Demonstration of Integrated Micro-Electro-Mechanical Relay Circuits for VLSI Applications

Matthew Spencer, Fred Chen, Cheng C. Wang, *Student Member, IEEE*, Rhesa Nathanael, *Student Member, IEEE*, Hossein Fariborzi, *Student Member, IEEE*, Abhinav Gupta, Hei Kam, *Member, IEEE*, Vincent Pott, *Member, IEEE*, Jaeseok Jeon, Tsu-Jae King Liu, *Fellow, IEEE*, Dejan Marković, *Member, IEEE*, Elad Alon, *Member, IEEE*, and Vladimir Stojanović, *Member, IEEE*

***Abstract*—This work presents measured results from test chips containing circuits implemented with micro-electro-mechanical (MEM) relays. The relay circuits designed on these test chips illustrate a range of important functions necessary for the implementation of integrated VLSI systems and lend insight into circuit design techniques optimized for the physical properties of these devices. To explore these techniques a hybrid electro-mechanical model of the relays' electrical and mechanical characteristics has been developed, correlated to measurements, and then also applied to predict MEM relay performance if the technology were scaled to a 90 nm technology node. A theoretical, scaled, 32-bit MEM relay-based adder, with a single-bit functionality demonstrated by the measured circuits, is found to offer a factor of ten energy efficiency gain over an optimized CMOS adder for sub-20 MOPS throughputs at a moderate increase in area.**

***Index Terms*—Adders, digital circuits, MEM relays, microelectromechanical devices, minimum energy point, very-large-scale integration.**

I. Introduction

ALTHOUGH CMOS technology scaling has historically enabled significantly reduced energy-per-operation in integrated circuits, today's designs are increasingly power limited in ways that technology scaling cannot alleviate. This has occurred because the threshold voltage of the transistors has already been scaled to the value that optimally balances leakage energy and dynamic energy, and hence further reductions in the threshold voltage would actually increase the amount of energy consumed per operation. With the threshold voltage pinned because of sub-threshold leakage, further supply voltage scaling comes at the expense of per-core performance, forcing a trend towards increasingly parallel circuit implementations as the only means to efficiently improve throughput.

Manuscript received April 20, 2010; revised July 05, 2010; accepted July 16, 2010. Date of publication November 01, 2010; date of current version December 27, 2010. This paper was approved by Guest Editor Alison Burdett. This work was supported in part by the DARPA NEMS program, the C2S2 and MSD FCRP centers, MIT CICS, UCLA ICSL, BWRC, and NSF Infrastructure Grant No. 040327.

M. Spencer, R. Nathanael, A. Gupta, H. Kam, V. Pott, J. Jeon, T.-J. King Liu, and E. Alon are with the University of California, Berkeley, CA 94720 USA (e-mail: mespence@eecs.berkeley.edu).

F. Chen, H. Fariborzi, and V. Stojanović are with the Massachusetts Institute of Technology, Cambridge, MA 02142 USA.

C. C. Wang and D. Marković are with the University of California, Los Angeles, CA 90095 USA.

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

Digital Object Identifier 10.1109/JSSC.2010.2074370

Unfortunately, even this parallelism will eventually become ineffective as each CMOS functional unit approaches its throughput-independent, minimum achievable energy. The minimum energy in CMOS is limited by the sub-threshold leakage of the transistors because once CMOS circuits enter the sub-threshold regime, an increase in the threshold voltage decreases the leakage current by exactly the same amount that it increases the delay. The only mechanism left to tune both the leakage and dynamic energy components is therefore the supply voltage, which cannot be reduced below a certain value (set by $k_{\mathrm{B}}T/q$ and the circuit's activity factor and logic depth) without increasing the total energy [1].

If a device with significantly improved leakage characteristics (i.e., steeper sub-threshold slope) were available, major improvements in energy efficiency over CMOS could be achieved [2]. Many researchers have therefore been exploring new switching device concepts to achieve sub-threshold slopes steeper than the limit set by $k_{\mathrm{B}}T/q$ in field-effect or bipolar-junction transistors [3], [4]. However, many of these devices achieve sharp sub-threshold slope over only a limited range of supply voltage, leading to relatively poor *on*-to-*off* current ratios and/or very low *on*-state current at low supply voltages.

In this context, micro-electro-mechanical (MEM) relays appear very attractive, having recently demonstrated *on*-to-*off* current ratios of ten orders of magnitude over input swings of one millivolt and immeasurably low leakage currents [5]. These relays are four-terminal devices that are functionally similar to CMOS transistors. Despite their nearly ideal I-V characteristics, the time required to mechanically switch a relay from the *off*- to the *on*-state is significantly longer than the electrical switching delay of an equivalent CMOS device. Specifically, relays fabricated in a 90 nm technology node are predicted to have delays of 10's of nanoseconds compared to 100's of picoseconds for transistors in a threshold optimized CMOS process operating in the sub-threshold regime, or up to nanoseconds in a more generically available CMOS process. Although the large mechanical delay of MEM relays suggests that MEM relay circuits would have very poor performance, we have proposed circuit architectures that significantly mitigate this by implementing logic as large, complex gates that minimize the number of mechanical delays on the critical path [6].

To verify these circuit principles and the feasibility of the MEM relay technology, this paper describes a test chip demonstrating several functioning MEM-relay circuits [7]. Results from the test chip confirm MEM-relay delay characteristics and

demonstrate functionality of key circuit components needed for the development of an integrated VLSI system: logic, latches, memory, and I/O circuits.

To quantify the impact of device topology and dimensions on MEM circuits, scalable models of the MEM relay behavior are developed and used to drive a comparison between the predicted capabilities of scaled MEM relays and CMOS in an equivalent 90 nm technology node. A comparison of optimized 32-bit adders built from CMOS and MEM relays indicates that despite their large mechanical delay, MEM relays can achieve energy-delay characteristics that are nearly an order of magnitude better than CMOS over a wide range of frequencies, while requiring only three times as much area. This further validates the notion that MEM relays are viable candidates for future energy-efficient digital integrated circuits.

The remainder of the paper is organized as follows. In Section II we first describe the structure, operation, and modeling of the MEM relays on the test chip. In Section III we report measured results of circuits constructed from MEM relays illustrating a variety of functions. Then, in Section IV we study how this MEM relay's performance would scale with device dimensions down to a 90 nm process node, including a few possible layout optimizations. Section V uses the conclusions about scaled devices to compare MEM relay and CMOS circuit performance and show that adders similar to those that have already been implemented appear promising for use in the construction of low-power VLSI blocks. Finally, Section VI concludes the paper.

II. Structure, Operation and Modeling of a MEM Relay

In this section we describe the structure and operation of the MEM relay in order to provide background for the circuit and scaling explorations that are the focus of this paper. We also develop a hybrid electro-mechanical model of the MEM relay, which we will use for MEM relay circuit design and analysis.

A. MEM Relay Structure and Operation

Fig. 1 shows an SEM image and diagram of the MEM relay device used in this work [5]. The four-terminal device consists of a movable poly-SiGe gate structure suspended by spring-like folded flexures above the tungsten body, drain, and source electrodes. The channel consists of a strip of tungsten attached to an insulating oxide layer on the bottom of the gate. The channel and gate have vertical deformations, referred to as dimples, which define the regions where contact is made between the channel and the source and drain. To improve device reliability, a thin titanium oxide (TiO_2) coating is applied to the device to reduce current flow at the contacts and slow the formation of tungsten native oxides [8].

The basic operation and switching states of the MEM relay are also shown in Fig. 1. When a voltage is applied between the gate structure and the body electrode, the applied electrostatic force pulls against the mechanical spring force of the flexures and displaces the gate vertically. When sufficient electrostatic force is applied, the relay is turned *on* by the channel coming into contact with the source and drain electrodes and creating a

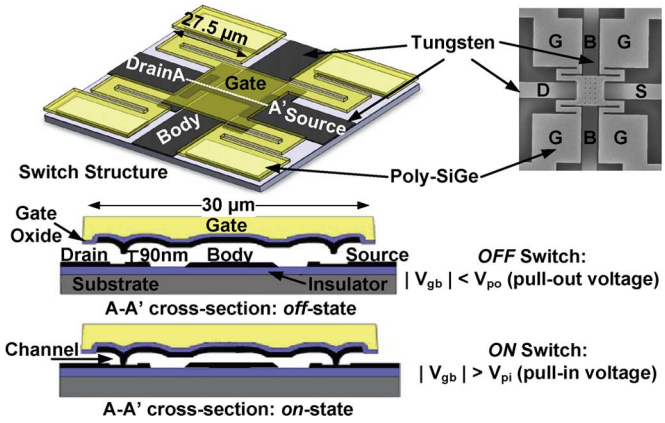

Fig. 1. SEM, diagram, and operating states of the MEM relay device.

conduction path. To prevent the structure from collapsing catastrophically, the dimples in the channel restrict the gate's motion once they make contact with the drain and source. When the gate-to-body voltage is reduced sufficiently, the springs return to their unflexed position, pulling the channel out of contact with the drain and source. These MEM relay dynamics can be accurately captured with a classical hybrid electro-mechanical model, described next, which significantly aids design intuition and enables fast circuit simulation and verification.

