Analysis and Demonstration of MEM-Relay Power Gating

Hossein Fariborzi[1], Matthew Spencer[2], Vaibhav Karkare[3], Jaeseok Jeon[2], Rhesa Nathanael[2], Chengcheng Wang[3], Fred Chen[1], Hei Kam[2], Vincent Pott[2], Tsu-Jae King Liu[2], Elad Alon[2], Vladimir Stojanović[1], Dejan Marković[3]

[1]MIT, Cambridge, MA; [2]UC Berkeley, Berkeley, CA; [3]UCLA, Los Angeles, CA

***Abstract* – This paper shows that due to their negligibly low leakage, in certain applications, chips utilizing power gates built even with today's relatively large, high-voltage micro-electro-mechanical (MEM) relays can achieve lower total energy than those built with CMOS transistors. A simple analysis provides design guidelines for *off*-time and savings estimates as a function of technology parameters, and quantifies the further benefits of scaled relay designs. Finally, we demonstrate a relay chip successfully power-gating a CMOS chip, and show a relay-based timer suitable for self-timed operation.**

I. Introduction

Power gating has become ubiquitous in ICs to reduce the power consumed by inactive CMOS logic circuits. However, the finite I_{on}/I_{off} ratio of MOSFET power gates limits their ability to reduce *off*-state leakage. In contrast, as described in [1], MEMS-based power gates that mechanically make or break electrical contact can completely eliminate *off*-state leakage. However, the leakage benefits of MEMS-based power gates may be outweighed by increased switching energy and voltage droop due to relatively large device dimensions and/or operating voltages and *on*-state resistance. In this paper we present a simple analysis that predicts the conditions under which electrostatically-actuated MEM relays can achieve energy savings over MOSFETs for power gates. Furthermore, we utilize recently developed MEM relays [2,3] to experimentally demonstrate that these switches can successfully power-gate a functional CMOS chip.

II. Relay Structure and Operation

Before comparing the energy-efficiency of MOSFET and relay-based power gates, it is instructive to examine the basic structure and operation of the MEM relay used in this study. As shown in Fig. 1 [3], the movable poly-SiGe gate is anchored to the substrate at four corners, and a tungsten channel electrode is attached to the gate via a dielectric layer (Al_2O_3).

The state of the switch is set by the gate-to-body voltage. When the gate-to-body voltage is below the release (pull-out) voltage (V_{PO}), the relay is in the *off*-state and an air gap separates the channel from the tungsten source/drain electrodes so that no current flows. When the gate-to-body voltage exceeds the pull-in voltage (V_{PI}), the channel contacts the source/drain electrodes and allows current to flow – i.e., the relay is in the *on*-state. As tungsten is a hard metal (which improves reliability [2]) that deforms minimally at the point of contact, the *on*-resistance of each relay is typically ~2 kΩ. A lower effective resistance can be attained by connecting multiple relays in parallel.

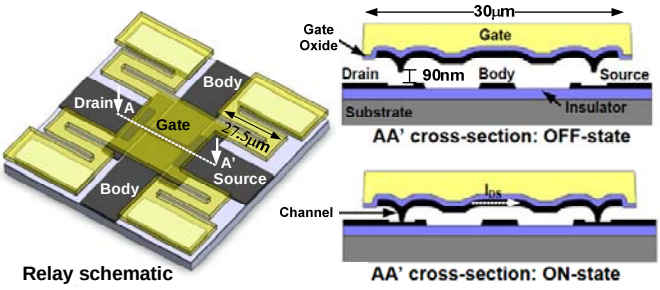

Fig. 1. Structure of a 4-terminal MEM relay device and its cross-section in the *off*- and *on*-states.

