan chain bits listed in the order that they are connected.

Chain	Field	Width	Description
Select 1	selected_address_0	3	Selects 1 of 8 rows w/in column 0 to activate.
Select 1	selected_address_1	3	Selects 1 of 8 rows w/in column 1 to activate.
Select 1	selected_address_2	3	Selects 1 of 8 rows w/in column 2 to activate.
Select 1	selected_address_3	3	Selects 1 of 8 rows w/in column 3 to activate.
Select 2	selected_address_4	3	Selects 1 of 8 rows w/in column 4 to activate.
Select 2	selected_address_5	3	Selects 1 of 8 rows w/in column 5 to activate.
Select 2	selected_address_6	3	Selects 1 of 8 rows w/in column 6 to activate.
Select 2	selected_address_7	3	Selects 1 of 8 rows w/in column 7 to activate.

Control Interface

The "control" interface is stitched through all cell controllers starting w/ cell controller 0 and ending with cell controller 63. A summary of all bits stored in the control scan chain are shown in Table 3.4. Even greater detail into the individual cell controller bits is given in Section 3.2. The only other global configuration bit is the *scan_out_sel* field which controls the select interface output mux. The total length of the scan chain is 6,465 bits. This scan chain is not intended to be configured during normal operation. Programming of this chain should be done after power up and before any rows are selected through the select interfaces.

Cell Controller

Fig. 3.10 shows a simplified schematic of a single cell controller. The "cell" in this context refers to a single switch element located on the SiPh chip. The cell controller itself contains 5 total circuit elements. The most basic element, discussed in the previous section, is the control scan chain. All bits of this scan chain are detailed in Table 3.5. Each cell controller

Table 3.4: Global "Control" scan chain bits listed in the order that they are connected.

Field	Width	Description
scan_out_sel	1	Sets the output mux shown in Fig. 3.9. This mux drives the output data signal for the select scan chains. 0 selects select chain 1 and 1 selects select chain 2.
Cell Controller 0	101	All configuration bits for 0th cell controller.
...		
Cell Controller 63	101	All configuration bits for 63rd cell controller.

has its own set of bits that allow for the enabling and configuring of the remaining 4 cell controller circuit elements. These remaining elements include a static current DAC (IREFP), a dynamic current DAC (6-bit IDAC), a finite state machine (FSM) and a high voltage driver (HVDRIVER). The IREFP circuit accepts a reference current and generates a new internal reference current for use by the 6-bit IDAC. The 6-bit IDAC pulls current through the HVDRIVER circuit to produce the high voltage that will control the MEMS switching element. The 6-bit IDAC combined with the HVDRIVER can be thought of as the basic 6-bit HV DAC. The FSM controls the HV DAC to apply the correct voltages based on whether that particular cell controller's row has been selected.

Table 3.5: List of cell-controller control scan chain fields. Each cell controller has its own unique set of these fields.

Field	Width	Description
address	3	Sets the address, or row ID, of the cell controller.
iref_cfg	9	Configuration bits for the cell controller's current reference DAC.
idac_cfg	2	Configuration bits for the cell controller's current steering DAC.
fsm_cfg	80	Configuration bits for the cell controller's finite state machine.
bypass_val	6	The bypass DAC code. Applied to DAC only when bypass_enb is 0.
bypass_enb	1	Enables bypassing the cell controller FSM when set to 0.

Resistively Loaded HV Driver - HVDRIVER

Perhaps the most important circuit within the cell controller is the HVDRIVER, shown on the right side of Fig. 3.10 as a simple resistor connected to a high voltage supply (HVDD). This circuit is responsible for directly driving the MEMS switching element and is the only high voltage circuit in the entire chip. In reality, the HVDRIVER is composed of a current mirror network that pulls current through a resistor to generate an analog output voltage.

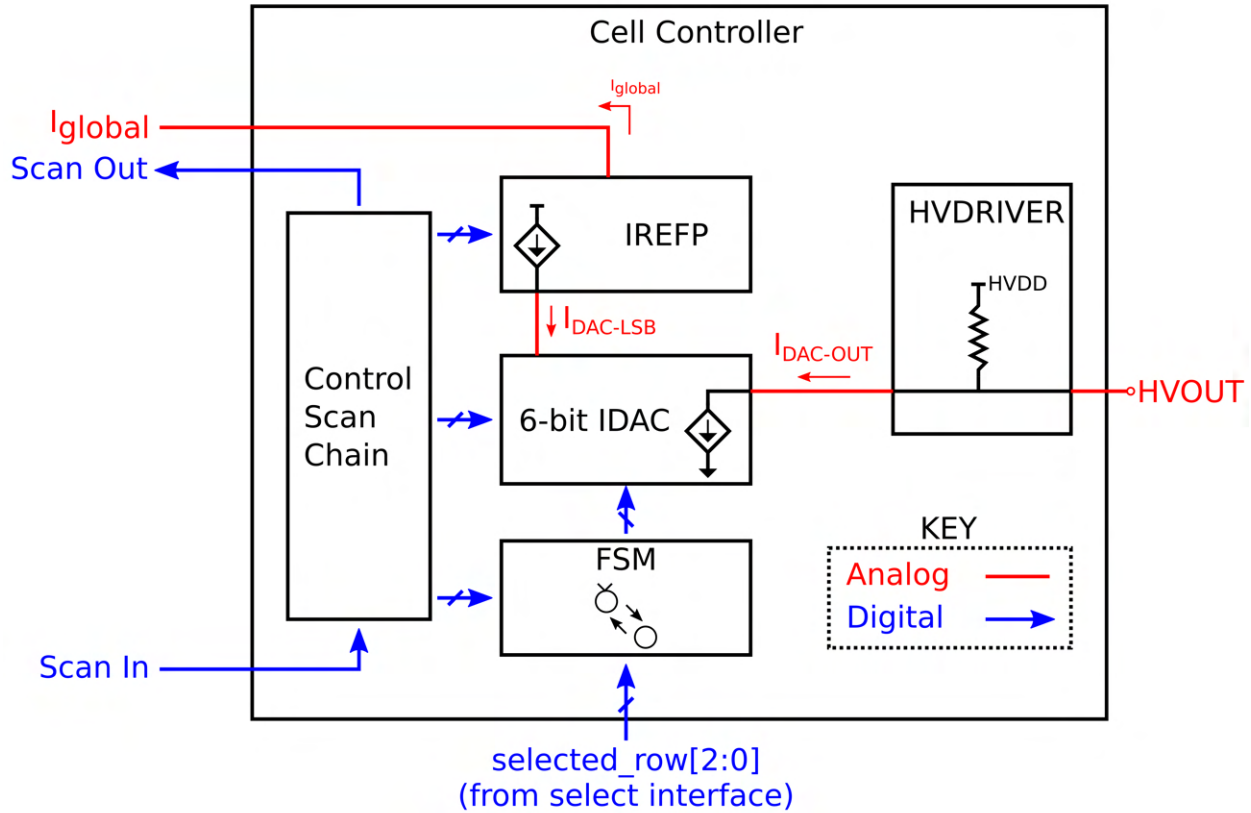


Figure 3.10: Simplified schematic of a single cell controller

The schematic of the HVDRIVER is shown in Fig. 3.11 (a). V_{DDH} is nominally 5V and HV_{DD} is nominally 70V. Transistors M_1 - M_3 are rated for 5V while M_4 and M_5 are rated up to 70V (across V_{ds}). Only M_5 is subject to high voltage but a high voltage transistor is also used for M_4 to give better current matching between M_4 and M_5 . A total current mirror ratio of 10 is used to minimize the current and power required to drive the input of the HVDRIVER. Current mirror ratios beyond 10 proved difficult to maintain across all operating points. The resistor was intentionally chosen to be connected to an NMOS sink instead of a PMOS source. This is because the majority of the switching elements will be at a high voltage during normal operation as previously shown by Fig. 3.2 (a). An NMOS-sink HVDRIVER outputs a high voltage at low current while a PMOS-source HVDRIVER outputs a high voltage only at high current. Using the chosen style of HVDRIVER gives us a much lower total power consumption, something that will be discussed further in Section 3.5.

The basic sizing parameters are given in Fig. 3.11 (b). The first length, L_3 , is chosen such that the r_o of all 5V transistors is sufficiently large. The width W_3 is calculated using

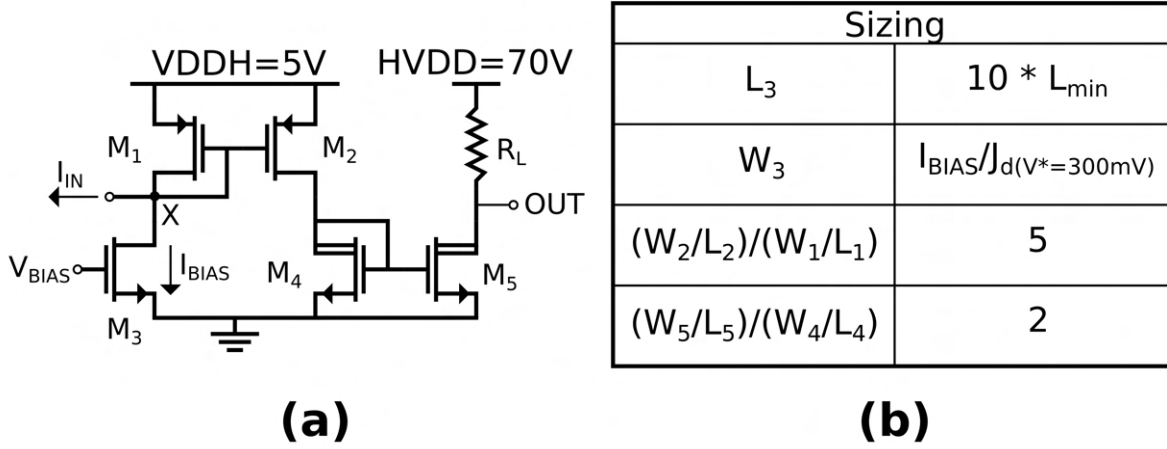


Figure 3.11: (a) Schematic of high voltage driver cell - M_4 and M_5 are high voltage transistors. (b) Sizing guide for driver given a particular I_{BIAS} .

the total input bias current, I_{BIAS} , and the current density corresponding to a V^* of 300mV. V^* is a transistor metric that is equal to the overdrive voltage assuming that the transistor perfectly follows the basic square law model (so $V^* = \frac{2I_D}{g_m}$ which is $\approx V_{ov}$ even for non-square law devices). The input bias current, I_{BIAS} , is calculated after the minimum and maximum input currents are known. I_{BIAS} needs to be large enough such that all transistors remain in saturation regardless of the input current to the HVDRIVER.