B. Mechanical Modeling of the MEM Relay

A non-linear second-order differential equation is used to model the movement of the gate under the applied electrostatic force [6], [8]. In this model, which is summarized in Fig. 2, the mass and flexures are modeled as a spring-mass-damper system which controls the variable resistances between channel, drain, and source

$$m\ddot{x} = F_{\text{elec}}(x) - b\dot{x} - kx \tag{1}$$

where x is the displacement of the gate, b is the damping coefficient associated with the motion of the gate structure, k is the effective spring constant of the gate structure, and $F_{\text{elec}}(x)$ is the non-linear electrical force between the gate and the body. Both the spring and damping coefficients in this dynamical model must be evaluated by finite-element simulations, but if we neglect fringing fields, the electrical force can be modeled simply as the electrostatic force between the gate and body electrodes

$$F_{\text{elec}}(x) = \frac{\varepsilon_0 A_{\text{ov}} V_{\text{gb}}^2}{2(g_0 - x)^2} \tag{2}$$

where ε_0 is the permittivity of free space, A_{ov} is the area of the overlap between the gate and body electrodes, g_0 is the normal gap between electrodes in the absence of electrical force, and V_{gb} is the voltage between the gate and the body.

During relay actuation, the spring force (kx) varies linearly with the displacement of the gate, while the electrical force is inversely quadratic in gate displacement. This results in a range of displacements where the electrical force is always larger than the spring restoring force, and in this range the difference in forces causes the device to unconditionally snap shut. It can be shown that the critical displacement at which the spring and

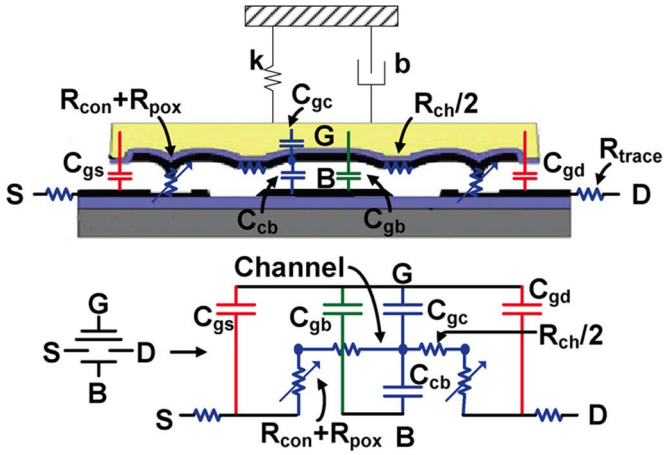

Fig. 2. Schematic indicating the relevant components in the MEM relay Verilog-A model.

electrical forces are equal is one third of the nominal gap distance. In the case that the total displacement defined by the dimples (g_d) is larger than a third of the nominal gap (g_0), the voltage necessary to snap the structure shut—called the pull-in voltage $(V_{\rm pi})$ [9]—can be derived to be

$$V_{\rm pi} = \sqrt{\frac{8}{27} \cdot \frac{k g_0^3}{\varepsilon_o A_{\rm ov}}}. \tag{3}$$

The delay required for the relay to turn *on* is governed by the relay dynamics described in (1). Similar to the reduced delay from increasing gate overdrive in CMOS, a larger applied gate-to-body voltage $(V_{\rm gb})$ in relays results in a larger electrical force, larger acceleration, and thus a shorter mechanical delay. In addition, the mass (m) and spring constant (k) of the device affect how rapidly the electrostatic force displaces it. As developed in [10] and [11], under typical operating conditions, these effects on the mechanical turn-*on* delay $(t_{\rm mech})$ of the relay can be mathematically modeled by

$$t_{\rm mech} \propto \sqrt{\frac{m}{k}} \cdot \left(\frac{V_{\rm pi}}{|V_{\rm gb}|} \right). \tag{4}$$

If the gate-to-body voltage is held constant as the gate is pulled in, then the amount of force applied to the gate will be higher than it was in the *off*-state since the effective gap is smaller. This creates hysteretic behavior in the MEM relay: once the device is pulled in, the absolute value of the gate-to-body voltage must be lowered below the pull-in voltage in order to cause the device to release or pull out. This hysteresis effect is also increased by surface forces that attract the channel to the source/drain electrodes in the contact regions. In spite of the hysteresis, the mechanical turn-*off* delay is typically much faster than the turn-*on* delay because electrical contact is broken as soon as the channel dimple moves more than $\sim$1 nm away from the surface when the device is released. In contrast, turning the device *on* requires that the gate must travel nearly the entire gap between the dimple and the source/drain electrode (which in our devices is $\sim$90 nm) before making electrical contact.

Having described the mechanical model of the relay, it is interesting to note that concerns related to environmental vibrations causing spurious actuation of the relay are largely misplaced. The low mass of the structure implies that it has low inertia and ensures that the device can only be actuated by external forces in the presence of very large accelerations. For instance, the devices fabricated in this work would require an external acceleration of 20 000 g to balance against the spring force (kx) such that contact between the channel and the source/drain electrodes is made. If all dimensions are scaled equally, the device tends to become even more immune to vibration as it scales because the mass shrinks in a cubic fashion while the opposing spring force scales only linearly.

Similarly, thermal energy has little impact on device operation at the scales considered in this study. Since the structure is effectively confined to one (vertical) degree of freedom, it is subject to $k_B T/2$ joules of thermal energy. Comparing this to the $kx^2/2$ joules of stored spring energy, it can be shown that the σ of displacement due to thermal energy is only $\sim$8 pm for the fabricated relays. This is 0.005% of the size of the current actuation gap and causes negligible variation in the pull-in voltage. Though scaling the devices will reduce the spring constant and the gap size, leading to a relative increase in the effect of the thermal energy displacement, the σ of the variation in pull-in voltage remains less than 0.5% at the scaled dimensions considered in Sections IV and V.

C. Electrical Modeling of the MEM Relay

Although mechanical motion tends to dominate the switching delay of a single relay, the overall circuit switching delay is affected by the electrical delay as well. Predicting the amount of time required for a channel in the *on*-state to discharge a load capacitance requires accurate modeling of both the *on*-state resistance and the capacitances of the device.

The *on*-state resistance is comprised of the resistance of the tungsten wires leading to and from the device $(R_{\rm trace})$, the resistance of the channel $(R_{\rm ch})$, the resistance of the contacts between the channel and the source/drain $(R_{\rm con})$, as well as the resistance of the passivating oxide used to improve the endurance of the device $(R_{\rm pox})$. Of these components, the last two are by far the most significant contributors to the total *on*-state resistance. As summarized in [12], the resistance of the contact at each side of the channel depends on the conditions under which the contact is made and the properties of the material

$$R_{\rm con} = \frac{4\rho\lambda}{3A_r} \tag{5}$$

where ρ is the resistivity of the contacting material, λ is the mean free path of electrons in the contact material, and A_r is the effective contact area as given by

$$A_r \approx \frac{F_{\rm elec}(g_d)}{\xi H}. \tag{6}$$

In (6), H is the hardness of the material and ξ is the deformation coefficient. In the case of our contacts, the material is tungsten and the contact is elastic. The parameter values of $\rho = 55$ nΩ-m, $\lambda = 33$ nm, $\xi = 0.3$, and $H = 1.1$ GPa model this

TABLE I
SCALED AND CURRENT MEM RELAY DEVICE MODEL PARAMETERS. MEASURED DATA IS INDICATED WITH AN ASTERISK

Parameter	Current Devices	Scaled Model
A_{ov} [μm^2]	450	0.77
g_0 [nm]	180	10
g_d [nm]	90	5
R_{con} [Ω/contact]	~0.1	40-400
R_{pox} [Ω/contact]	500	500
C_{gc} [fF]	128.3	0.9
C_{gb}(x=0) [fF]	32.3	1.46
C_{gd}(x=0) and C_{gs}(x=0) [fF]	6.6	0.6
k [N/m]	62.5	0.07
m [fg]	3700	0.86
b [μN/m/s]	50.785	0.0078
t_{mech} [μs]	34.0*	0.02-0.08
$t_{elec} = R_{on}(C_{gc}+ C_{gb}+2C_{gd/s})$ [ps]	304.4	2.5-3.5
V_{pi} [V]	~8-10*	0.04

situation and using these values in (6) leads to the range of resistances given in Table I.

The overall resistance is also modified by the sub-1 nm thick TiO_2 electrode coating that limits current flow and mitigates native oxide formation to improve device reliability [8]. Devices with this coating have been operated for over 60 billion cycles and show no surface wear after this cycling [5], [8]. The resistance of this oxide coating is included in the model by comparing measured device resistance to the theoretical contact resistance model introduced above and the results are noted in Table I.