III. Energy-efficiency Analysis and Comparisons

Fig. 2 shows the basic structures and relevant parameters for CMOS logic power-gated by MOSFET or MEM relay headers. For a given amount of time that the CMOS logic is in sleep (T_{off}) or active (T_{on}) mode, the energy per power gate switching cycle for MOSFET (E_M) and relay (E_R) gating is:

$$E_M = (I_{on}T_{on} + U_M I_{off} T_{off})V_{EXT,M} + U_M C_M V_{IO}^2 + (U_M \gamma C_M + \beta C_L)V_{DD}V_{EXT,M} \quad (1)$$

$$E_R = I_{on}T_{on}V_{EXT,R} + U_R C_R V_{R,eff}^2 + (U_R \varepsilon C_R + \beta C_L)V_{DD}V_{EXT,R}, \quad (2)$$

where the value of the external supply $V_{EXT,M(R)}$ is set by the desired on-die supply V_{DD} and the IR drop through the power gate – i.e.:

$$V_{EXT,M(R)} = V_{DD} + I_{onp}R_{on,M(R)}/U_{M(R)}. \quad (3)$$

$R_{on,M(R)}$ is the *on*-resistance of a unit MOSFET (relay) switch, $C_{M(R)}$ is the gate capacitance of the unit switch, and $U_{M(R)}$ the upsizing factor (i.e. width for the MOSFET, number of parallel switches for the MEM relay power gate). V_{IO} is the I/O voltage driving the gate capacitance C_M, and $V_{R,eff}$ the effective relay gate voltage obtained from pull-in (V_{PI}) and pull-out (V_{PO}) voltages [2]: $V_{R,eff}^2 = V_{PI}\cdot(V_{PI} - V_{PO})$. I_{off} is the leakage current of a unit MOSFET switch. The average and peak load currents, I_{on} and I_{onp}, are equal to $\alpha f C_L V_{DD}$ and $\alpha_p f C_L V_{DD}$, respectively.

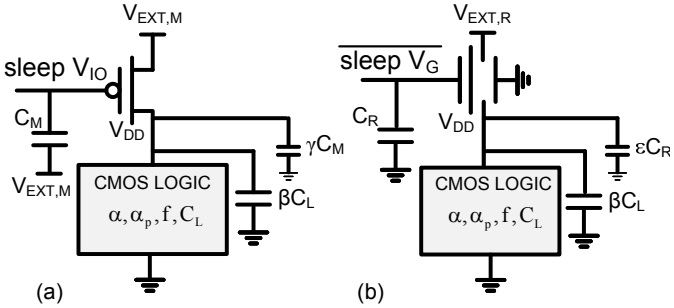

Fig. 2. MOSFET (a) and MEM relay (b) power gates. α and α_p are the average and peak activity factors of the CMOS logic, f is the operation frequency and C_L the CMOS logic load capacitance. β is the ratio of supply capacitance (at V_{DD}) and CMOS logic load capacitance, while γ and ε are the ratios of drain and gate capacitances of MOSFET (C_M) and MEM relay (C_R) headers.

TABLE I

CURRENT AND SCALED MEM RELAY TECHNOLOGY PARAMETERS [1,2]

Parameter	Current MEM Relay	Scaled Relay
$R_{on,R}$	2 kΩ	2 kΩ
C_R	300 fF	30 fF
V_{PI}	7 V	2.33 V
V_{PO}	5 V	2 V
$V_{R,eff}$	3.75 V	0.88 V
Pitch	100 μm	20 μm

TABLE II

CMOS TECHNOLOGY PARAMETERS FROM STANDARD AND PREDICTIVE MODELS [5]

CMOS process (nm)	C_M (fF/μm)	$R_{on,M}$ (kΩ·μm)	I_{off} (A/μm)		V_{NOM} (V)
			$V_{IO}=V_{NOM}$	$V_{IO}=V_{EXT,M}$	
32	0.9	0.55	336n	336n	1
45	1.2	1	115n	115n	1
65	1.7	1.1	1.3n	37.1n	1.1
90	2.2	1.6	54p	3.15n	1.2
130	2.5	2.4	22p	378p	1.2
180	1.6	2.8	3p	48p	1.8
250	1.3	4.5	1.5p	9.6p	2.5

Tables I and II provide parameters for current and scaled relays and a range of CMOS technologies for a target V_{DD} of 1 V.