The first step in the design process is to calculate the required resistive load, R_L , in order to guarantee a specific settling time. Before this can be done, the total load capacitance needs to be estimated. Fig. 3.12 (a) shows the total extracted parasitic capacitance for all 64 cell controllers. This capacitance is compared against the calculated capacitance of the switching element model plotted in Fig. 3.12 (b). The MEMS capacitance is a few femto Farads compared to a worst case parasitic capacitance of nearly a pico Farad. For this reason, the total capacitance, $C_{load-max}$, will be assumed to be 1pF. This will ensure that the target settling time is met for all HVDRIVER circuits. The parasitic capacitance is much worse than in SuperSwitch 1 due to large wirebond pads and long routes from the output of the HVDRIVER circuits to the west or east side of the chip. Eq. 3.4 calculates the required resistance R_L assuming a desired rise/fall time of 100ns. This value was chosen as it is $> 10x$ smaller than the calculated rise time of the switch element itself ($\approx 1.7\mu s$ from Fig. 3.8 (d)).

$$R_L = \frac{\tau}{C_{load-max}} \approx \frac{t_{settle}/3}{C_{load-max}} = \frac{100ns/3}{1pF} \approx 33.3k\Omega \quad (3.4)$$

Now that R_L is known, the maximum output current, $I_{output-max}$, can be easily calculated assuming a specific full scale range. The model developed in Section 3.1 suggests that 70V,




Figure 3.12: (a) Histogram of parasitic output cap from all 64 HV DAC outputs to their corresponding pads. (b) Capacitance of MEMS device as voltage is stepped.

the highest voltage supported in the process, likely isn't required but is chosen anyway to provide the most design flexibility. Eq. 3.5 uses Ohm's Law to calculate the required maximum output current to give us a full scale range of 70V.

$$I_{output-max} = \frac{HVDD}{R_L} = \frac{70V}{33.3k\Omega} \approx 2.1mA \quad (3.5)$$

The output current, I_{output} , can also be expressed as a function of the input current, I_{IN} , and the bias current, I_{BIAS} . I_{IN} is simply the output current of the 6-bit IDAC shown in Fig. 3.10 and is equal to the 6-bit code multiplied by the LSB current ($I_{IN} = code I_{LSB}$). The bias current can also be expressed as a multiple of the LSB current to make the design simpler ($I_{BIAS} = F_{BIAS} I_{LSB}$). Eq. 3.6 defines the output current for any DAC code. M_{mirror} is the total current mirror ratio of 10 as shown in Fig. 3.11 (b).

$$I_{output} = M_{mirror}(I_{IN} + I_{BIAS}) = M_{mirror}(code + F_{BIAS})I_{LSB} \quad (3.6)$$

By using the maximum code, $2^N - 1$ where $N = 6$ for our 6-bit DAC, the equations 3.5 and 3.6 can be set equal to solve for the LSB current, I_{LSB} .

$$I_{LSB} = \frac{HVDD}{R_L M_{mirror}((2^N - 1) + F_{bias})} \quad (3.7)$$

Before Eq. 3.7 can be solved, F_{BIAS} must be selected. In order to do so, one must understand the effects of F_{BIAS} on the internal node voltages of the HVDRIVER circuit. Node X, shown in Fig. 3.11 (a), will approximately be $V_{X-max} \approx V_{DDH} - (V_{OV-min} + |V_{th-M1}|)$ when $I_{IN} = 0$ (i.e. V_{OV-min} corresponds to just I_{BIAS} flowing through M_1 and M_3). Designing for a V^* of 300mV gives us a $V_{X-max} \approx 5V - (300mV + 700mV) = 4V$. When $I_{output} = I_{output-max}$, the change in voltage at node X, ΔV_X , needs to be small enough that M_3 remains in saturation. An F_{BIAS} of 4 gives a ΔV_X of only 1.2V, plenty small enough

to keep M_3 in saturation. Eq. 3.8 shows how this number is calculated. N is the number of DAC bits, 6, and $2^N - 1$ represents the maximum DAC code.

$$\begin{aligned} \Delta V_X &\approx V^* \left(\sqrt{\frac{I_{IN-max} + I_{BIAS}}{I_{BIAS}}} - 1 \right) = V^* \left(\sqrt{\frac{2^N - 1 + F_{bias}}{F_{bias}}} - 1 \right) \\ &= 300\text{mV} \left(\sqrt{\frac{2^6 - 1 + 4}{4}} - 1 \right) \approx 928\text{mV} \end{aligned} \quad (3.8)$$

F_{BIAS} is then selected to be 4 as it sufficiently reduces ΔV_X while only requiring $4I_{LSB}$ of static current through the input stage of the HVDRIVER. Eq. 3.7 can now be solved:

$$I_{LSB} = \frac{70\text{V}}{33.3\text{k}\Omega * 10 * (63 + 4)} \approx 3.13\mu\text{A} \quad (3.9)$$

6-bit Current Steering IDAC

The input to the HVDRIVER circuit is provided by a 6-bit current DAC as shown in Fig. 3.10. This DAC is designed as a fully segmented current-steering DAC. The schematic for the DAC is shown in Fig. 3.13 (a). The 6 bit input code, labelled as *dac_in*, is first translated into thermometer code. The output thermometer code has a total number of bits set to 1 that is equal to the input binary code. For example, an input code of 0 would generate an output code of 63 0s while an input code of 4 would output 4 1s and 59 0s. This thermometer-coded DAC is guaranteed to be monotonic as each additional segment turned on is guaranteed to increase the output current. After converting to thermometer code, the bits are registered such that on the following clock edge all segments can be turned on or off at nearly the exact same time. The segments consist of a single differential pair and one tail transistor as shown in Fig. 3.13. Depending on the value of the segment's input, labelled as *IN*, current should flow in either the left or right side of the differential pair. The tail transistor will remain in saturation regardless of the input, thus allowing current to quickly steer from one side of the segment to the other. The segment input voltage is limited to 1.8V to ensure that the inverted transistor within the differential pair is kept in saturation. More details about the current steering topology and common design techniques can be found in [34].

The tail transistor bias for all segments, V_X in Fig. 3.13 (a) and (b), is generated from a reference current supplied by an internal static current DAC (labelled as IREFP in Fig. 3.10 and discussed in Section 3.2). The LSB current, I_{LSB} , was calculated in the previous section and is equal to 3.13 μ A. Each differential pair transistor was minimum sized while the tail transistor was designed to have a $V^* = 300\text{mV}$ at $I_{tail} = 3.13\mu\text{A}$. V_X from Fig. 3.13 (a) is also used to generate the bias current in Fig. 3.11 (a) by connecting V_X and V_{BIAS} .

Table 3.6 shows all digital configuration registers for the cell controller's current steering IDAC. The field *idac_en* will force V_X to 0V when set to 0. *idac_inv* will invert the output of the binary-to-thermometer block in Fig. 3.13 (a) when set to 1. This is useful because

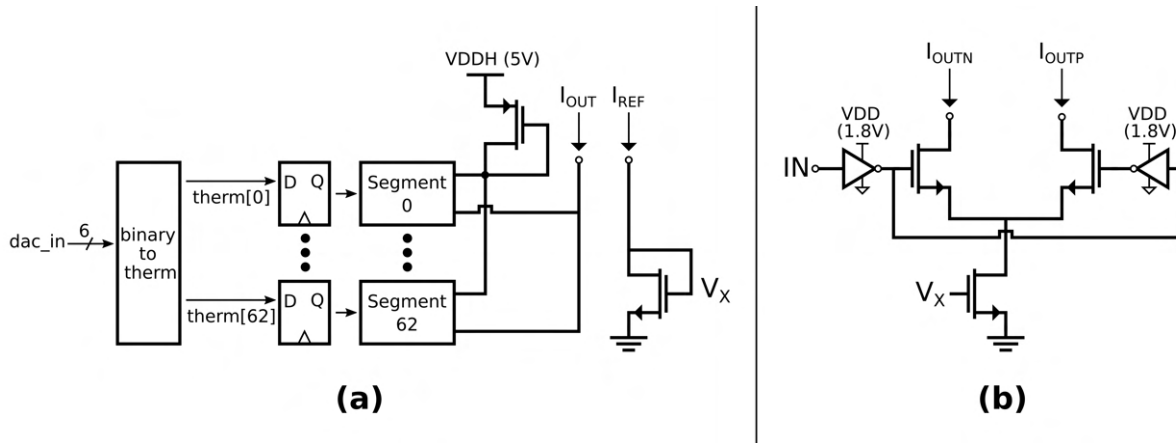


Figure 3.13: (a) Schematic of 6-bit segmented current-steering DAC. (b) Schematic of single current-steering DAC segment.

the HVDRIVER circuit will by default output a high voltage given a low code. By setting $idac_inv$ to 1, a low code will correspond to a low output voltage.