Finally, the electrical delay is also determined by the load capacitance seen by the device. In our intended VLSI applications the load capacitance is dominated by wire parasitics and the load presented by other relay devices. The load presented by the devices consists of many parasitic capacitors, but the largest of these are the capacitors formed by the air gap between the gate and the body across which the device is actuated $(C_{\rm gb})$, and the parasitic capacitance between the gate and the channel of the device $(C_{\rm gc})$. The channel terminal of the gate-to-channel capacitance is floating when the relay is in the *off*-state, so it only contributes to the total capacitance when the relay is in the *on*-state. Both of these are well-modeled as standard parallel plate capacitors

$$C_{\rm gb}(x) = \frac{\varepsilon_o A_{\rm ov}}{g_0 - x} \text{ and } C_{\rm gc} = \frac{\kappa_{\rm gox}\varepsilon_o A_{\rm ch}}{t_{\rm gox}} \tag{7}$$

where $\kappa_{\rm gox}$ is the relative permittivity of the gate oxide, $A_{\rm ch}$ is the area of the gate to channel overlap, and $t_{\rm gox}$ is the thickness of the gate oxide. Even though $A_{\rm ch}$ is relatively small compared to $A_{\rm ov}$, the gate-to-channel capacitance can be a significant contributor to the overall load capacitance because of the high relative permittivity of the gate oxide—Al_2O_3 in our devices. The channel-to-body capacitance, $C_{\rm cb}$, is not significant because it is formed over the same air gap as the gate to body capacitance but has the smaller area of overlap, $A_{\rm ch}$.[1]

All of the other capacitors in the device model arise from the overlap between the gate and the other electrodes. The overlap with the drain and source electrodes introduces gate-drain $(C_{\rm gd})$ and gate-source $(C_{\rm gs})$ parasitic capacitors. These are also modeled as parallel-plate capacitors, with gap size and permittivity that are the same as for the gate-to-body capacitance. Consequently the ratio of $C_{\rm gb}$ to the $C_{\rm gd}$ and $C_{\rm gs}$ capacitances is set by the ratio of gate-to-body overlap area to gate-to-drain/source overlap area and, by design, the gate-to-drain and gate-to-source overlaps are smaller than the gate-to-body overlap. As shown in Table I, the total capacitance from the gate to the source or to the drain in our measured devices $(C_{\rm gs} + C_{\rm gd})$ is roughly 40% of the main actuation capacitance $(C_{\rm gb})$, with this ratio reducing to $\sim$27% in the scaled device with a somewhat improved layout. Since they are not negligible, the drain and source parasitic capacitors must be included in both the delay and energy analyses.

By combining the electrical and mechanical models, performance estimates of a general MEM-relay device have been obtained and verified by experiment [6]. These estimates can in turn be used to predict the performance of devices with different dimensions. To expedite circuit design with the relays these models have also been implemented in Verilog-A. The performance of the analytical and computer models is well correlated with measured data for a similar device [8]. Notably, this model accurately captures several important phenomena of device operation: the switching voltages, the mechanical delay, and the electrical delay.

III. MEM-RELAY CIRCUITS

Although the previously described models are valuable for predicting the behavior of MEM-relay circuits, experimental verification of their functionality is clearly needed. Previous work has characterized a complementary MEM-relay inverter [5], but in order to demonstrate the feasibility of implementing more highly integrated MEM-relay circuits, and to study their properties, a range of example circuits have been fabricated on

[1]In a properly designed device, the gate-to-body and gate-channel capacitors will dominate the total capacitance of the device. Our initial devices utilized a simplified process that required large anchors to ensure that the anchors wouldn't be released, and thus the parallel plate capacitance from the gate anchors to the substrate was significant. Fortunately however, the substrate plays essentially no role in the operation of the device, and thus the thickness of the substrate oxide could in principle be increased. Furthermore, in a scaled process the anchor dimensions can be directly set by lithography and hence the anchor capacitance of a scaled device would be negligible.

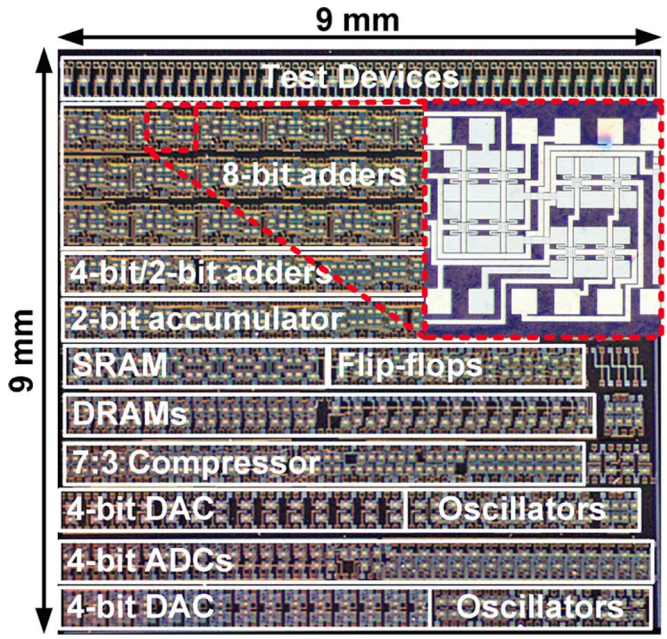

Fig. 3. Die photo with relay-based, propagate-generate-kill circuit shown in inset.

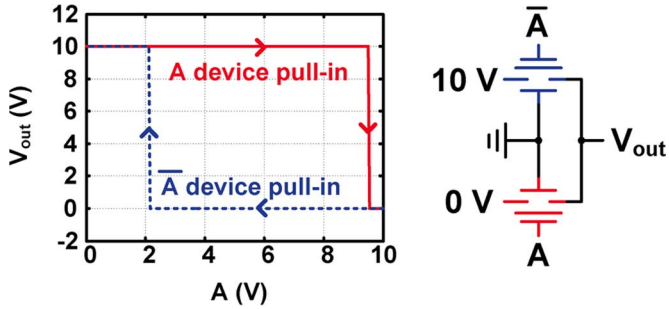

Fig. 4. MEM relay based inverter and measured VTC illustrating full-rail swing at the output and digital gain.

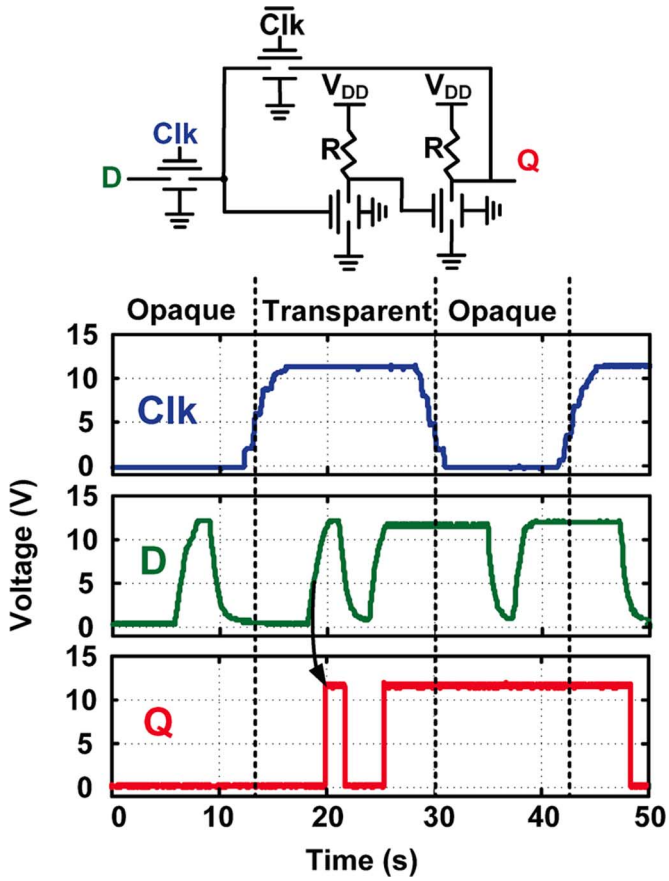

Fig. 5. Latch composed of MEM relay devices and waveforms showing its operation. Functionality of the latch illustrates that MEM relay logic stages are composable.

a demonstration chip (Fig. 3) using a 1 μm lithographic process [7]. In this section we show measurements of circuits composed only of relays with grounded body terminals, which limited the number of parasitic effects that could affect device operation.

The static switching characteristics of a MEM relay-based logic gate are demonstrated in Fig. 4, which shows the schematic and measured voltage transfer characteristic (VTC) of a pass-gate style MEM inverter/XOR. The VTC of the XOR with one input held static at 10 V highlights the expected hysteresis of the relay-based XOR; this hysteresis window determines the minimum voltage swing required to switch both devices *on* and *off*. The switching voltages of the inverter depend on the pull-in voltage of the devices and there is some device-to-device pull-in voltage variability on the test chip; this variability was caused by an uneven film deposition during the manufacturing process. This is evidenced by the difference in the magnitude of the switching voltage in the A pull-in and $\bar{A}$ (10 V minus the A voltage) pull-in paths. Unlike CMOS, the conduction of the MEM relay is ideally independent of the drain and source voltages, which enables the "NMOS"-style pass-gate to swing full-rail at the output. This VTC shows that the gate is capable of driving the necessary output voltages to overcome its own hysteresis window and switch another gate. Like in CMOS, this is a critical requirement to enable the composition of multiple digital logic gates.

The composability of MEM-relay circuits is verified in Fig. 5, where the measured MEM-relay latch demonstrates both transparent and opaque states. Like in CMOS, a cascade of two latches could be used to create a flip-flop. The latch's circuit topology is directly ported from an equivalent pseudo-NMOS structure. Although the circuit is functional, it suffers several mechanical delays during operation—i.e., the performance of this implementation is much worse than it would be if it were re-optimized for the characteristics of MEM relays.

