Low Threshold CMOS Circuits with Low Standby Current

Mircea R. Stan
University of Virginia, Electrical Engineering Department
Charlottesville, VA 22903
mircea@virginia.edu

Abstract

Multi-Voltage CMOS (MVCMOS) is a design methodology for very low power supply voltages that uses low-threshold transistors in series with the supply rails. The control voltages on the gating transistors need to be outside of the V_{dd} - V_{ss} range (hence the name MVCMOS) in order to reduce the standby current, but the resulting circuits operate at lower supply voltages and have a lower area overhead than the previously proposed Multi-Threshold CMOS (MTCMOS).

1 Introduction

Low-power design techniques try to reduce the power *dissipation* of high-performance systems (which impacts their packaging and heat removal costs) or the power *consumption* of portable equipment (which directly relates to size and battery life). Lowering the power supply voltage is one of the preferred methods as it has a quadratic effect on reducing power in CMOS [3] but this unfortunately increases the gate delay. Reducing V_{th} of the transistors in order to preserve performance presents no problem in theory and has been reported [2, 6], but obtaining such small V_{th} values with small variance remains a manufacturing challenge.

Even if the threshold can be accurately controlled, a big problem for low-power design is the exponential increase in the "off" current with reduced V_{th}. For circuits that operate *continuously* it has been shown that an optimal power-delay product can be obtained by reducing V_{th} such that the DC power consumption becomes approximately equal to the dynamic power [2]. For *event-driven* applications though, which are idle for long periods of time, the large current in sleep mode becomes unacceptable and negatively impacts the overall power budget. Multi-threshold CMOS (MTCMOS) [7, 8] has been proposed as a general technique for reducing the standby current by using two types of transistors, low-threshold transistors for fast operation in normal mode and high-threshold transistors for reducing the sleep mode "off" current. The high-threshold transistors are placed as gating transistors in series with the supply, creating *virtual* power rails for the rest of the circuit being implemented with low-threshold transistors [7]. The high-threshold gating transistors are always "on" for normal operation and, if sized correctly, don't have a significant impact on the circuit speed. With the circuit in sleep mode, the gating transistors are turned "off" and the overall leakage is limited by their small subthreshold current.

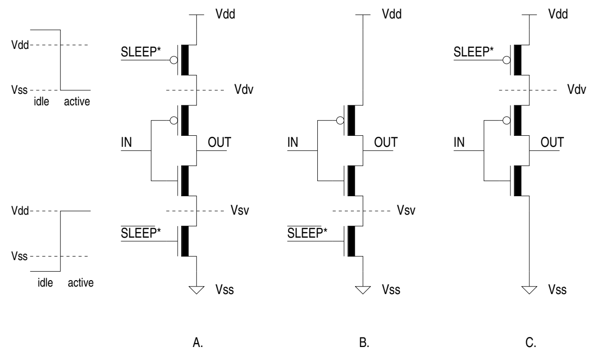

Figure 1: MVCMOS inverters: A. with virtual power and ground, B. only with virtual ground, C. only with virtual power. Only low threshold transistors are used but the SLEEP signals need to be outside the V_{ss} - V_{dd} range.

The rest of the paper shows how a straightforward circuit technique similar to MTCMOS but using only low-threshold transistors has better performance and less area by using control voltages at values outside of the V_{ss} - V_{dd} range.

2 Multi-Voltage CMOS (MVCMOS)

The proposed Multi-Voltage CMOS (MVCMOS) has the same circuit topology as MTCMOS (see figure 1) but uses low-threshold instead of high-threshold gating transistors and this has many implications for the gate performance. Similar to MTCMOS, in normal mode the gating transistors are "on", while in sleep mode they are turned "off", but in order to reduce the "off" current, MVCMOS needs to use control voltages larger than V_{dd} for the PMOS and smaller than V_{ss} for the NMOS. The advantages of MVCMOS can be summarized as:

- For MVCMOS the supply voltage can get much smaller than for MTCMOS due to the absence of high-threshold transistors, correct operation at $0.5V$ and lower being possible. The same effect could be obtained for MTCMOS by using a separate network driving the gating transistors with a separate (larger) V_{dd}.
- Because the gating transistors for MVCMOS have low V_{th} their size can be smaller than the high V_{th} MTCMOS transistors for the same current drive. Furthermore, the MVCMOS gating transistors can be *distributed* which results in a very efficient layout with minimal area overhead by using diffusion sharing. A similar diffusion sharing may be more difficult for MTCMOS depending on how well the implant registration mask can be positioned between two transistors.