Table 3.6: List of all IDAC control scan chain registers. These bits are labelled as $idac_cfg$ in table 3.5

Field	Width	Description
$idac_en$	1	Enables the cell controller's current steering DAC.
$idac_inv$	1	Inverts the DAC code before the registers in Fig. 3.13.

LSB Current DAC - IREFP

The LSB current DAC, labelled IREFP in Fig. 3.10, is responsible for generating the I_{REF} current shown in Fig.3.13 (a). This will set the maximum current for each DAC segment and thus will also set the LSB current. The IREFP circuit is a simple current mirror w/ PMOS switches to configure the total output current I_{OUT} . A schematic of this circuit is shown in Fig. 3.14.

Table 3.7 lists all configuration registers for the IREFP circuit. $iref_en$ will force the diode-connected transistor's gate to $VDDH$, turning off all output PMOS current sources. $iref_coarse$ is a 4-bit field that turns on 4 separate transistors each with width $1/2$ of the reference transistor. Each bit corresponds to a different transistor. $iref_fine$ works in a similar manner but controls transistors with widths equal to $1/4$ of the reference transistor's width. The input transistor is sized for $V^* = 300mV$ at $I_{REF} = 3.13mA$. This means that the nominal IREFP configuration should correspond to a 1:1 current mirror ratio such that the reference current and output current both equal 3.13mA.

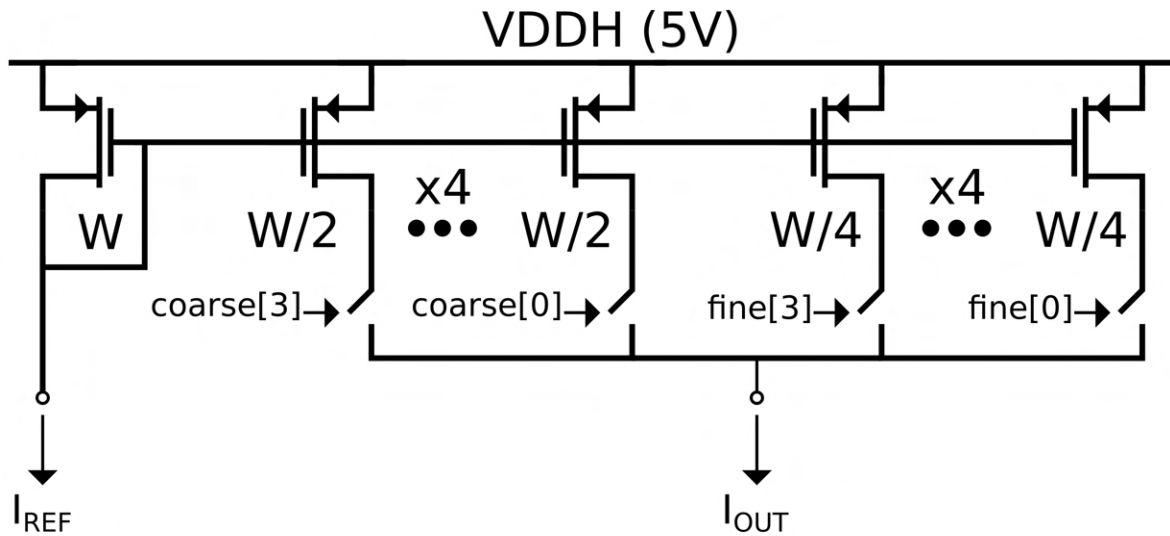


Figure 3.14: Schematic for current reference to set LSB current of current-steering DAC.

Table 3.7: List of all IREFP control scan chain registers. These bits are labelled as irefp_cfg in table 3.5

Field	Width	Description
iref_en	1	Enables the cell controller’s internal current reference.
iref_coarse	4	Each bit set to 1 will add iref/2 to IDAC reference current.
iref_fine	4	Each bit set to 1 will add iref/4 to IDAC reference current.

Selected/Deselected FSM

Each switching element has two separate states: selected or deselected. When selected, the switching element should go from a high voltage to a low voltage such that light is now coupled from its input to output. When deselected, a low-to-high transition should be applied such that light is no longer coupled through that element’s output port. Figures 3.15 (a) and (b) show the voltage and displacement for two switching elements in the same column. Figures 3.15 (c) and (d) show the digital control signals that tell the FSM when it should start applying new voltages. Fig. 3.15 (c) shows the scan enable signal for the select scan chain. The scan chain is loaded while scan enable is asserted. Once deasserted, the "Load" signal, generated from the circuit shown in Fig. 3.16 (c), forces the FSM to start applying the selected or deselected voltage waveform that is programmed into the FSM control registers.

A list of all FSM control registers is shown in Table 3.8. The selected and deselected waveforms consist of 4 separate voltages, *selected_value_[0-3]* and *deselected_value_[0-3]*, and

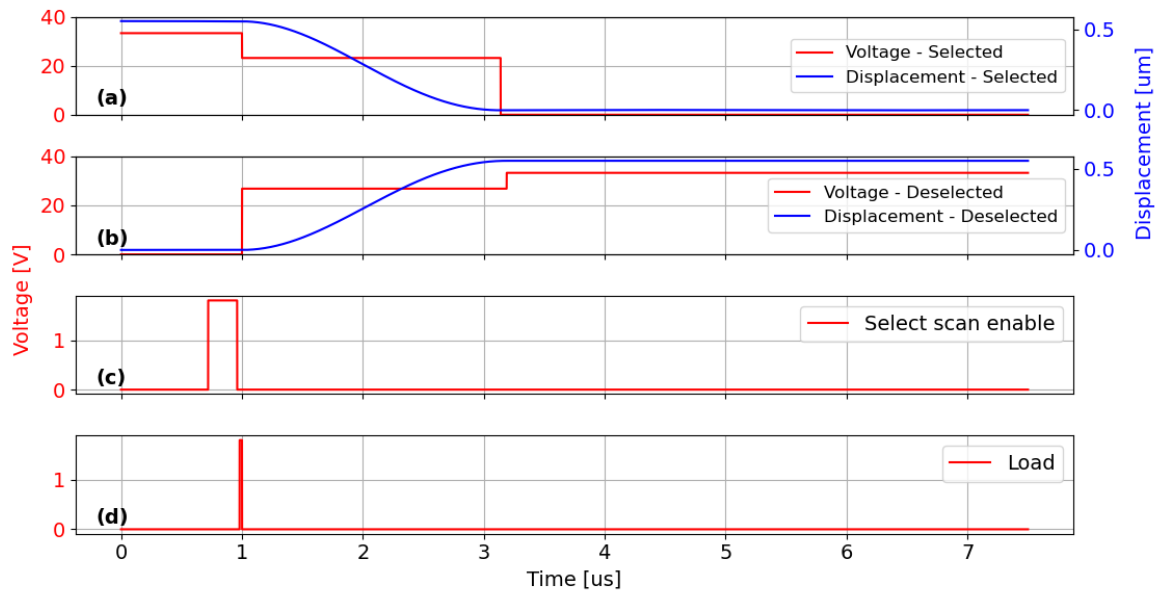


Figure 3.15: Selected and deselected waveforms. The select scan enable signal is shown as well as the load signal that will trigger each cell controller’s FSM to either start cycling through the programmed selected or deselected waveform.

4 related counter values, *selected_hold_[0-3]* and *deselected_hold_[0-3]*. When a switch element becomes either selected or deselected, the corresponding voltages will be applied in series (value *_0* first and value *_3* last). Each voltage value will be applied for the number of clock cycles specified by its corresponding hold control register. The clock frequency is nominally 50MHz giving the FSM a minimum resolution of 20ns for applying voltages at specific times. Once the entire selected/deselected waveform has been applied, the voltage will remain stable at the final value until the switch element’s state changes and a new waveform gets applied. The 4 separate values for the selected/deselected waveforms make it easy to implement the dual-step control strategy discussed in Section 3.1 as well as multi-step waveforms for further optimized dynamic responses.

An FSM diagram is shown in Fig. 3.16 (a). The actions for the various FSM states are detailed in fig. 3.16 (b). The FSM will remain in the *IDLE* state until the *load* signal is equal to 1. Fig. 3.16 (c) shows how the load signal is generated. A digital comparator evaluates the address of the cell controller, programmed through the *address* field in table 3.5, against the selected address for its column, programmed via one of the select interfaces detailed in Section 3.2. Whenever the output of the comparator changes, 0-to-1 or 1-to-0, the *load* signal will be asserted for 1 clock cycle. The FSM will trigger off this *load* signal and will begin applying the selected waveform if the output of the comparator is 1 and the deselected waveform if the output is 0.

Table 3.8: List of all FSM control scan chain registers. These bits are labelled as fsm_cfg in table 3.5

Field	Width	Description
selected_value_0	6	The first value to be applied to MEMS cell when device is selected.
selected_value_1	6	The second value to be applied to MEMS cell when device is selected.
selected_value_2	6	The third value to be applied to MEMS cell when device is selected.
selected_value_3	6	The fourth value to be applied to MEMS cell when device is selected.
selected_hold_0	4	The number of clock cycles to apply selected_pattern_0.
selected_hold_1	4	The number of clock cycles to apply selected_pattern_1.
selected_hold_2	4	The number of clock cycles to apply selected_pattern_2.
selected_hold_3	4	The number of clock cycles to apply selected_pattern_3.
deselected_value_0	6	The first value to be applied to MEMS cell when device is deselected.
deselected_value_1	6	The second value to be applied to MEMS cell when device is deselected.
deselected_value_2	6	The third value to be applied to MEMS cell when device is deselected.
deselected_value_3	6	The fourth value to be applied to MEMS cell when device is deselected.
deselected_hold_0	4	The number of clock cycles to apply deselected_pattern_0.
deselected_hold_1	4	The number of clock cycles to apply deselected_pattern_1.
deselected_hold_2	4	The number of clock cycles to apply deselected_pattern_2.
deselected_hold_3	4	The number of clock cycles to apply deselected_pattern_3.