The principal variable that a designer can adjust to minimize the energy consumed in each power-gating cycle is $U_{M(R)}$, whose optimum value can be obtained as follows:

$$\frac{dE_M}{dU_M}=0\rightarrow U_{M,OPT}\approx I_{on}\sqrt{\frac{kR_{on,M}T_{on}}{I_{off}T_{off}V_{DD}+C_MV_{IO}{}^2+\gamma C_MV_{DD}{}^2}} \quad (4)$$

$$\frac{dE_R}{dU_R}=0\rightarrow U_{R,OPT}\approx I_{on}\sqrt{\frac{kR_{on,R}T_{on}}{C_RV_{R,eff}{}^2+\varepsilon C_RV_{DD}{}^2}}, \quad (5)$$

where $k = I_{onp}/I_{on}$ and it was assumed that $\alpha f T_{on} >> \beta$. In order to clearly compare the performance of the two power gating schemes, it is useful to examine the normalized amount of energy each one loses relative to the *on*-state energy with ideal power gating ($E_{on}=V_{DD}I_{on}T_{on}$):

$$\delta E_M\approx 2\sqrt{\left(\frac{kR_{on,M}(P_{leak,M}T_{off}+E_{G,M})}{T_{on}V_{DD}{}^2}\right)}\left(1+\frac{\beta}{2\alpha fT_{on}}+\sqrt{\frac{kR_{on,M}I_{off}{}^2T_{off}{}^2}{4T_{on}(P_{leak,M}T_{off}+E_{G,M})}}\right)+\frac{\beta}{\alpha fT_{on}} \quad (6)$$

$$\delta E_R\approx 2\sqrt{\left(\frac{kR_{on,R}E_{G,R}}{T_{on}V_{DD}{}^2}\right)}\left(1+\frac{\beta}{2\alpha fT_{on}}\right)+\frac{\beta}{\alpha fT_{on}}, \quad (7)$$

where $\delta E_{M(R)} = (E_{M(R)} - E_{on})/E_{on}$, $E_{G,M} = C_M(V_{IO}{}^2 + \gamma V_{DD}{}^2)$, $E_{G,R} = C_R(V_{R,eff}{}^2 + \varepsilon V_{DD}{}^2)$, and $P_{leak,M} = V_{DD}I_{off}$.

Using (6-7) with parameters from a standard 90 nm CMOS process and our current relays [2,3] (Tables I and II), Fig. 3(a) shows the energy ratio of designs with MOSFET and MEM relay power gates versus T_{off} for fixed T_{on}. For short T_{off} (< 1 ms), the increased switching energy and the energy lost due to IR drop of the relay-based power gate outweigh its leakage reduction benefit. However, even with current relays, for T_{off} > 1 ms and T_{on} > 100 ns, the relay's negligible leakage continuously reduces the total energy as *off*-time is increased.

As shown in Fig. 3(b), scaling the relays to dimensions comparable to current mass-produced MEMS [6] in order to reduce their capacitance and operating voltages enables the relays to begin accruing energy savings at a substantially lower