The most significant relay characteristic, as predicted by the model as well as earlier works, is a large discrepancy between the intrinsic mechanical and electrical delays of a MEM relay [6]. These predictions are verified by the measurements in Fig. 6, which show the schematic and two waveforms for a single relay pseudo-NMOS style oscillator that were measured at the same time but at different voltage resolutions.

In Fig. 6, the rising edge is set by the RC time constant of the 74 kΩ external load resistor and the capacitance due to the test infrastructure (probecard and oscilloscope), which is estimated to be 55 pF. The mechanical delay of the device can be measured from the time the rising edge reaches the previously characterized $V_{\rm pi}$ to the time when the relay actuates—i.e., when

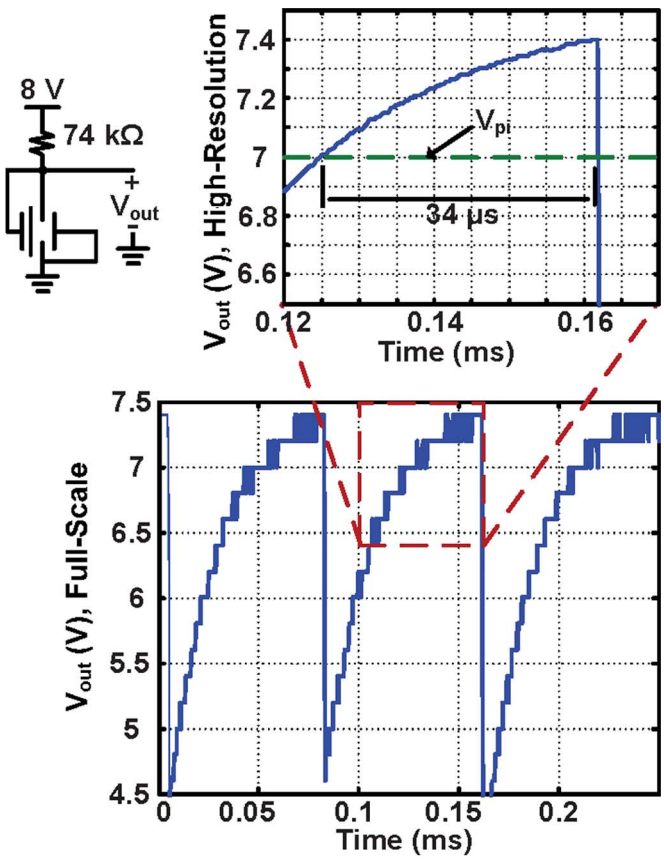

Fig. 6. Two waveforms from a single-stage, MEM relay-based oscillator that were collected at the same time but at two different voltage resolutions, one for full-scale waveforms and one for a precise timing measurement. These waveforms can be used to estimate load capacitance (55 pF), device *on*-resistance ($\sim$1 kΩ), and device mechanical delay (34 μs).

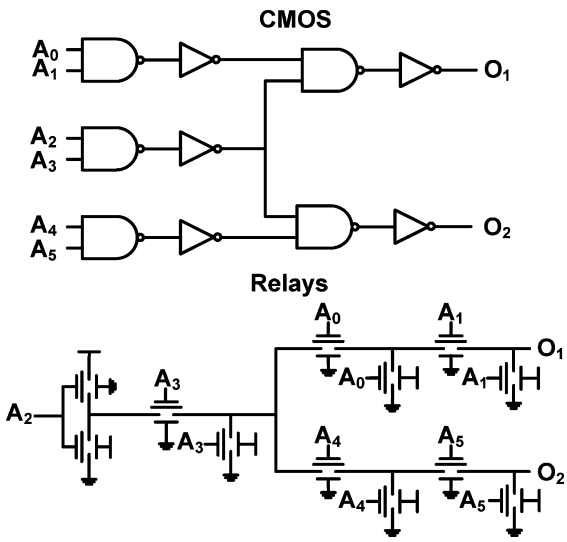

Fig. 7. CMOS to MEM relay logic mapping.

the output voltage begins to drop. Here, die-to-die variability resulted in a lower pull-in voltage than the previously measured inverter. Despite the lower pull-in voltage, 8 V of gate-to-body potential for the oscillator device is a relatively low "gate overdrive" (i.e., $V_{\rm dd}/V_{\rm pi}$). When the gate overdrive is low the mechanical delay is especially sensitive to changes in the overdrive, and thus the measured mechanical delay varies between $\sim$25–35 μs on different cycles.

Once the relay actuates, the output is discharged based on the electrical delay of the MEM relay. The oscillator's falling edge sees the same load capacitance as the rising edge, but its delay is set by the *on*-resistance of the relay, which is estimated as 1 kΩ based on the sub-300 ns electrical delay. The electrical delay measurement in Fig. 6 does not reflect the intrinsic delay of the device because of the parasitic capacitances due to the testing infrastructure. If the actual load capacitance was only the gate of another relay, which is estimated as $\sim$300 fF, the resulting electrical time constant would be on the order of $\sim$0.3 ns.

Even with the high capacitance of the probe card and oscilloscope, these measurements indicate that there is at least a difference of several orders of magnitude between the device's mechanical and electrical delays, which has implications for both the device fabrication and circuit design. Despite the many sources of contact resistance discussed in Section II, the relatively small electrical delay means that there is significant margin for contact resistance variation without dramatically impacting the dynamic device performance. This flexibility may even be exploited to enable material tradeoffs that sacrifice the electrical delay by introducing a slightly higher contact resistance for improved reliability, such as the previously described oxide coating.

The imbalance between electrical and mechanical delays suggests that optimized relay-based circuit designs should minimize the impact of the mechanical delay by arranging all mechanical movement to happen simultaneously within each circuit. The resulting design paradigm, as shown in Fig. 7, is that MEM relay logic should be designed as single complex gates wherever possible [6]. Only once the complexity of the logic gate has grown to the extent that the electrical delay due to device stacking becomes larger than the mechanical delay should the gate be partitioned into multiple stages. Therefore, regardless of the final output load, the worst case electrical delay for a logic stage should be roughly balanced with the mechanical delay. For reference, this balance is achieved with a stack of 600 relays when driving a single device (fanout of one) using the currently fabricated devices. A stack of 290 relays balances the delays when using the scaled devices described later in this work.

This design paradigm is demonstrated in the carry-generation circuit shown in Fig. 8. This circuit is a key component of the Manchester carry chain adder, which is particularly favorable to MEM relays since the carry signal does not incur additional mechanical delays as it propagates through each adder stage. In principle, the XOR function can be implemented efficiently using a single MEM relay with the input signals tied to gate and body terminals [6], but because only devices with the body terminal tied to ground were measured on this test chip, the propagate XOR is implemented with four devices [7]. The waveforms in Fig. 8 also illustrate the operation of this circuit in propagate, generate, and carry modes, showing that MEM relays are amenable to logic implemented using complex gates.

Given its immeasurably low leakage current [4] (and hence the potential for very large retention times), the MEM relay is

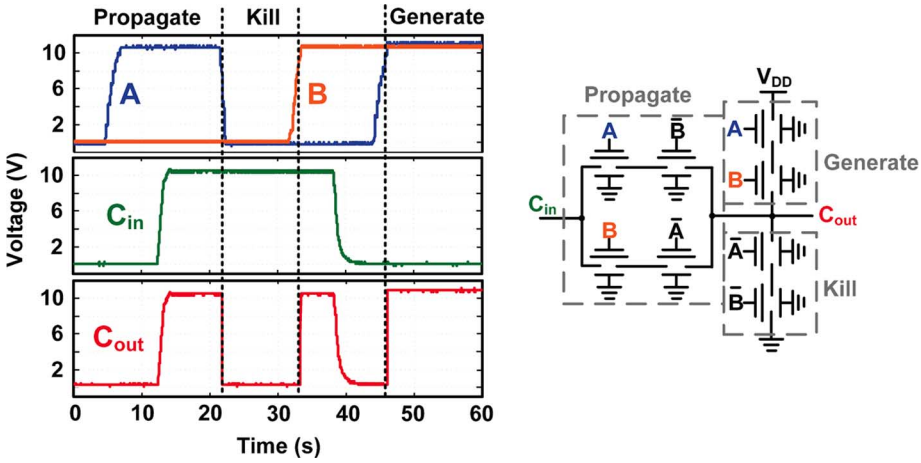

Fig. 8. MEM relay based carry generation circuit and measured waveform demonstrating operation in propagate, generate and kill modes.

also very well suited for implementing DRAMs. The structure of this circuit is illustrated in Fig. 9, which shows a 10-bit DRAM column composed of MEM relays constructed in a NAND configuration. Each DRAM cell consists of three devices as seen in Fig. 9: a storage device which can either short or open the read bit-line ($\mathrm{BL_{RD}}$), an access device that gates the signal between the write bit-line ($\mathrm{BL_{WR}}$) and the gate of the storage device, and a bypass device that is in parallel with the storage device and also attached to the read bit-line.

During a read operation, the read bit-line output is pre-discharged to a low voltage and the read word-line of the cell being interrogated ($\mathrm{WL_{RD}}[n]$) is lowered from its normally high voltage. Because the read word-line is normally high, the bypass devices in every cell except for the cell being read are turned *on*, shunting/bypassing the corresponding storage devices. Thus, if the gate capacitance of the storage device in the interrogated cell contains enough charge so that device will be turned *on*, a conducting path from the supply to the bottom of the read bit-line stack will be formed and $\mathrm{BL_{RD}}$ will be pulled high. Otherwise, $\mathrm{BL_{RD}}$ will remain low. During a write operation $\mathrm{WL_{WR}}[n]$ is raised high to allow $\mathrm{BL_{WR}}$ to charge or discharge the gate of the storage device.