Global Current Distribution

One final circuit is required to provide the reference current, I_{REF} in Fig. 3.14, for each of the 64 cell controllers. This global current distribution circuit is shown in Fig. 3.17. The circuit consists of 65 transistors all with the same width and all sized for a V^* of 300mV at a current of 3.13 μ A. One transistor is used to mirror an external reference current, $I_{REF-EXT}$, onto 64 NMOS current sources to produce reference currents $I_{REF-CC0}$ through $I_{REF-CC63}$.

Post Layout Simulations

Post layout transient simulations are shown in Fig. 3.18 (a). The simulations include all analog circuit elements from the global current distribution circuit to a single HV Driver. A

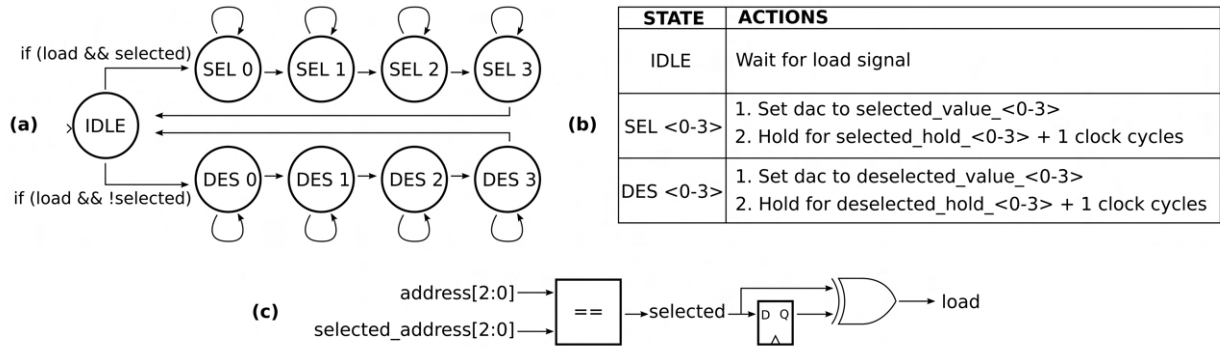


Figure 3.16: (a) Simplified FSM diagram for stepping through selected and deselected waveforms. (b) Description of FSM state actions. (c) Logic to generate the load signal that triggers the FSM.




Figure 3.17: Global current distribution circuit.

reference current of $3.1\mu A$ is fed into the global current circuit from the simulation testbench. Figures 3.18 (a) and (b) show the extracted rise and fall times for a $1pF$ load across all process and temperature corners. The testbench initially drives the current DAC at minimum code, 0, then transitions to maximum code, 63, and then back to minimum. All rise and fall times are close to the desired value of $100ns$, validating the design procedure detailed in Section 3.2.




Figure 3.18: (a) Output voltage of HV DAC going from minimum to maximum code driving a 1pF load. (b) Extracted rise times. (c) Extracted fall times.

3.3 Wirebond-based SiPh-CMOS Packaging

Because the SiPh chip has wirebond connections on both its east and west sides, 2 CMOS chips are required to control 1 SiPh chip. Fig. 3.19 shows a wirebond diagram for this 3-chip package. CMOS 0 controls the east side of the SiPh chip using only cell controllers 0-31. The other 32 cell controllers can be disabled via the *iref_en* register detailed in Table 3.7. CMOS 1 controls the east side of the SiPh chip and similarly enables only cell controllers 32-63. The fiber array will approach the SiPh chip from the south side of Fig. 3.19.

As of writing this dissertation, the SiPh chip has yet to undergo the MEMS release process. Fig. 3.20 (a) shows a picture of a wirebonded package that only contains the CMOS 1 chip from Fig. 3.19. This CMOS-only test vehicle allows us to characterize the performance of the CMOS separate from the SiPh chip. Figures 3.20 (b) and (c) show closeups of the 1 mil gold wires used to connect the CMOS pads and the PCB.

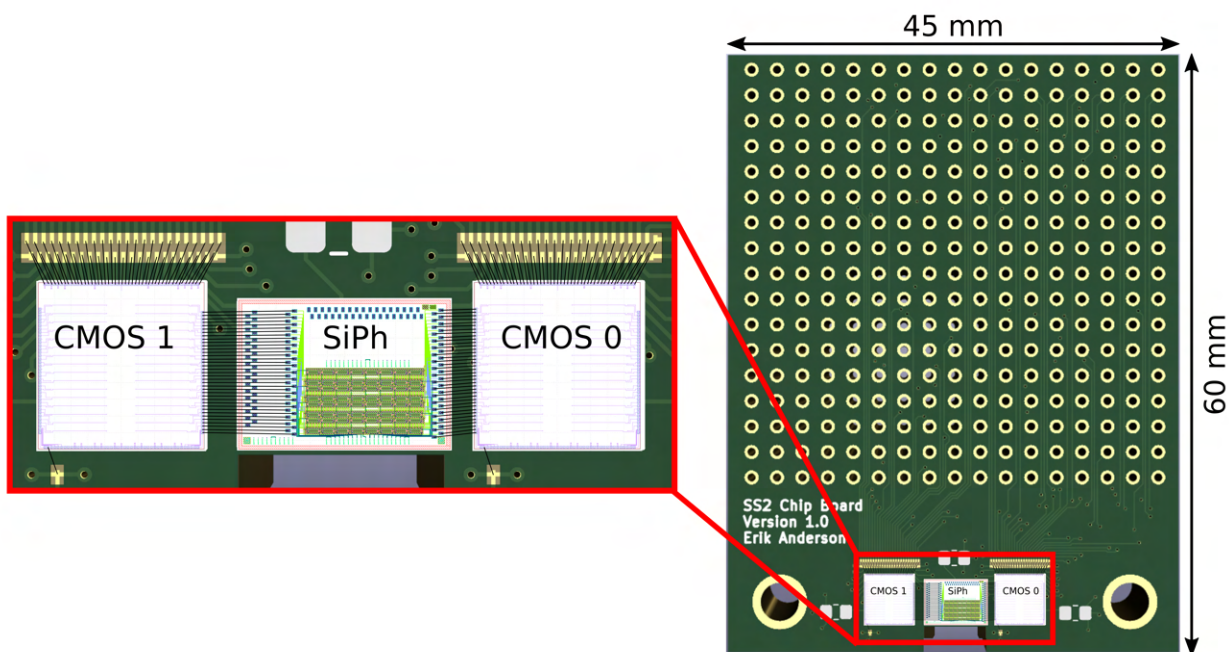


Figure 3.19: Wirebond diagram for SuperSwitch 2 using a 3-chip package for controlling the 8x8 OCS.




Figure 3.20: CMOS only test package with (a) CMOS 1, the west side control chip from Fig. 3.19, (b) side-view of wirebonds, and (c) top-view of wirebonds.

3.4 Test Setup

The automated test setup from SuperSwitch 1, shown in Figures 2.20 and 2.21, can be easily repurposed for SuperSwitch 2. Only 64 switching elements need to be characterized making the test setup even simpler than in SuperSwitch 1. Fig. 3.21 shows a top-view of the SuperSwitch 2 test setup. Similar to SuperSwitch 1, an FPGA is used to drive signals

on a host board that contains a standard PGA socket. The chip board, containing only CMOS chiplets for now, is plugged into the PGA socket for testing. The FPGA is used to communicate to the CMOS chiplets in a similar to what is shown in Fig. 2.19. The FPGA image is updated to reflect the fact that SuperSwitch 2 only has 3 scan chains per chiplet (1x control chain, 2x select chain) as opposed to the 9 scan chains per chiplet for SuperSwitch 1. An electrical probe, shown in the lower left of Fig. 3.21, is used to probe the outputs of each cell controller to characterize the CMOS performance.