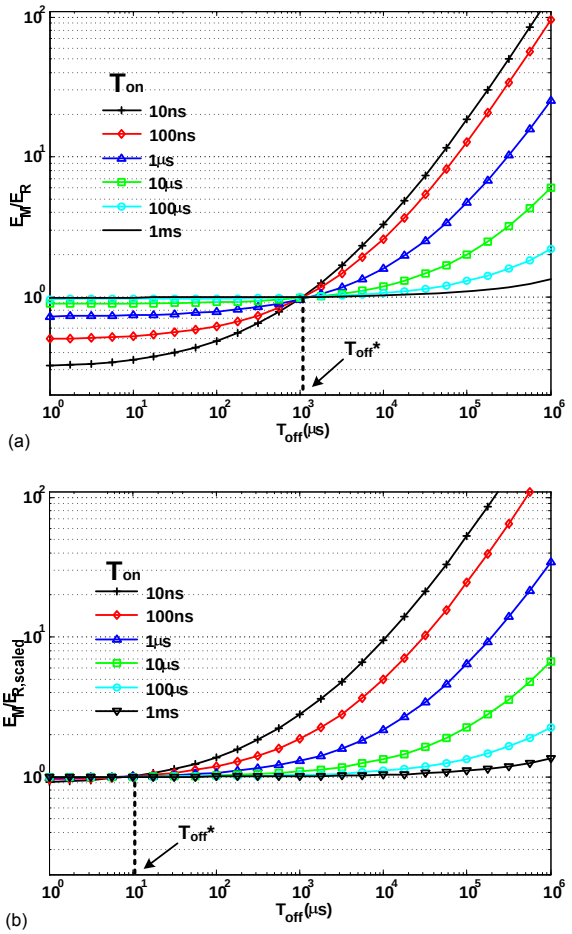

Fig. 3. Energy ratio vs. T_{off}, for various T_{on}, for designs power gated with (a) 90 nm MOSFETs and current MEM Relays, and (b) 90 nm MOSFETs and scaled MEM Relays. Here, α and α_p are 0.1 and 0.5, respectively, f = 1GHz, β = 5, γ = 1 and ε = 0.1.

T_{off} of 10μs (a 100x improvement). In fact, the minimum T_{off} at which relay power gating provides savings over MOSFET power gating (i.e., $T_{off}{}^*$ in Fig. 3) is well predicted by:

$$T_{off}^*=\frac{R_{on,R}C_R}{R_{on,M}C_M}\left(\frac{1+\frac{\beta}{2\alpha fT_{on}}}{1+\frac{\beta}{2\alpha fT_{on}}+\sqrt{\frac{kR_{on,M}I_{off}{}^2T_{off}{}^2}{4T_{on}(P_{leak,M}T_{off}+E_{G,M})}}}\right)^2\frac{E_{G,R}}{P_{leak,M}}-\frac{E_{G,M}}{P_{leak,M}}\approx \quad (8)$$

$$\approx\left(\frac{R_{on,R}C_R}{R_{on,M}C_M}\cdot\frac{E_{G,R}}{E_{G,M}}-1\right)\frac{E_{G,M}}{P_{leak,M}},$$

which states that the cross-over time constant is set by the ratio of switching energy overheads and MOSFET leakage power. The value of T_{off} for achieving a required energy improvement can also be obtained by revisiting (6-7) to calculate E_M/E_R:

$$G=\frac{E_M}{E_R}\approx\frac{\frac{kR_{on,M}I_{off}T_{off}}{T_{on}V_{DD}}+\frac{\beta}{\alpha fT_{on}}+1}{2\sqrt{\left(\frac{kR_{on,R}E_{G,R}}{T_{on}V_{DD}{}^2}\right)}\left(1+\frac{\beta}{2\alpha fT_{on}}\right)+\frac{\beta}{\alpha fT_{on}}+1}, \quad (9)$$

$$\approx\frac{kR_{on,M}I_{off}}{V_{DD}}\cdot\left(\frac{T_{off}}{T_{on}}\right)$$

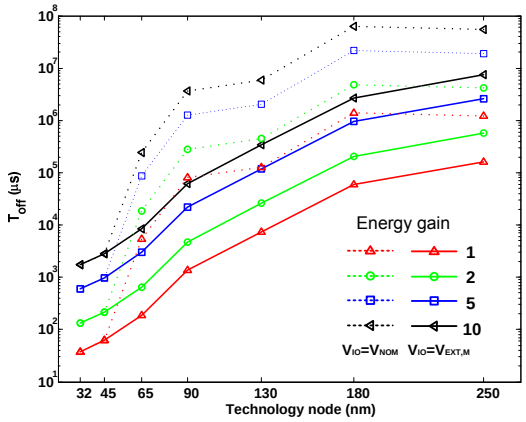

Fig. 4. Required T_{off} for specific energy gains of relay power gating over MOSFET power gating for different CMOS processes (T_{on}=10 µs).

where we have assumed that $T_{off} > 1$ ms, $T_{on} > 100$ ns, and parameters α, k, f, and β as in Fig. 3.