This configuration allows the memory to perform a read operation in a single mechanical turn-*on* delay (for decoding the address) plus a mechanical turn-*off* delay (to turn *off* the bypass device of the cell being read). As previously mentioned, the turn-*off* delay is much smaller than turn-*on* delay, and hence the read latency of this DRAM design is substantially lower than designs similar to CMOS implementations that would require 2 or more turn-*on* delays.

An experiment illustrating the operation of the DRAM is shown in Fig. 9, where the waveforms show a simultaneous memory read and write. The 10-bit memory cell was fully integrated except for (due to lack of vias) the wires between the drain of the access device and the gate of the storage device.[2] The pre-discharge of the read bit line was accomplished using

[2]This external connection meant that the storage node of the device was exposed to the capacitance and leakage of the probe card and external wires. Accurate measurements of the retention time of the DRAM could therefore not be taken since the measured retention time would be completely set by external leakage current discharging the external capacitance.

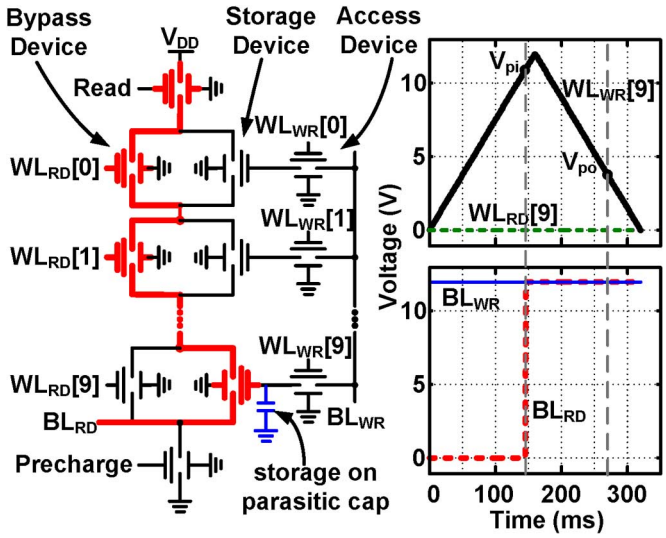

Fig. 9. Read and write of a single bit in a MEM-relay-based NAND-style DRAM column. The $\mathrm{WL_{RD}}[0]$ to $\mathrm{WL_{RD}}[8]$ signals are high so that the highlighted devices attached to them are turned *on*, creating a conductive path from $\mathrm{V_{DD}}$ to the tenth DRAM cell.

a 100 kΩ pull-down resistor that bypassed the pull-down relay in order to reduce the number of control signals needed to test the circuit as well as enable the simultaneous read and write operation.

Finally, the fact that each of the circuits described were able to drive the probecard and the oscilloscope suggests that the devices are capable of driving large capacitive loads. The ability to drive these loads presents the MEM relay as a candidate for an I/O device. Fig. 10 demonstrates the MEM relay in that role by showing the operation of a 2-bit thermometer-coded DAC. The DAC is implemented using three MEM relay-based buffers which create a programmable resistor divider between the I/O voltage rail and ground. Note that because relay actuation ideally depends only on the gate-to-body voltage and not on the drain or source voltages, it is relatively straightforward to incorporate a level-shift into the output stage. This is demonstrated in Fig. 10, where the output full-scale voltage is 3 V while the relay actuation voltage is 10 V.

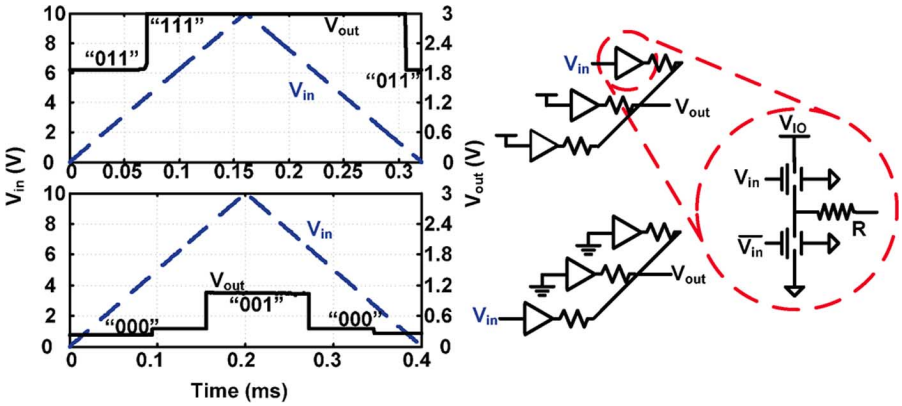

Fig. 10. A MEM relay based 2-bit thermometer coded DAC. The outputs for "11X" and "00X" inputs are shown, where X= V_{in}. A level shift is built into the DAC because the V_{IO} voltage is different than the actuation voltage.

IV. Effects of Scaling MEM Relays

Even though the previously described circuits demonstrate the functionality of MEM relay-based circuits, the size and high switching voltages of the current devices do not immediately demonstrate the benefits of these MEM circuits over current CMOS implementations. However, like CMOS transistors, MEM relays are expected to achieve substantially lower switching energy and lower delay as their physical dimensions are scaled.

The modifications to the device necessary to achieve these benefits can be analyzed by looking at the sources of energy consumption in a MEM relay. Energy is primarily spent driving the parasitic capacitances, and thus the energy/per operation is improved by decreasing the load capacitance—achieved by reducing the actuation area—and lowering the relay operating voltage. Both of these effects reduce the electrical force available to actuate the device, and hence a corresponding decrease in the spring constant—achieved by thinning the flexures—and reduction of the gap thickness is necessary to reduce the pull-in voltage. Even though the pull-in voltage and applied force are reduced, the mass of the device scales more quickly than the spring constant, which results in faster actuation at smaller dimensions.

Of course, this improved performance and particularly the reduction in energy cannot be extended forever. Ultimately, the spring must be able to restore the relay to its un-actuated position after the electrical force is removed. The surface forces holding the relay in place—primarily hydrogen bonds, capillary forces, and Van der Waals force [13]—therefore provide a lower bound on the strength of the spring and set the minimum amount of energy for switching the device *on* and *off*. For large contact dimple areas, these forces are proportional to the area of the contact dimple. Thus, even though the spring force will become weaker as the device is scaled, scaling the contact dimensions will allow the spring force to dominate surface forces even as the device becomes smaller.

However, eventually both the surface forces and contact resistance will be determined by a handful of bonds between the channel and the drain/source electrode. At this point, further scaling of the contact dimple area leaves the surface forces largely unchanged [13]. This implies that the minimum stored spring energy must be large enough to overcome the energy of a small number of bonds; given five bonds, each of which have an energy of ~0.2 aJ [14], the minimum switching energy would be 4 aJ per device switching cycle [8]. This is roughly a factor of ten lower than the minimum energy per switching cycle of a single, minimum-sized 65 nm CMOS transistor [2], [8]. It is important to re-iterate however that since relay-based circuits will be constructed in a significantly different manner than their CMOS counter-parts, the true energy benefit of relays must be evaluated at the circuit level.

V. Comparison of Scaled MEM-Relay and CMOS Logic

To realistically project the benefits of a MEM relay it is necessary to compare modern CMOS circuits against their counterparts built from scaled versions of fabricated MEM relays. However, in measuring the original test chip, several parasitic effects that affected the operation of the devices were discovered. Specifically, the relatively large gate-to-drain and gate-to-source capacitances lead to parasitic actuation forces between the source or drain electrodes and the gate. Through this mechanism the voltage on the source and drain (relative to the gate) could change the pull-in voltage, as seen in Fig. 11(a).

In addition to the parasitic effects caused by the drain and source-to-gate overlaps, the overlap between the channel and the body could result in an electrostatic attractive force between the channel and the body. If this force is sufficiently large, it may keep the device in the *on* state after the gate-to-body voltage is removed. The combination of these effects made the use of the body electrode for device control unpredictable. These undesirable effects must be eliminated in a practical scaled device, and as a first step in that process, devices with an improved layout were fabricated and tested. Fig. 12 shows an image of the devices [15] and of the test chip utilizing these devices. In these devices the drain and source areas were reduced, the channel size was minimized, and the channel-to-body overlap was shrunk. These changes resulted in a drastic reduction in parasitic capacitances and forces as evidenced by Fig. 11(b).