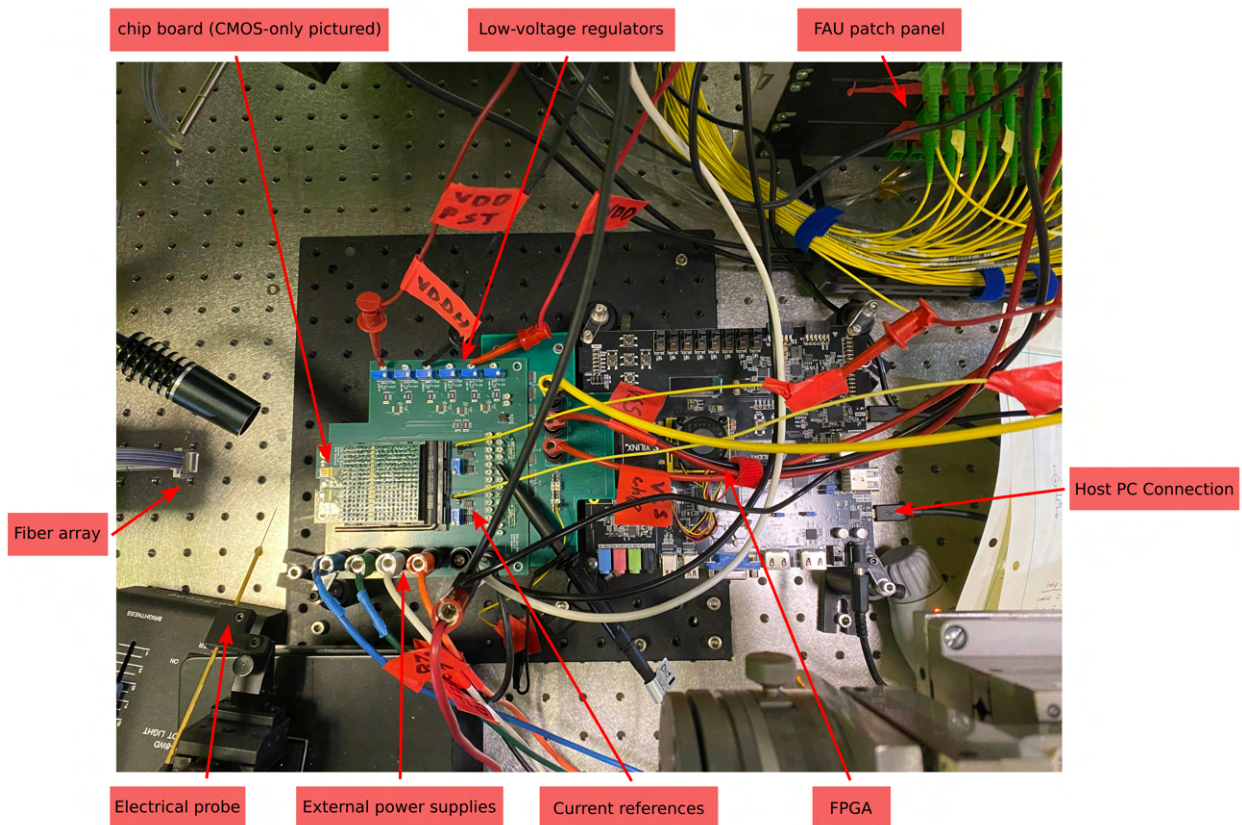


Figure 3.21: The test setup for SuperSwitch 2.

3.5 Results

The static performance of the CMOS has been characterized using the CMOS-only wirebond package shown in both Figures 3.19 and 3.21.

DNL/INL Calculation

Fig. 3.22 (a) shows the measured output voltages for a single cell controller for all 64 DAC codes. The chip is provided a reference current of $\approx 3.13\mu\text{A}$ to its global current reference circuit. The cell controller's IREFP circuit is configured to set the IDAC LSB current to the same value. Fig. 3.22 (a) also shows the ideal and simulated output voltages in blue and orange respectively. The ideal voltages are calculated from the Equations detailed in Section 3.2 which assume that I_{BIAS} is set to exactly $4I_{LSB}$ and that the current mirror ratio through the HVDRIVER circuit is always $M_{mirror} = 10$ no matter the bias point. The simulated output voltage takes into account the non-ideal current mirror ratio as well as the non-ideal bias current. The DNL and INL calculations are done for the measured and simulated results relative to the ideal values.

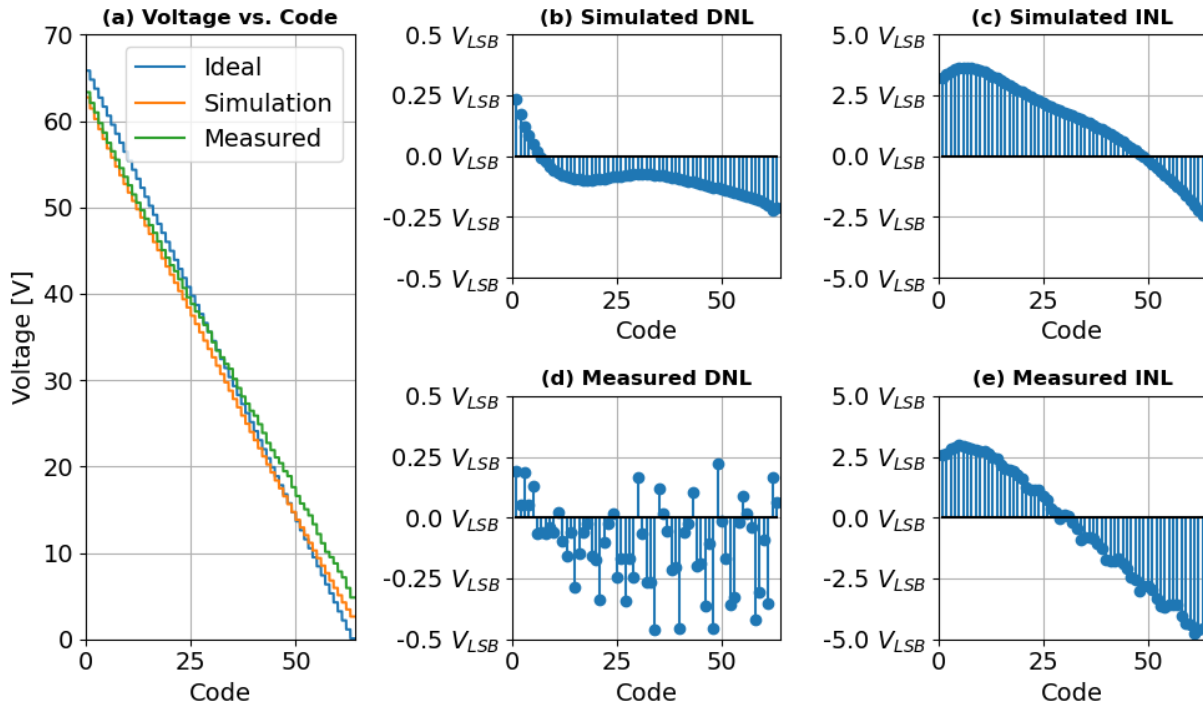


Figure 3.22: Simulated and measured DNL and INL.

Figures 3.22 (b) and (d) show the high voltage DAC's simulated and measured differential nonlinearity (DNL). Both plots show a similar trend of mostly negative DNL. This can be explained by a mirror ratio of less than $M_{mirror} = 10$ such that the change in output voltage for each pair of adjacent DAC codes is smaller in amplitude than the ideal value. Despite this, and more importantly, the maximum DNL for the measured circuit is shown to be less than 1 LSB, proving that the design is indeed monotonic (fully expected as described in Section 3.2). The accuracy of each HV DAC is not as important as the full scale range

and monotonicity. Each cell controller can be programmed with independent selected and deselected waveforms (see Table 3.8) as well as different LSB currents (see Table 3.7) to obtain output voltages as close as possible to the ones plotted in Fig. 3.6 for a given target displacement. To calculate these DNL values, the following procedure was followed:

1. Calculate the ideal LSB voltage (using I_{LSB} from Eq. 3.9):

$$V_{LSB} = -I_{LSB}M_{mirror}R_L \approx -1.04V \quad (3.10)$$

2. Calculate DNL for DAC codes $k = \{1, 2, \dots, 63\}$:

$$DNL[k] = \frac{V[k] - V[k-1]}{V_{LSB}} - 1 \quad (3.11)$$

The simulated and measured integral nonlinearity (INL) plots are shown in Figures 3.22 (c) and (e). An initial INL of about 2.5 LSB is observed for both plots. This is due to the nonideal bias current mirroring causing our input bias current to be greater than the expected $4I_{LSB}$. Unfortunately, both the bias current for the high voltage driver and the LSB current for the current-steering DAC are set by the same static current DAC (i.e. circuit from Fig. 3.14 sets both currents). This means that the maximum output voltage, $V_{out-max}$, cannot be set independently from the LSB voltage. Future designs could include two static current DACs to set these values separately. This would help achieve more uniform $V_{out-max}$ values and full scale ranges across cell controllers. The following procedure is done to calculate the integral nonlinearity (INL) of the high voltage DAC:

1. Calculate the ideal maximum output voltage:

$$V_{out-max} = HVDD - (F_{BIAS} * R_L * M_{mirror} * I_{LSB}) \approx 65.8V \quad (3.12)$$

2. Calculate INL for DAC codes $k = \{1, 2, \dots, 63\}$:

$$INL[k] = \frac{V[0] - V_{out-max}}{V_{LSB}} + \sum_{i=1}^k DNL[i] \quad (3.13)$$

Expected Response

The expected response can be calculated using the measured values plotted in Fig. 3.22 (a) and the model defined in Section 3.1. This represents the expected or predicted response of a single MEMS device if the presented HV driver was used to control it. The MEMS+CMOS packages have not yet been completed and thus the actual response cannot be measured. Fig. 3.23 (a) shows the applied voltage and displacement for ideal and measured DAC values for a target displacement of 550nm. The ideal values are taken directly from Fig. 3.6, and the measured values are the measured voltages from Fig. 3.22 (a) that are closest to the

ideal values. The voltages are applied w/ a resolution of 20ns to model the 50MHz maximum clock frequency of the CMOS chip. The back bias of the modelled switch element is assumed to be tied to the minimum HV DAC voltage, around $\approx 5V$ for this particular HV DAC. This reduces the maximum voltage by the same amount but still gives us enough range to apply $\approx 33V$, the maximum required voltage as shown by Fig. 3.6, across the switch element. Fig. 3.23 (b) shows the settling time for 2%, 5%, and 10% error tolerance. The settling time for the expected response is essentially the same as the ideal case for both 5% and 10% error tolerance. The settling time increases quite significantly for lower error tolerances as small voltage offsets can still cause small oscillations as seen by the expected displacement in Fig. 3.23 (a). 5% settling time is likely a good enough metric as Fig. 3.4 shows $< 0.5dB$ difference between 95% (523nm) and 100% of the target displacement (550nm).




Figure 3.23: (a) Ideal voltage vs. measured DAC voltages. (b) Ideal displacement and displacement using measured DAC values.

Power Consumption

The typical power consumption for the SuperSwitch 2 8x8 CMOS control chiplet is dominated by two supplies:

1. High voltage supply (HVDD): 8 cell controllers draw maximum current through HVDD, while 56 cell controllers draw just the bias current.
2. Low voltage analog supply (VDDH): All 64 cell controllers draw the same current, $63I_{LSB}$, through their current-steering DACs at all times. The nature of the current-steering DAC requires that the same amount of current is always drawn through the DACs regardless of their input DAC codes.