Equation (9) makes it clear that as should be expected, in applications with large *off-on* ratios, the switching energy overhead of even today's relays becomes negligible. Therefore, the energy reduction from relay power gates is set entirely by the removed leakage energy, and this energy savings grows linearly with T_{off}/T_{on}.

It is interesting to note that in this regime, the energy improvement for relay power gates over MOSFETs is also a function of the peak-to-average ratio k of the *on*-state current. This stems from the fact that both relay and MOSFET power gates must be sized to achieve a certain IR drop under worst-case load conditions. In the case of MOSFET power gates, the increased transistor width (vs. a power gate sized for average load) leads to linearly increased leakage. In contrast, the number of parallel relays used to implement the power gate can be increased without impacting the *off*-state leakage.

We next examine the *off*-times across different technology nodes, assuming that even for the most advanced CMOS designs, I/O transistors with the properties of an older technology node can be utilized as power-gates in order to exploit their potentially lower leakage properties. In Fig. 4, parameters from standard and predictive CMOS models [5] and (8-9) are used to find the T_{off}^{*} and required T_{off} for energy gains between 1 and 10 for different MOSFET technology nodes, and for current relay technology.

If a separate rail (other than $V_{EXT,M}$) is used for V_{IO} and matched to the available CMOS power-gate devices, lower I_{off} and hence higher T_{off} are expected for a given energy gain, as in Fig. 4. This is especially true for 0.25 µm and 0.18 µm power gates, as their nominal voltages (2.5 V and 1.8 V) are substantially higher than the ~1 V $V_{EXT,M}$, resulting in substantial gate under-drive for the PMOS headers. However, the leakage suppression with these long-channel/thick-oxide I/O devices is limited by junction leakage.

Fig. 5(a) illustrates the energy gain of designs with MEM relay gating over designs with MOSFET gating as a function of T_{off}/T_{on} for different power gate technologies. In Fig. 5(b), a separate, higher voltage rail $V_{IO} = V_{NOM}$ is assumed to be available, resulting in lower energy gains for relay-gated vs. MOSFET-gated designs. Interestingly, as pointed out in

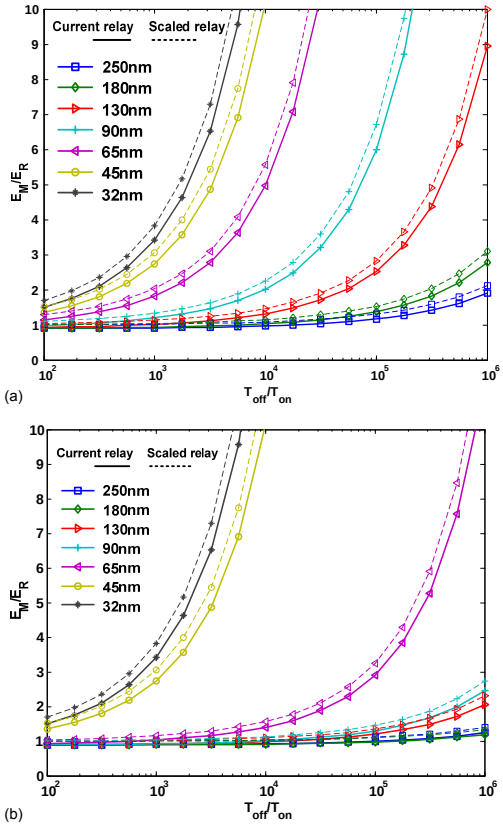

Fig. 5. Energy gain vs T_{off}/T_{on} for designs with MEM relay vs. MOSFET gating, for MOSFET power gates implemented in different CMOS processes. (a) $V_{IO} = V_{EXT,M}$, and (b) $V_{IO} = V_{NOM}$.