Utilizing the improved device design, a one-bit adder was implemented on the second test chip in a style that took advantage of the ability to use both gate and body as logic inputs. As shown in Fig. 13, since the relay is actuated when the absolute value of gate potential is greater than the pull-in voltage, it is possible to

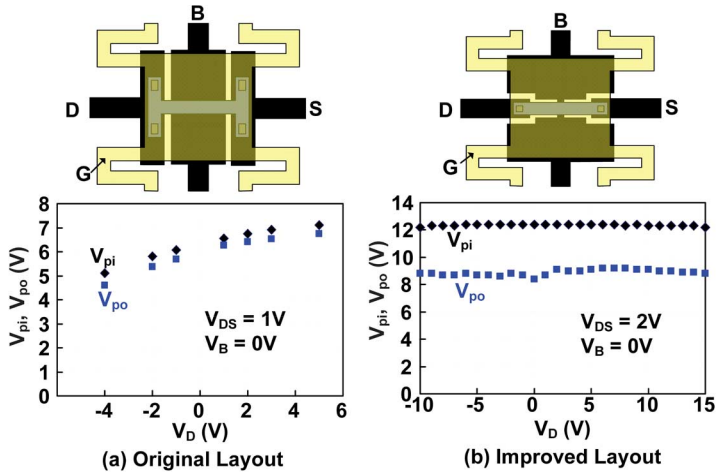

Fig. 11. Dependence of pull-in and release voltages (which are measured between gate and body) on the drain to body voltage of (a) the original device design, and (b) the revised design with reduced source/drain to gate overlap.

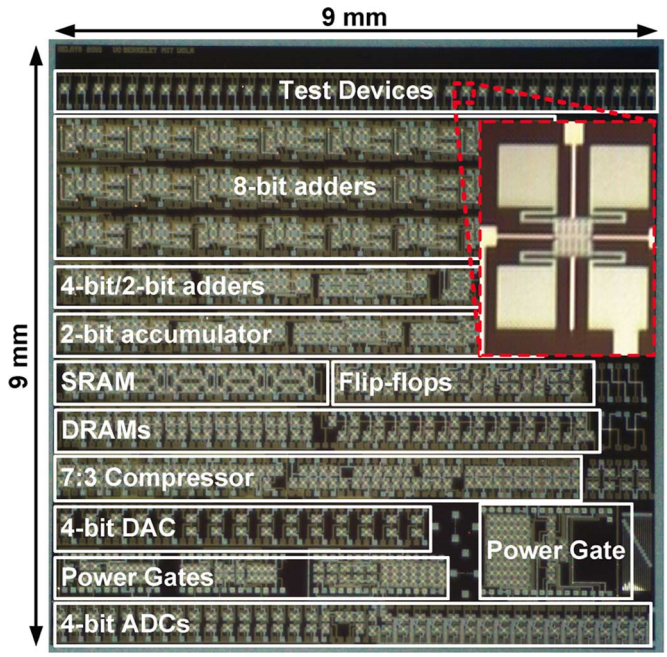

Fig. 12. Die photo of the test chip using the revised device layout and a zoomed-in image of a single device.

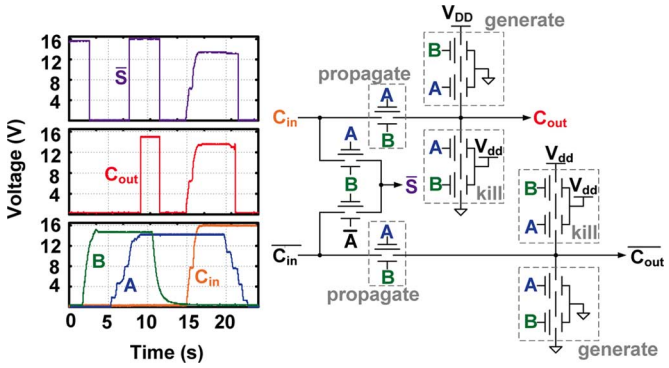

Fig. 13. Schematic and measured waveforms of a one-bit adder implemented using the revised layout for the MEM relays.[3]

use the back-gate as an input to implement the XOR of two signals (for both the *propagate* and *sum* calculations) rather than the four-device XOR used on the original test chip. Sum is implemented as a wired-XOR gate and both true and complement versions of the carry signal are computed to avoid the additional mechanical delay that might be required to invert the incoming carry signal.

Fig. 13 also shows the measured results from this full-adder circuit. The operating voltage used in these results was higher than on the previous demonstration chip because the as-fabricated gap thickness was larger than expected, raising the pull-in voltage of the device. We expect that this thickness would be better controlled in an industrial fabrication facility.

In an effort to reduce the footprint of the device while preserving the original device topology, a number of evolutionary optimizations are incorporated into the device layout to maximize area efficiency and minimize parasitic capacitances. Specifically, the device anchors were moved in line with the flexures, the flexures were adjusted so that they contact the top/bottom of the actuation area rather than the sides, the actuation area was extended across the entire length of the device, and the length of the channel was minimized by extending the drain and source. The first three of these optimizations reduce the overall footprint by enabling the actuation area to be increased without stretching The improved device layout from this second test chip was used as the starting point for the design of a scaled device used to analyze MEM relay logic performance. Fig. 14 presents a device layout optimized for a 90 nm equivalent MEM relay technology; in this device most dimensions are directly scaled by a factor of 50 compared to the measured devices. the device footprint vertically, while the fourth minimized the load capacitance by limiting the channel overlap area.

Though better performance, energy, and/or area could possibly be achieved by more radical changes in the device

[3]The sum and carry-out signals don't precisely track the carry-in signal in the measured waveforms because of a source resistance that was attached to the supply generating C_{in} in order to limit the current to the devices. This source resistance formed a resistive divider with the resistance of the oscilloscope probe. C_{in} was measured on the source side of the source resistance and thus did not see the same voltage division as S and C_{out}.

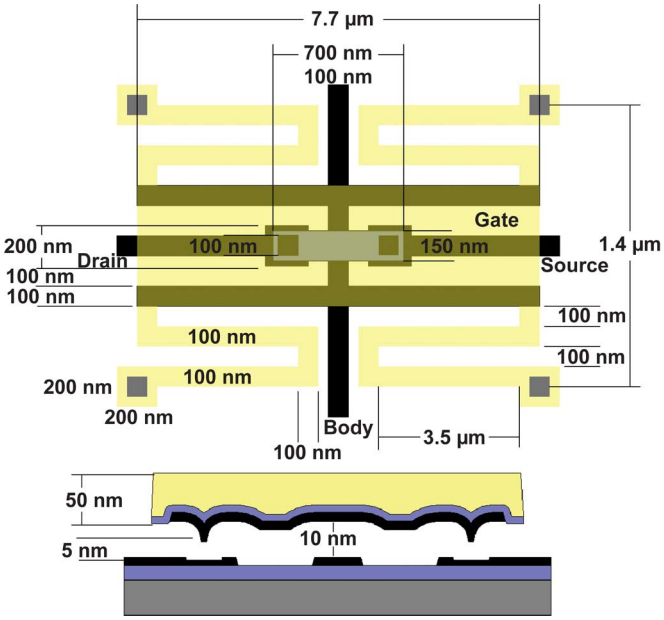

Fig. 14. Possible layout for a 90 nm technology MEM relay that incorporates optimizations to improve area efficiency and reduce parasitic capacitance.

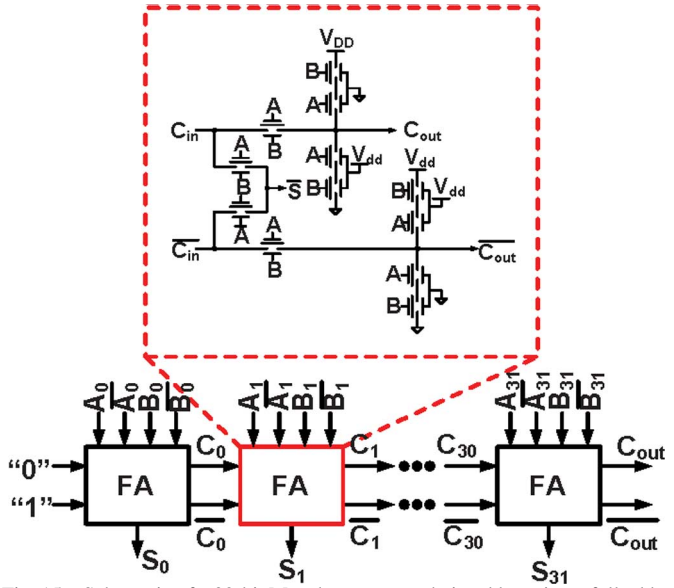

Fig. 15. Schematic of a 32-bit Manchester carry chain adder using a full adder cell implemented as a single compound gate built from MEM relays.

structure (e.g., by switching to the simple cantilevers analyzed in [6]), this scaled version of the fabricated devices is analyzed because of the similarities with the current devices that have already achieved good yield and reliability and are thus closer to large-scale realization. To be more specific, although the number of devices fabricated was insufficient for a thorough statistical study, all of the devices that were tested on the fabricated chips were functional, and as mentioned previously, the devices remained functional after more than 60 billion cycles.

Table I shows the consequences of this scaling on the device parameters that directly impact relay performance. The values in the table are calculated by using the previously established electrical and mechanical models along with the new device dimensions. It is important to notice that the switching delay of a MEM relay is dominated by the mechanical delay of the structure even at these scaled dimensions. Specifically, as seen in Table I, the mechanical delay is on the order of tens of nanoseconds, while the intrinsic electrical delay of the device is less than five picoseconds. This remains true even for stacks of 32 series relays as long as the contact resistance is less than 15 kΩ, and hence the precise scaling behavior of the contact resistance is unlikely to affect the design style of the circuit.