These currents can be easily estimated using the equations presented in Section 3.2. The high voltage supply, HVDD, can be adjusted to minimize the total power consumption at

the cost of reduced maximum output voltage. Fig. 3.24 (a) plots the estimated power for HVDD=50V and HVDD=70V next to the measured power from the package shown in Fig. 3.21. 8 cell controllers are set to maximum code while the remaining 56 are set to minimum code. This corresponds to the typical power consumption for a given HVDD. 70V is used as this is the maximum operating voltage and thus represents the maximum typical power consumption. 50V is used as this is the minimum voltage that can be used to obtain the output DAC voltages needed to implement a dual-step response with a target displacement of 550nm (see Fig. 3.23 (a)).




Figure 3.24: (a) Estimated and measured power for an HVDD of 50V and 70V. (b) The estimated average power per high voltage DAC (HVDAC) vs. radix.

The measured data in Fig. 3.24 (a) also contains the power for the 1.8V digital logic supply, VDD, and the 5V IO supply, VDDPST. Seeing as how neither supply's contribution can be seen, it's obvious that the power consumption of these supplies is insignificant compared to that of VDDH and HVDD. The total measured power is $\approx 30\%$ more than the estimate for HVDD=50V and $\approx 20\%$ more than the estimate for HVDD=70V. These percentages are used to adjust the estimated total power for the plot shown in Fig. 3.24 (b). This plot shows the estimated power per cell controller (i.e. per HV DAC) versus the radix of the switch. The power per cell drops dramatically as the radix is increased from 2 to 64 and then begins to saturate around 6.8mW for HVDD=50V and 11.8mW for HVDD=70V. At low radices, the total switch power is dominated by the cell controllers consuming the maximum amount of HVDD current (i.e. the *selected* cell controllers of which there should only ever be k for a $k \times k$ OCS). At higher radices, the power consumption of the *deselected* cell controllers will begin to dominate the total power. This is because the number of *selected* cell controllers scales linearly with the radix ($num_{selected} = k$) while the number of *deselected* cell controllers scales with the square of the radix ($num_{deselected} = k^2 - k$). Once $k^2 - k \gg k$, the power per cell controller will be approximately equal to the power of a single *deselected* cell controller (i.e. the minimum power for a single cell controller).

The power per cell values in Fig. 3.24 (b) can be directly compared to the power consumption of amplifier-based HV DACs, such as those in [1, 29, 28]. These types of HV DACs will have nearly all of their power scale with the square of the radix. Put more simply, any additional amplifier-based HV DAC will draw the same amount of power regardless of the output voltage (at least for the capacitive loads that must be driven for this OCS). This will be discussed in more detail in Chapter 4.

Chapter 4

Conclusion

4.1 Final Results

Table 4.1 compares SuperSwitch1 to Google’s Palomar OCS [45] and Polatis’ 576x576 piezo-based OCS [17]. Both of these switches represent the most advanced OCS products currently in use. A 240x240 digital SiPh MEMS switch is also included in the table as it represents the current lowest loss OCS that uses essentially the same architecture as SuperSwitch1 [37]. Both SuperSwitch1 and the OCS from [37] suffer from approximately 9dB of coupling loss. This is due to the use of vertical grating couplers, lossy couplers chosen for their testability in spite of their high insertion loss. More efficient couplers ($< 0.1\text{dB}/\text{coupler}$) have been proposed [27] and could be used in a loss-optimized design. SuperSwitch1 also suffers from an additional 3.8dB of loss due to the 32x32 OCS being located in the upper left quadrant of a larger 128x128 OCS. This adds unnecessary waveguide and MMI crossing loss. Unfortunately, even if these additional losses are ignored, the on-chip loss of both SiPh MEMS switches is much higher than the total loss from the piezo and 3D MEMS switches. This fact makes it clear that the insertion loss must be lowered before these OCSs can be introduced into datacenters and HPC clusters. Despite this, the potential of these SiPh MEMS OCSs remains strong as SuperSwitch1 demonstrated orders of magnitude faster switching at an incredibly low static power consumption.

As the SuperSwitch2 SiPh MEMS chip has yet to be tested, Table 2 4.2 compares the driver IC against other similar high-voltage drivers. Alameh et al [2] uses a digitally-reconfigurable charge-pump based driver designed specifically for MEMS capacitive loads. This driver burns much less power than SuperSwitch2 but is far slower for the same load capacitance of 1pF. Additionally, this particular driver can only reach a maximum voltage of 10V, much too low to control the SuperSwitch2 MEMS devices. Other charge-pump based MEMS drivers have been proposed to reach higher voltages [4] and faster rise times [35], but have yet to be fabricated and tested. It remains to be seen whether or not these types of drivers can reach sub-microsecond rise and fall times at the voltages required by SuperSwitch2. Ning et al [28] details a high voltage amplifier based design that achieves a low

Table 4.1: Comparison table for optical circuit switches.

	SuperSwitch1 (this work)	Google [45]	Polatis [17]	OFC 2019 (Seok et al [37])
Architecture	SiPh MEMS	3D MEMS Mirror	Piezo	SiPh MEMS
Ports	32x32	136x136	572x572	240x240
Max Insertion Loss	23dB	2dB	3dB	18.8dB
Static Power Consumption	3.61mW	108W*	175W*	n/a

*Power reported for full system. Includes power for misc items (e.g. MCUs, OPMs, etc...).

power consumption w/ a moderately fast rise/fall time ($\approx 5\mu\text{s}$ rise/fall time for 1pF given the reported slew rate). While the SuperSwitch2 IC has a much higher maximum power consumption (178.7mW vs 10.95mW), the majority of the HV driver cells will operate near the minimum power consumption when operating a $k \times k$ OCS of the style presented in Fig. 3.2 (a). For larger values of k , as evidenced by the plot in Fig. 3.24 (b), the average power consumption of each HV driver becomes comparable to the power consumption of the high voltage amplifier in [28] (11.8mW vs. 10.95mW). This allows the SuperSwitch2 IC to burn a similar amount of total power to a high-voltage amplifier based design, while achieving much faster rise and fall times for the same load capacitance (122ns/136ns vs. 5 μs).

Table 4.2: Comparison table for HV driver circuits.

	SuperSwitch2 (this work)	TVLSI 2017 (Alameh et al[2])	DDECS 2014 (Ning et al[28])
Technology	180nm	130nm	350nm
Architecture	current-steering DAC	charge pump	LV DAC + HV amplifier
Max Output Voltage	64V	10.1V	120V
DAC Resolution	6 bits	< 6 bits*	8 bits
Output Load	1pF	1pF	20pF
Max Rise/Fall Time	122ns/136ns	7 μ s/611 μ s	100 μ s
Output Channels	64	1	16
Power/Channel	11.8mW [†] /178.7mW [§]	252 μ W [†] /864 μ W [§]	10.95mW

*Average resolution across FSR. † Min power consumption. § Max power consumption.

4.2 Thesis Contribution

This thesis proposed and implemented two novel designs for controlling SiPh MEMS optical circuit switches. SuperSwitch1 demonstrated a low power and scalable architecture that can be used as a blueprint for future digital SiPh MEMS switches. Additionally, SuperSwitch1’s 3D-packaged 32x32 OCS proved the feasibility of crossbar style OCSs, demonstrating individual control of 1024 MEMS devices and providing a pathway to control up to 16,384 (128 * 128). SuperSwitch2 provides a similar blueprint but for SiPh MEMS OCSs that require *analog* high voltage control. The SuperSwitch2 IC implements custom digital logic combined with a current-steering high voltage DAC, allowing for fast control of crossbar SiPh MEMS switches while limiting the power consumption per MEMS device.

4.3 Future of SiPh MEMS OCS

Both SuperSwitch1 and SuperSwitch2 were designed with future fast-switching applications in mind. Table 4.1 shows that digital SiPh MEMS (i.e. SuperSwitch1) greatly outperforms the incumbent OCS technologies when it comes to switch reconfiguration times and static control power. As HPC clusters continue to scale to support larger and larger AI models, the importance of fast and low-power OCSs will also grow. Unfortunately, even the lowest loss SiPh MEMS designs, such as the 240x240 switch in [37], are nowhere close to where they need to be in terms of total insertion loss to displace the current leading technologies. Work has already been done to identify ways to minimize the loss and realize the full potential

of these SiPh MEMS switches [44] [12] [39]. Future work will look to incorporate the best of these solutions, alongside the CMOS control innovations demonstrated in this thesis, to provide a truly competitive OCS platform.