Fig. 3(b), the scaled relay technology does not significantly increase the energy gain (since the gain is mostly determined by the MOSFET leakage characteristics and *off-on* ratio), but does affect the *off-on* ratio at which the relay-gated designs begin to show savings over MOSFET-gated designs.

As mentioned previously, relay reliability is improved by the use of hard metals, which results in relatively high contact

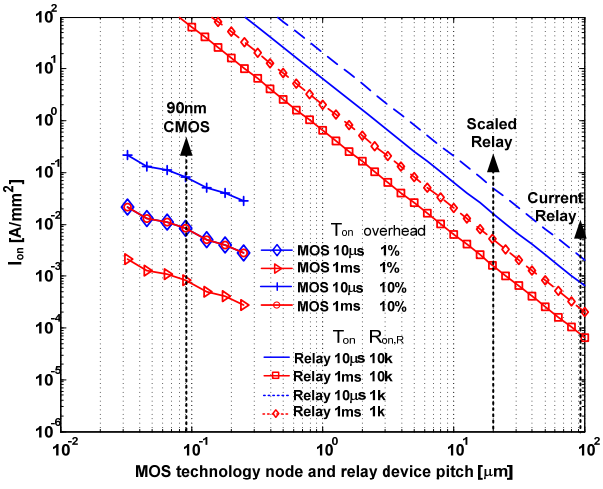

Fig. 6. Maximum logic current density as a function of MOSFET power-gate technology node and MEM relay device pitch. For MOSFET power gates, cases with power-gate device area overhead of 1% and 10% of the design area are shown. MEM relay power gates are assumed to be fabricated in the backend metallization layers, and thus incur no active area penalty.

resistance. For a given relay size, this resistance limits the current density that an array of relay power gates can deliver while maintaining the optimal voltage drop. As shown in Fig. 6, using current relays with ~100 µm device pitch, MEM power gating can be applied to CMOS chips with up to ~1 mA/mm^2 current density.

However, power gates built from moderately scaled relays with a device pitch of 17 µm [6], would support > 10-100 mA/mm^2 and would still fit into the same area as the CMOS chip they are driving. The relays could therefore be post-fabricated on top of the chip or integrated into the backend metallization layers with no penalty in the overall die area. For comparison, MOSFET power gates in older technology nodes (0.18 µm and 0.25 µm) that are most competitive to relays in terms of leakage suppression would provide similar current density, while consuming ~10% of the active area of the design.

IV. Experimental Demonstration

To experimentally demonstrate the feasibility of power-gating with current relay technology, we applied MEM relay power gating to a 90 nm CMOS chip [4] operating at V_{DD} = 0.6-1 V (I_{on} = 10-25 µA). Fig. 7 illustrates the waveforms of the MEM relay power-gating this chip with MEM gate voltages $V_{G,M}$ swinging between 5 and 7 V, with the inset indicating the chip's correct I/O activity during T_{on}. As this CMOS chip was not originally designed to support MEM relay power gating, in the *off*-state the chip's supply is limited by I/O ESD diode clamps to ~ 300 mV.

Beyond driving the power switches with an externally generated signal, Fig. 8 shows a simple MEM relay-based timer

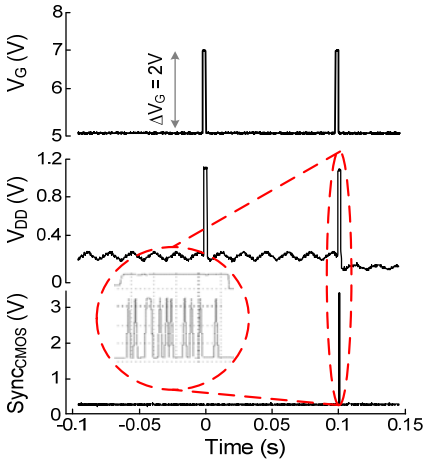

Fig. 7. Measured waveforms for MEM relay-gated CMOS chip [4]: MEM Relay gate voltage (top), supply voltage of the CMOS chip (middle), and synchronization signal from the CMOS chip (bottom).