In order to illustrate the potential benefits of scaled MEM relays while keeping in mind the differences in circuit design paradigms, we next compare the energy-performance characteristics of an optimized 32-bit MEM relay-based adder against those of an optimized CMOS implementation [16]. The implementation of the 32-bit relay adder (which utilizes the full-adder cells measured in Fig. 13) is shown in Fig. 15. Note that in total, only 12 relays (as compared to 24 transistors in CMOS) are used to implement a full-adder cell, which helps to amortize the area penalty stemming from the fact that each individual MEM relay is significantly larger than a minimum-sized CMOS transistor.

In order to implement the complete multi-bit adder, the MEM relay-based full-adder cell is used in a ripple-carry configuration. Overall, the structure implements a complete 32-bit add function in one compound gate and requires 384 relays, each of which occupies slightly less than 12 μm^2. Assuming a wiring overhead of $\sim$30%, this results in a total area of $\sim$ 7000 μm^2 for the MEM-relay adder.

In order to benchmark the NEM relay-based adder, it was compared to a 32-bit CMOS Sklansky adder [17]. Sklansky adders were identified as the minimum energy topology across a wide range of delays in [16]. At the minimum energy/maximum delay, the adder uses 836 gates and occupies an area of $\sim$2000 μm^2 when placed and routed using standard cells with no delay constraint. The energy per operation of the CMOS adder is based upon the results from [16], and Verilog-A simulations of the whole relay adder were used to estimate its energy.

Fig. 16 shows the energy-delay tradeoffs in CMOS and MEM-relay adders where the adders have been designed to drive a load capacitance of either 25 fF or 100 fF. In Fig. 16(a) the energy spent driving the load is omitted in order to facilitate comparisons between only the adders, while Fig. 16(b) includes the load energy to compare the overall performance. The delay of all of the adders includes the effects of driving the load capacitance at each of their outputs.

The delay from driving these load capacitances increases with the size of the load. When the electrical delay becomes comparable to the mechanical delay, it is beneficial to add a buffering stage to drive the load and reduce the electrical delay of the circuit. Consequently, the results for relay adders with 100 fF load in Fig. 16 assume a single buffer is added at the output to drive the loads. This modification roughly doubles the delay of the unloaded circuit because two mechanical delays are incurred, but essentially eliminates the electrical delay and thus improves performance in the most energy-efficient way.

The energy of the CMOS adder reaches its minimum energy point [1], [16] for delays above $\sim$ 1 ns. Thus, at delays of $\sim$ 50 ns, a single MEM-relay adder would offer an improvement of $\sim$ 10x or more in energy efficiency with an additional area overhead of $\sim$3.5x compared to the CMOS adder. Applications requiring throughputs of $\sim$20 MOPS or less would immediately benefit from deployment of such a technology.

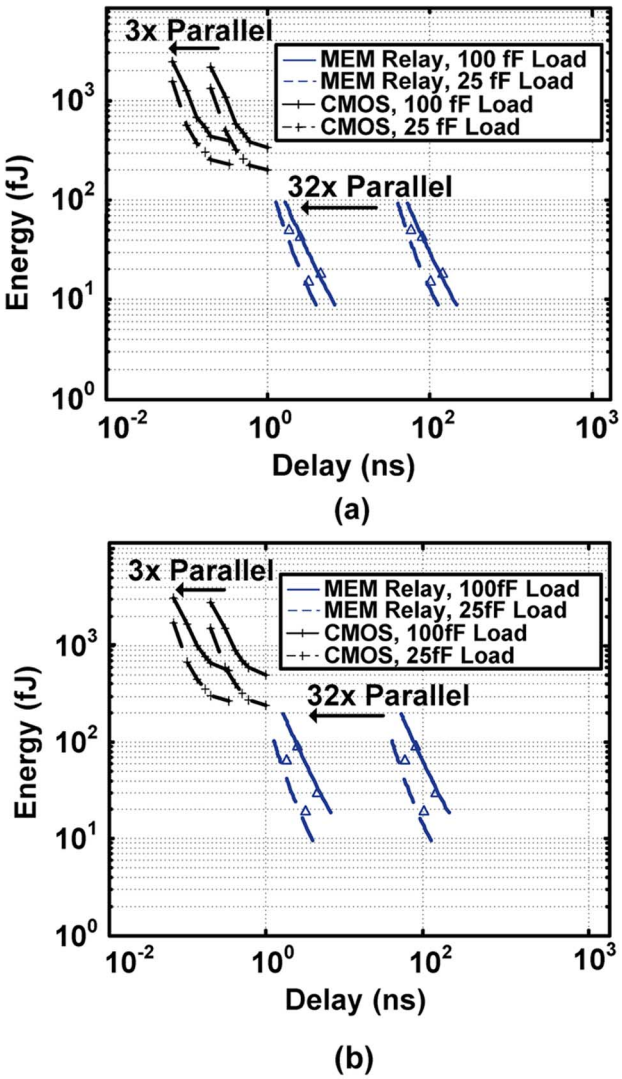

Fig. 16. Energy-throughput comparison of 32-bit adders after sizing and $V_{\rm dd}$ scaling of static CMOS Sklansky design [17] versus MEM relay adder. Comparisons excluding the load energy (a) and including the load energy (b) are both shown. The discrete points marked on the relay curves are extracted from Verilog-A using the relay model highlighted in Fig. 2.

The improved energy-efficiency offered by relays can also be extended to higher throughputs by trading-off increased area to make use of parallelism (as also indicated in Fig. 16). For example, a 32-parallel MEM relay adder implementation operating at 0.5 GOPS, which has 10x lower energy per operation than CMOS, would require 100x the area. As previously mentioned, this area penalty can be reduced by more radical optimization of the process and device layout—e.g., cantilever-based relays like those analyzed in [6].

VI. Conclusion

This work has demonstrated the functionality of several MEM-relay circuits that are key VLSI system building blocks. The presented analyses suggest that scaled versions of MEM relays have the potential to reduce the energy consumption of adders, and possibly many other VLSI blocks, over a wide range of performance points.

In addition to the demonstration of relay-based circuits, this work contains an analysis showing that relay-based circuits generally achieve minimum delay by assembling large, complex gates where all of the devices are actuated at the same time. Using this design style, simulations predict that relay-based adders in a 90 nm technology node can achieve energy savings of $\sim$10x over a minimum-energy-point CMOS design for up to 20 MOPS. Though these energy savings have not been experimentally verified, the building blocks necessary to construct an adder and other VLSI blocks with this style have been demonstrated.

These results, taken together, point to MEM relays as interesting candidates for the implementation of VLSI circuits as CMOS designs become increasingly power limited. Though the functionality demonstrations here are far from illustrating the energy gains promised by comparative analysis of scaled relays vs. CMOS transistors, they confirm the feasibility of implementing logic using MEM relays and suggest that further scaling and integration of the relays could eventually lead to gains in energy efficiency over CMOS.

References

[1] B. H. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in subthreshold circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1778–1786, Sep. 2005.
[2] H. Kam, T. King-Liu, E. Alon, and M. Horowitz, "Circuit-level requirements for MOSFET-replacement devices," in *Proc. IEEE Int. Electron Device Meeting Tech. Dig.*, Dec. 2008, p. 473.
[3] K. Gopalakrishnan, P. B. Griffin, and J. D. Plummer, "I-MOS: A novel semiconductor device with a subthreshold slope lower than kT/q," in *Proc. IEEE Int. Electron Device Meeting Tech. Dig.*, Dec. 2002, pp. 289–292.
[4] W.-Y. Choi, B.-G. Park, J. D. Lee, and T. King-Liu, "Tunneling field-effect transistors (TFETs) with subthreshold swing (SS) less than 60 mV/dec," *IEEE Electron Device Lett.*, vol. 28, no. 8, pp. 743–745, Aug. 2007.
[5] R. Nathanael, V. Pott, H. Kam, J. Jeon, and T.-J. K. Liu, "4-terminal relay technology for complementary logic," in *Proc. IEEE Int. Electron Device Meeting Tech. Dig.*, Dec. 2009, pp. 223–226.
[6] F. Chen, H. Kam, D. Markovic, T.-J. K. Liu, V. Stojanovic, and E. Alon, "Integrated circuit design with NEM relays," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, Nov. 2008, pp. 750–757.
[7] F. Chen, M. Spencer, R. Nathanael, C. Wang, H. Fariborzi, A. Gupta, H. Kam, V. Pott, J. Jeon, T.-J. K. Liu, D. Markovic, V. Stojanovic, and E. Alon, "Demonstration of integrated micro-electro-mechanical switch circuits for VLSI applications," in *Proc. IEEE Int. Solid-State Circuits Conf. Tech. Dig.*, Feb. 2010, pp. 150–151.
[8] H. Kam, R. Nathanael, J. Jeon, E. Alon, and T.-J. K. Liu, "Design and reliability of a micro-relay technology for zero-standby-power digital logic applications," in *Proc. IEEE Int. Electron Device Meeting Tech. Dig.*, Dec. 2009, pp. 809–812.
[9] S. D. Senturia, *Microsystem Design*. Norwell, MA: Kluwer, 2001.
[10] J. B. Muldavin and G. M. Rebeiz, "Nonlinear electro-mechanical modeling of MEMS switches," in *Proc. Microw. Symp. Dig.*, 2001, vol. 3, pp. 2119–2122.
[11] G. M. Rebeiz, *RF MEMS: Theory, Design, and Technology* 2003.
[12] R. Holm, *Electric Contact*. New York: Springer Verlag, 1967.
[13] B. D. Jensen, K. Huang, L. L. W. Chow, and K. Kurabyashi, "Adhesion effects on contact opening dynamics in micromachined switches," *J. Appl. Phys.*, vol. 97, no. 10, pp. 103–535, May 2005.
[14] G. Rubi-Bollinger, S. R. bahn, N. Agrait, K. W. Jocobsen, and S. Vieira, "Mechanical properties and formation mechanics of a wire of single gold atoms," *Phys. Rev. Lett.*, vol. 87, no. 2, pp. 0261011–0261014, Jun. 2001.
[15] J. Jeon, R. Nathanael, V. Pott, and T.-J. K. Liu, "Four-terminal relay design for improved body effect," *Electron Device Lett.*, vol. 81, no. 5, pp. 515–517, May 2010.
[16] D. Patil, O. Azizi, M. Horowitz, R. Ho, and R. Ananthraman, "Robust energy-efficient adder topologies," in *Proc. 18th IEEE Symp. Computer Arithmetic (ARITH'07)*, Jun. 2007, pp. 16–28.
[17] J. Sklansky, "Conditional-sum addition logic," *IRE Trans. Electronic Comput.*, vol. EC-9, pp. 226–231, 1960.