Bibliography

- [1] Lasse Aaltonen, Mikko Saukoski, and Kari Halonen. “On-chip Digitally Tunable High Voltage Generator for Electrostatic Control of Micromechanical Devices”. In: *IEEE Custom Integrated Circuits Conference 2006*. IEEE Custom Integrated Circuits Conference 2006. ISSN: 2152-3630. Sept. 2006, pp. 583–586. DOI: 10.1109/CICC.2006.320825. URL: <https://ieeexplore.ieee.org/document/4115027> (visited on 11/05/2024).
- [2] Abdul Hafiz Alameh and Frederic Nabki. “A 0.13- μm CMOS Dynamically Reconfigurable Charge Pump for Electrostatic MEMS Actuation”. In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 25.4 (Apr. 2017). Conference Name: IEEE Transactions on Very Large Scale Integration (VLSI) Systems, pp. 1261–1270. ISSN: 1557-9999. DOI: 10.1109/TVLSI.2016.2629439. URL: <https://ieeexplore.ieee.org/document/7776924> (visited on 11/15/2024).
- [3] José Roberto de Almeida Amazonas, Germán Santos-Boada, and Josep Solé-Pareta. “Who shot Optical Packet switching?” In: *2017 19th International Conference on Transparent Optical Networks (ICTON)*. 2017 19th International Conference on Transparent Optical Networks (ICTON). ISSN: 2161-2064. July 2017, pp. 1–4. DOI: 10.1109/ICTON.2017.8025164. URL: <https://ieeexplore.ieee.org/document/8025164> (visited on 11/18/2024).
- [4] Philippe-Olivier Beaulieu et al. “A 360 V high voltage reconfigurable charge pump in 0.8 μm CMOS for optical MEMS applications”. In: *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*. 2016 IEEE International Symposium on Circuits and Systems (ISCAS). ISSN: 2379-447X. May 2016, pp. 1630–1633. DOI: 10.1109/ISCAS.2016.7538878. URL: <https://ieeexplore.ieee.org/document/7538878> (visited on 11/14/2024).
- [5] John E. Bowers and Alan Y. Liu. “A comparison of four approaches to photonic integration”. In: *2017 Optical Fiber Communications Conference and Exhibition (OFC)*. 2017 Optical Fiber Communications Conference and Exhibition (OFC). Mar. 2017, pp. 1–3. URL: <https://ieeexplore.ieee.org/document/7936791> (visited on 11/19/2024).
- [6] Broadcom. *Ethernet Network Switches — Data Center Switches — Tomahawk 5*. URL: <https://www.broadcom.com/products/ethernet-connectivity/switching/strataxgs/bcm78900-series> (visited on 10/28/2024).

- [7] Dritan Celo et al. “32×32 silicon photonic switch”. In: *2016 21st OptoElectronics and Communications Conference (OECC) held jointly with 2016 International Conference on Photonics in Switching (PS)*. 2016 21st OptoElectronics and Communications Conference (OECC) held jointly with 2016 International Conference on Photonics in Switching (PS). July 2016, pp. 1–3. URL: <https://ieeexplore.ieee.org/document/7718577> (visited on 11/19/2024).
- [8] Qixiang Cheng et al. “Photonic switching in high performance datacenters [Invited]”. In: *Optics Express* 26.12 (June 11, 2018). Publisher: Optica Publishing Group, pp. 16022–16043. ISSN: 1094-4087. DOI: 10.1364/OE.26.016022. URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-26-12-16022> (visited on 11/13/2024).
- [9] Nicolas Dupuis et al. “Design and Fabrication of Low-Insertion-Loss and Low-Crosstalk Broadband 2\textimes 2 Mach-Zehnder Silicon Photonic Switches”. In: *Journal of Lightwave Technology* 33.17 (Sept. 2015). Conference Name: Journal of Lightwave Technology, pp. 3597–3606. ISSN: 1558-2213. DOI: 10.1109/JLT.2015.2446463. URL: <https://ieeexplore.ieee.org/document/7126923> (visited on 11/19/2024).
- [10] Nesbitt Hagood et al. “Beam-steering optical switching apparatus”. U.S. pat. Continuum Photonics Inc. Feb. 10, 2005. URL: <https://patents.google.com/patent/US20050030840A1/en> (visited on 11/14/2024).
- [11] Sangyoon Han et al. “32 × 32 silicon photonic MEMS switch with gap-adjustable directional couplers fabricated in commercial CMOS foundry”. In: *Journal of Optical Microsystems* 1.2 (Mar. 2021). Publisher: SPIE, p. 024003. ISSN: 2708-5260, 2708-5260. DOI: 10.1117/1.JOM.1.2.024003. URL: <https://www.spiedigitallibrary.org/journals/journal-of-optical-microsystems/volume-1/issue-2/024003/32--32-silicon-photonic-MEMS-switch-with-gap-adjustable/10.1117/1.JOM.1.2.024003.full> (visited on 10/22/2024).
- [12] Amirmahdi Honardoost et al. “Low-Loss Wafer-Bonded SiPh MEMS Switches”. In: *2022 Optical Fiber Communications Conference and Exhibition (OFC)*. 2022 Optical Fiber Communications Conference and Exhibition (OFC). Mar. 2022, pp. 1–3. URL: <https://ieeexplore.ieee.org/document/9748536> (visited on 10/31/2024).
- [13] Yishen Huang et al. “Multi-Stage 8 × 8 Silicon Photonic Switch Based on Dual-Microring Switching Elements”. In: *Journal of Lightwave Technology* 38.2 (Jan. 2020). Conference Name: Journal of Lightwave Technology, pp. 194–201. ISSN: 1558-2213. DOI: 10.1109/JLT.2019.2945941. URL: <https://ieeexplore.ieee.org/document/8861317> (visited on 11/19/2024).
- [14] How Yuan Hwang et al. “Flip Chip Packaging of Digital Silicon Photonics MEMS Switch for Cloud Computing and Data Centre”. In: *IEEE Photonics Journal* 9.3 (June 2017). Conference Name: IEEE Photonics Journal, pp. 1–10. ISSN: 1943-0655. DOI: 10.1109/JPHOT.2017.2704097. URL: <https://ieeexplore.ieee.org/document/7927713> (visited on 10/22/2024).

- [15] Matthias Imboden et al. “High-Speed Control of Electromechanical Transduction: Advanced Drive Techniques for Optimized Step-and-Settle Response of MEMS Micromirrors”. In: *IEEE Control Systems Magazine* 36.5 (Oct. 2016). Conference Name: IEEE Control Systems Magazine, pp. 48–76. ISSN: 1941-000X. DOI: 10.1109/MCS.2016.2584338. URL: <https://ieeexplore.ieee.org/document/7569123> (visited on 10/31/2024).
- [16] Richard A. Jensen, Nick Parsons, and Rohit Kunjappa. “All-Optical Switching: Past, Present and Future”. In: *2023 Optical Fiber Communications Conference and Exhibition (OFC)*. 2023 Optical Fiber Communications Conference and Exhibition (OFC). Mar. 2023, pp. 1–3. DOI: 10.1364/OFC.2023.M4J.1. URL: <https://ieeexplore.ieee.org/document/10116206> (visited on 11/14/2024).
- [17] joel webster joel. *H+S Polatis - All-Optical SDN enabled Switches - Highest Performance, Lowest Loss, Configurable from 8x8 to 576x576 ports*. URL: <http://www.polatis.com/> (visited on 11/18/2024).
- [18] Norm Jouppi et al. “TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings”. In: *Proceedings of the 50th Annual International Symposium on Computer Architecture*. ISCA '23. New York, NY, USA: Association for Computing Machinery, June 17, 2023, pp. 1–14. ISBN: 9798400700958. DOI: 10.1145/3579371.3589350. URL: <https://dl.acm.org/doi/10.1145/3579371.3589350> (visited on 10/28/2024).
- [19] Mehrdad Khani et al. “SiP-ML: high-bandwidth optical network interconnects for machine learning training”. In: *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. SIGCOMM '21. New York, NY, USA: Association for Computing Machinery, Aug. 9, 2021, pp. 657–675. ISBN: 978-1-4503-8383-7. DOI: 10.1145/3452296.3472900. URL: <https://dl.acm.org/doi/10.1145/3452296.3472900> (visited on 11/19/2024).
- [20] J. Kim et al. “1100 x 1100 port MEMS-based optical crossconnect with 4-dB maximum loss”. In: *IEEE Photonics Technology Letters* 15.11 (Nov. 2003), pp. 1537–1539. ISSN: 1041-1135, 1941-0174. DOI: 10.1109/LPT.2003.818653. URL: <http://ieeexplore.ieee.org/document/1237580/> (visited on 11/13/2024).
- [21] Ikuo Kurachi et al. “Intelligent Three-Dimensional Chip-Stacking Process for Pixel Detectors for High Energy Physics Experiments”. In: *Proceedings of the 29th International Workshop on Vertex Detectors (VERTEX2020)*. Vol. 34. JPS Conference Proceedings 34. Journal of the Physical Society of Japan, June 9, 2021. DOI: 10.7566/JPSCP.34.010010. URL: <https://journals.jps.jp/doi/abs/10.7566/JPSCP.34.010010> (visited on 11/06/2024).
- [22] Kyungmok Kwon et al. “128×128 Silicon Photonic MEMS Switch with Scalable Row Column Addressing”. In: *Conference on Lasers and Electro-Optics (2018), paper SF1A.4*. CLEO: Science and Innovations. Optica Publishing Group, May 13, 2018, SF1A.4. DOI: 10.1364/CLEO_SI.2018.SF1A.4. URL: https://opg.optica.org/abstract.cfm?uri=CLEO_SI-2018-SF1A.4 (visited on 10/22/2024).