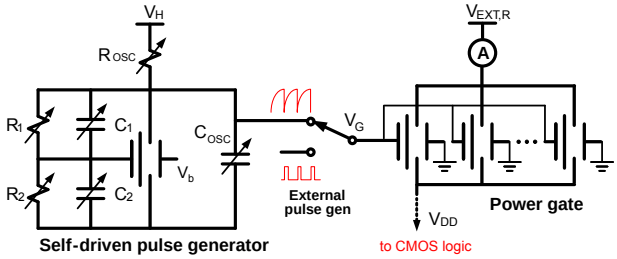

Fig. 8. Test setup for self-driven and external power gate pulse generation for the relay gating switch illustrated in Fig. 1.

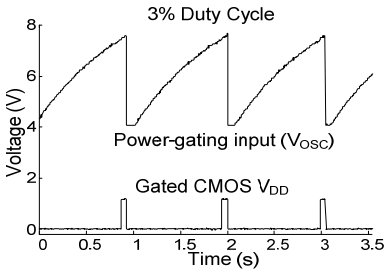

Fig. 9. Measured waveforms for MEM relay timer-gated CMOS chip [4].

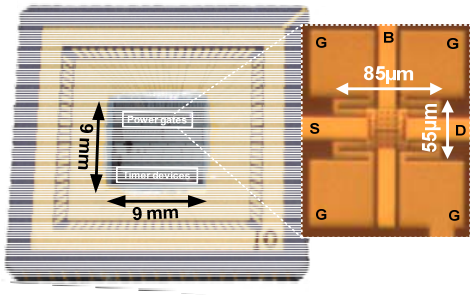

Fig. 10. Packaged MEM relay die photo and relay device micrograph.

circuit. The timer is based on a single-relay oscillator that enables autonomous operation, and was used to gate a CMOS chip as shown in Fig. 9. Although explicit capacitors would be removed in an optimized timer, current design uses adjustable RC elements to allow for tuning of period and duty cycle.

Fig. 10 illustrates the packaged MEM relay chip with devices configured to implement relay power gate and timer circuits. The inset shows a zoomed-in view of one of the 4-terminal relay devices from this chip.

V. Conclusion

An analytical comparative analysis and experimental demonstration illustrate that MEM relays even in their current state of technology (~ 7 V voltage and 100 µm pitch) can provide energy-reduction benefits over MOSFET power gates for *off*-periods > 1 ms. With relays scaled to current mass-produced MEMS device dimensions (~ 20 µm), the minimum *off*-period for which MEM relays provide energy-reduction benefit reduces to 10 µs and current densities greater than 10-100 mA/mm^2 can be supported.

Acknowledgments

This work was supported in part by the DARPA NEMS program, the C2S2 and MSD FCRP centers, BWRC sponsors, MIT CICS, and NSF Infrastructure Grant No. 0403427.

References

[1] A. Raychowdhury, J.-I. Kim, D. Peroulis, and K. Roy, "Integrated MEMS Switches for Leakage Control of Battery-Operated Systems," *CICC*, Sept. 2006.

[2] H. Kam *et al.*, "Design and reliability of a micro-relay technology for zero-standby-power digital logic applications," *IEDM*, Dec. 2009.

[3] R. Nathanael *et al.*, "4-Terminal Relay Technology for Complementary Logic," *IEDM*, Dec. 2009.

[4] F. Chen, A. P. Chandrakasan, and V. Stojanović, "A Signal-agnostic Compressed Sensing Acquisition System for Wireless and Implantable Sensors," (unpublished)

[5] Predictive Technology Models, http://ptm.asu.edu

[6] P. van Kessel *et al*, "A MEMS-based Projection Display," *Proceedings of the IEEE*, vol. 86, no. 8, Aug. 1998.