Matthew Spencer received the B.S. and M.Eng. degrees from The Massachusetts Institute of Technology, Cambridge, in 2007 and 2008, respectively and is currently pursuing the Ph.D. degree at the University of California, Berkeley.

His research interests are energy-efficient circuits, emerging circuit technologies and, particularly, the integration of the two.

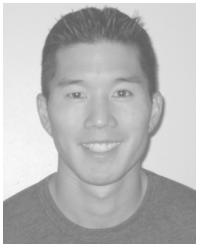

Fred Chen received the B.S. degree from the University of Illinois, Urbana-Champaign in 1997 and the M.S. degree from the University of California, Berkeley in 2000, both in electrical engineering. Since 2005 he has been at the Massachusetts Institute of Technology where is currently working toward the Ph.D. degree.

From 2000 to 2005 he joined Rambus Inc., Los Altos, CA where he worked on the design of high-speed I/O and equalization circuits. His current research interests include energy-efficient circuits and systems, and circuit design with novel devices.

Cheng C. Wang (S'05) received the B.S. degree in electrical engineering and computer sciences from the University of California, Berkeley, in 2005 and the M.S. degree in electrical engineering from the University of California, Los Angeles, in 2009, where he is currently working toward the Ph.D. degree.

Before joining UCLA, he worked as a VLSI design engineer at Zoran Corporation, Sunnyvale, CA. His current research interests include low-power design optimization and energy-efficient circuits and devices.

Rhesa Nathanael (S'08) received the B.S. (high hons) and M.S. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 2006 and 2008, respectively. He is currently working toward the Ph.D. degree also in the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley.

His research interests include advanced integrated-circuit device physics and technology, and the broad area of nanoelectronics. His current research focuses on Nano-Electro-Mechanical Systems (NEMS) devices and technology for ultra-low-power computation.

Hossein Fariborzi (S'08) received the B.S. degree from Sharif University of Technology, Iran, in 2006 and the M.S. degree from the University of Malaya, Malaysia, in 2008, both in electrical engineering.

He worked in Nokia, Iran, in 2006 as an RF Design Engineer. He is currently at the Massachusetts Institute of Technology, Cambridge, MA, working toward the Ph.D. degree. His research interests include digital/mixed-signal IC design and design and optimization of integrated systems with novel switching devices.

Abhinav Gupta received the B.S. degree in electrical engineering and computer sciences and the M.S. degree in electrical engineering from University of California, Berkeley in 2008 and 2009, respectively. His Master's research involved work on electro-mechanical relay devices for use in the design of digital logic blocks, mainly aimed at implementing memory structures using such devices.

Since July 2009, he has been working with the Windows Devices group at Microsoft Corporation, Redmond, Washington.

Hei Kam (M'05) received the B.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California at Berkeley (UCB) in 2004 and 2009, respectively.

From December 2009 to June 2010, he was a postdoctoral research fellow at UCB. He is the recipient of the Ford Motor Scholarship and the Graduate Assistance in Areas of National Need (GAANN) fellowship. His research interests are in novel logic devices with subthreshold swing steeper than 60 mV/dec, micro-electromechanical-systems (MEMS) and low temperature fabrication processes. He is currently with Intel Corporation.

Vincent Pott (M'06) received the M.S. degree in micro-engineering and micro-electronics and the Ph.D. degree in electrical engineering from the Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland, in 2002 and 2008, respectively.

From 2008 to 2010, he was with the University of California, Berkeley, as post-doctoral associate. He was the recipient of a Young Researcher Grant; awarded by the Swiss National Science Foundation. In June 2010, he joined the Institute of Microelectronics (IME), in Singapore. He currently holds a Senior Research Engineer position and works on various projects related to energy-efficient computing. His current research interests include micro and nano technology fabrication, MEMS, NEMS, nanowires and hybrid technologies. He has (co-)authored around 60 publications.

Dr. Pott's Ph.D. work was awarded the Dr. Dimitris N. Chorafas Award for the best EPFL thesis in computer sciences. He serves as a reviewer for various IEEE journals.

Jaeseok Jeon received the B.A.Sc. degree (first-class hons), in electronics engineering from Simon Fraser University, Burnaby, BC, Canada, in 2007. He is currently working toward the Ph.D. degree in electrical engineering at the University of California, Berkeley.

In 2006, he was an electronics designer intern at Kodak Graphic Communications Canada Company, where he involved in developing and testing precision electro-mechanical and optical devices for next-generation thermal heads. His research interests include nano/micro-electro-mechanical relays for ultra-low-power digital integrated circuit applications, surface enhanced Raman spectroscopy techniques, and micrometer-scale energy-scavenging devices.

Tsu-Jae King Liu (SM'00–F'07) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1984, 1986, and 1994, respectively.

In 1992 she joined the Xerox Palo Alto Research Center, Palo Alto, CA, as a Member of Research Staff, to research and develop polycrystalline-silicon thin-film transistor technologies for high-performance flat-panel displays. In August 1996 she joined the faculty of the University of California, Berkeley, where she is currently the Conexant Systems Distinguished Professor of Electrical Engineering and Computer Sciences and Associate Dean for Research in the College of Engineering.

Dr. Liu's awards include the DARPA Significant Technical Achievement Award (2000) for development of the FinFET, and the IEEE Kiyo Tomiyasu Award (2010) for contributions to nanoscale MOS transistors, memory devices, and MEMS devices. Her research activities are presently in nanometer-scale integrated-circuit devices and technology, as well as materials, processes, and devices for integrated microsystems. She has served on committees for many technical conferences, including the International Electron Devices Meeting and the Symposium on VLSI Technology. She was an Editor for the IEEE ELECTRON DEVICES LETTERS from 1999 to 2004.

Dejan Marković (S'96–M'06) received the Dipl.Ing. degree from the University of Belgrade, Serbia, in 1998 and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 2000 and 2006, respectively, all in electrical engineering.

In 2006, he joined the faculty of the Electrical Engineering Department at the University of California, Los Angeles as an Assistant Professor. His current research is focused on integrated circuits for emerging radio and healthcare systems, design with post-CMOS devices, optimization methods and CAD flows.

Dr. Marković was awarded the CalVIEW Fellow Award in 2001 and 2002 for excellence in teaching and mentoring of industry engineers through the UC Berkeley distance learning program. In 2004, he was a co-recipient of the Best Paper Award at the IEEE International Symposium on Quality Electronic Design. He received 2007 David J. Sakrison Memorial Prize from the Department of EECS, UC Berkeley.

Elad Alon (S'02–M'06) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2001, 2002, and 2006, respectively.

In 2007, he joined the University of California at Berkeley as an Assistant Professor of Electrical Engineering and Computer Sciences, where he is now a Co-Director of the Berkeley Wireless Research Center (BWRC). He has held positions at Sun Labs, Intel, AMD, Rambus, Hewlett Packard, and IBM Research, where he worked on digital, analog, and mixed-signal integrated circuits for computing, test and measurement, and high-speed communications. His research focuses on energy-efficient integrated systems, including the circuit, device, communications, and optimization techniques used to design them.

Dr. Alon received the IBM Faculty Award in 2008, the 2009 Hellman Family Faculty Fund Award, and the 2010 UC Berkeley Electrical Engineering Outstanding Teaching Award.

Vladimir Stojanovic (S'96–M'05) received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2005, respectively, and the Dipl. Ing. degree from the University of Belgrade, Belgrade, Serbia, in 1998.

He is currently the Emanuel E. Landsman Associate Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge. He was also with Rambus, Inc., Los Altos, CA, from 2001 to 2004. He was a Visiting Scholar with the Advanced Computer Systems Engineering Laboratory, Department of Electrical and Computer Engineering, University of California, Davis, during 1997–1998. His current research interests include design, modeling, and optimization of integrated systems, from novel switching and interconnect devices (such as NEM relays and silicon-photonics) to standard CMOS circuits.

Dr. Stojanovic was a recipient of the 2009 NSF CAREER Award.