- [23] Cedric Lam, Xiang Zhou, and Hong Liu. “200G per Lane for beyond 400GbE”. In: ().
- [24] Wenzhe Li et al. “Fast and scalable all-optical network architecture for distributed deep learning”. In: *Journal of Optical Communications and Networking* 16.3 (Mar. 2024). Conference Name: Journal of Optical Communications and Networking, pp. 342–357. ISSN: 1943-0639. DOI: 10.1364/JOCN.511696. URL: <https://ieeexplore.ieee.org/document/10444506> (visited on 11/19/2024).
- [25] Xin Li et al. “Ultra-low-loss multi-layer 8×8 microring optical switch”. In: *Photonics Research* 11.5 (May 1, 2023). Publisher: Optica Publishing Group, pp. 712–723. ISSN: 2327-9125. DOI: 10.1364/PRJ.479499. URL: <https://opg.optica.org/prj/abstract.cfm?uri=prj-11-5-712> (visited on 11/19/2024).
- [26] Hiroyuki Matsuura et al. “Accelerating switching speed of thermo-optic MZI silicon-photon switches with ”turbo pulse” in PWM control”. In: *2017 Optical Fiber Communications Conference and Exhibition (OFC)*. 2017 Optical Fiber Communications Conference and Exhibition (OFC). Mar. 2017, pp. 1–3. URL: <https://ieeexplore.ieee.org/document/7937475> (visited on 11/19/2024).
- [27] Andrew Michaels and Eli Yablonovitch. “Inverse design of near unity efficiency perfectly vertical grating couplers”. In: *Optics Express* 26.4 (Feb. 19, 2018). Publisher: Optica Publishing Group, pp. 4766–4779. ISSN: 1094-4087. DOI: 10.1364/OE.26.004766. URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-26-4-4766> (visited on 11/18/2024).
- [28] Jing Ning and Klaus Hofmann. “A 120V high voltage DAC array for a tunable antenna in communication system”. In: *17th International Symposium on Design and Diagnostics of Electronic Circuits & Systems*. 17th International Symposium on Design and Diagnostics of Electronic Circuits & Systems. Apr. 2014, pp. 65–70. DOI: 10.1109/DDECS.2014.6868765. URL: <https://ieeexplore.ieee.org/document/6868765> (visited on 10/22/2024).
- [29] Jing Ning and Klaus Hofmann. “A integrated high voltage controller for a reconfigurable antenna array”. In: *2013 NORCHIP*. 2013 NORCHIP. Nov. 2013, pp. 1–4. DOI: 10.1109/NORCHIP.2013.6702004. URL: <https://ieeexplore.ieee.org/document/6702004> (visited on 10/22/2024).
- [30] Gary O’Brien, David J. Monk, and Liwei Lin. “MEMS cantilever beam electrostatic pull-in model”. In: *Design, Characterization, and Packaging for MEMS and Microelectronics II*. Design, Characterization, and Packaging for MEMS and Microelectronics II. Vol. 4593. SPIE, Nov. 19, 2001, pp. 31–41. DOI: 10.1117/12.448834. URL: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/4593/0000/MEMS-cantilever-beam-electrostatic-pull-in-model/10.1117/12.448834.full> (visited on 10/31/2024).
- [31] AIM Photonics. *Base Active PIC*. <https://www.aimphotonics.com/base-active-pic>.

- [32] Leon Poutievski et al. “Jupiter evolving: transforming google’s datacenter network via optical circuit switches and software-defined networking”. In: *Proceedings of the ACM SIGCOMM 2022 Conference*. SIGCOMM ’22. New York, NY, USA: Association for Computing Machinery, Aug. 22, 2022, pp. 66–85. ISBN: 978-1-4503-9420-8. DOI: 10.1145/3544216.3544265. URL: <https://dl.acm.org/doi/10.1145/3544216.3544265> (visited on 10/28/2024).
- [33] Lei Qiao, Weijie Tang, and Tao Chu. “32 × 32 silicon electro-optic switch with built-in monitors and balanced-status units”. In: *Scientific Reports* 7.1 (Feb. 9, 2017). Publisher: Nature Publishing Group, p. 42306. ISSN: 2045-2322. DOI: 10.1038/srep42306. URL: <https://www.nature.com/articles/srep42306> (visited on 11/19/2024).
- [34] Behzad Razavi. “The Current-Steering DAC [A Circuit for All Seasons]”. In: *IEEE Solid-State Circuits Magazine* 10.1 (2018), pp. 11–15. ISSN: 1943-0582, 1943-0590. DOI: 10.1109/MSSC.2017.2771102. URL: <https://ieeexplore.ieee.org/document/8275558/> (visited on 11/04/2024).
- [35] Ahmed Saeed, Sameh Ibrahim, and Hani Fikry Ragai. “A Sizing Methodology for Rise-Time Minimization of Dickson Charge Pumps With Capacitive Loads”. In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 64.10 (Oct. 2017). Conference Name: IEEE Transactions on Circuits and Systems II: Express Briefs, pp. 1202–1206. ISSN: 1558-3791. DOI: 10.1109/TCSII.2017.2687864. URL: <https://ieeexplore.ieee.org/document/7887674> (visited on 11/15/2024).
- [36] Stephen D. Senturia. *Microsystem Design*. 1st ed. 2001. New York, NY: Imprint: Springer, 2001. 1 p. ISBN: 978-0-306-47601-3. DOI: 10.1007/b117574.
- [37] Tae Joon Seok et al. “240×240 Wafer-Scale Silicon Photonic Switches”. In: *Optical Fiber Communication Conference (OFC) 2019 (2019), paper Th1E.5*. Optical Fiber Communication Conference. Optica Publishing Group, Mar. 3, 2019, Th1E.5. DOI: 10.1364/OFC.2019.Th1E.5. URL: <https://opg.optica.org/abstract.cfm?uri=OFC-2019-Th1E.5> (visited on 10/22/2024).
- [38] Tae Joon Seok et al. “Large-scale broadband digital silicon photonic switches with vertical adiabatic couplers”. In: *Optica* 3.1 (Jan. 20, 2016). Publisher: Optica Publishing Group, pp. 64–70. ISSN: 2334-2536. DOI: 10.1364/OPTICA.3.000064. URL: <https://opg.optica.org/optica/abstract.cfm?uri=optica-3-1-64> (visited on 10/22/2024).
- [39] Mizuki Shirao et al. “High Efficiency Double Layer Grating Couplers Supporting Polarization Diversity for Photonic Switches”. In: *26th Optoelectronics and Communications Conference (2021), paper W1E.4*. Optoelectronics and Communications Conference. Optica Publishing Group, July 3, 2021, W1E.4. DOI: 10.1364/OECC.2021.W1E.4. URL: <https://opg.optica.org/abstract.cfm?uri=OECC-2021-W1E.4> (visited on 11/18/2024).

- [40] Matthew Spencer. “Design Considerations for Nano-Electromechanical Relay Circuits”. PhD thesis. Aug. 14, 2015. URL: <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-195.pdf>.
- [41] T. Tanzawa et al. “High-voltage transistor scaling circuit techniques for high-density negative-gate channel-erasing NOR flash memories”. In: *IEEE Journal of Solid-State Circuits* 37.10 (Oct. 2002). Conference Name: IEEE Journal of Solid-State Circuits, pp. 1318–1325. ISSN: 1558-173X. DOI: 10.1109/JSSC.2002.803045. URL: <https://ieeexplore.ieee.org/document/1035946> (visited on 10/28/2024).
- [42] Min Yee Teh et al. “Enabling Quasi-Static Reconfigurable Networks With Robust Topology Engineering”. In: *IEEE/ACM Trans. Netw.* 31.3 (Oct. 10, 2022), pp. 1056–1070. ISSN: 1063-6692. DOI: 10.1109/TNET.2022.3210534. URL: <https://dl.acm.org/doi/10.1109/TNET.2022.3210534> (visited on 11/18/2024).
- [43] Wenjing Tian et al. “Hybrid Photonic Integration for Optical Switches”. In: *2023 International Conference on Photonics in Switching and Computing (PSC)*. 2023 International Conference on Photonics in Switching and Computing (PSC). ISSN: 2166-8892. Sept. 2023, pp. 1–3. DOI: 10.1109/PSC57974.2023.10297289. URL: <https://ieeexplore.ieee.org/document/10297289> (visited on 11/19/2024).
- [44] Jean-Etienne Tremblay, Johannes Henriksson, and Ming C. Wu. “Polarization-Diversity Evanescent Coupler on Silicon with Integrated Polarization Splitter”. In: *2020 IEEE Photonics Conference (IPC)*. 2020 IEEE Photonics Conference (IPC). ISSN: 2575-274X. Sept. 2020, pp. 1–2. DOI: 10.1109/IPC47351.2020.9252558. URL: <https://ieeexplore.ieee.org/abstract/document/9252558> (visited on 11/18/2024).
- [45] Ryohei Urata et al. *Mission Apollo: Landing Optical Circuit Switching at Datacenter Scale*. Aug. 22, 2022. DOI: 10.48550/arXiv.2208.10041. arXiv: 2208.10041. URL: <http://arxiv.org/abs/2208.10041> (visited on 10/28/2024).
- [46] Weiyang Wang et al. “{TopoOpt}: Co-optimizing Network Topology and Parallelization Strategy for Distributed Training Jobs”. In: 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23). 2023, pp. 739–767. ISBN: 978-1-939133-33-5. URL: <https://www.usenix.org/conference/nsdi23/presentation/wang-weiyang> (visited on 11/18/2024).
- [47] Ke Wen et al. “Flexfly: Enabling a Reconfigurable Dragonfly through Silicon Photonics”. In: *SC '16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. SC '16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. ISSN: 2167-4337. Nov. 2016, pp. 166–177. DOI: 10.1109/SC.2016.14. URL: <https://ieeexplore.ieee.org/document/7877093> (visited on 11/18/2024).
- [48] Junfei Xia et al. “The Feasibility of Building 1024 & 4096-port Nanosecond Switching for Data Centre Networks Using Dilated Hybrid Optical Switches”. In: ().

- [49] Jing Zhang et al. “Lossless High-speed Silicon Photonic MZI switch with a Micro-Transfer-Printed III-V amplifier”. In: *2022 IEEE 72nd Electronic Components and Technology Conference (ECTC)*. 2022 IEEE 72nd Electronic Components and Technology Conference (ECTC). ISSN: 2377-5726. May 2022, pp. 441–445. DOI: 10.1109/ECTC51906.2022.00077. URL: <https://ieeexplore.ieee.org/document/9816408> (visited on 11/19/2024).
- [50] Wen-Ming Zheng et al. “Capacitive floating level shifter: Modeling and design”. In: *TENCON 2015 - 2015 IEEE Region 10 Conference*. TENCON 2015 - 2015 IEEE Region 10 Conference. ISSN: 2159-3450. Nov. 2015, pp. 1–6. DOI: 10.1109/TENCON.2015.7373013. URL: <https://ieeexplore.ieee.org/document/7373013> (visited on 10/28/2024).