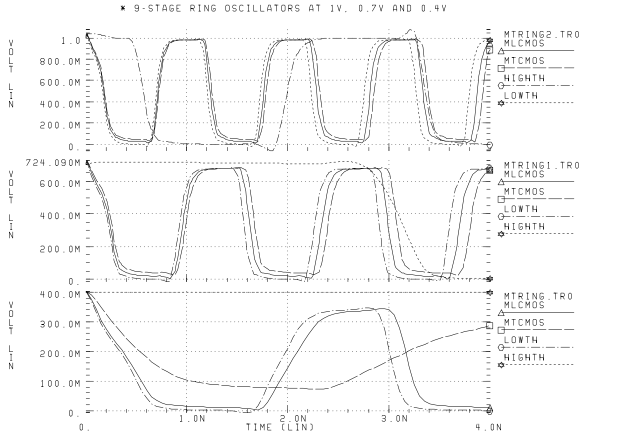

Figure 2: Regular high threshold, low threshold, MTCMOS and MVCMOS ring oscillators at different supply voltages: A. $1V$, B. $0.7V$, C. $0.4V$.

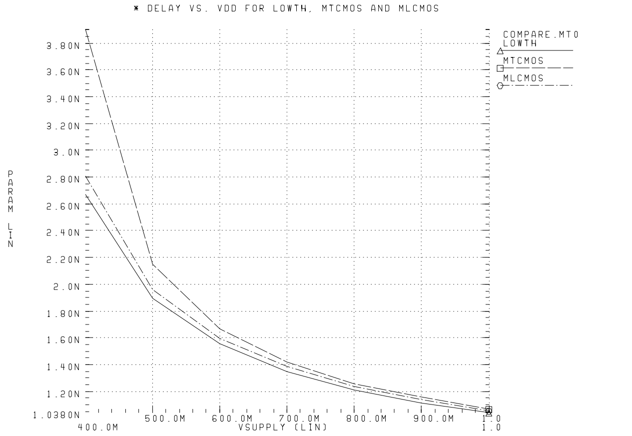

Figure 3: Regular low threshold, MVCMOS and MTCMOS delay at different supply voltages.

The simulations in this paper use HSPICE with BSIM3v3 models [1] for nine-stage ring oscillators in a 0.35μ CMOS technology. The models for high and low threshold transistors differ only in the `Vth0` parameter in the model which is $0.6322V$ and $0.2322V$ for high and low threshold NMOS and -0.6733 and -0.2733 for high and low threshold PMOS, respectively. The PMOS is sized three times larger than the NMOS (1.8μ vs. 0.6μ) for symmetric noise margins, but probably not optimal for low power.

3 MVCMOS Sub-Volt Operation

The major advantage of MVCMOS is that it can operate at very low supply voltages. Figure 2 shows the output of regular high threshold, low threshold, MTCMOS and MVCMOS ring oscillators (using only NMOS gating transistors nine times minimum size) at three different V_{dd} values, $1V$, $0.7V$ and $0.4V$. The regular circuit with high threshold transistors ceases to function at $0.7V$ and MTCMOS ceases to function at $0.4V$, but MVCMOS is just slightly slower than the regular low threshold circuit even at $0.4V$. The delays of MTCMOS, MVCMOS and regular low threshold oscillators (divide by 18 for inverter delay) as a function of the power supply voltage are shown in figure 3.

MVCMOS in sleep mode requires control voltages lower than V_{ss} for the gating NMOS and higher than V_{dd} for the gating PMOS. Fortunately there is no hot-electrons problem since we only apply the larger voltages to the *gates* of the gating transistors. Latchup should not be a problem either if we can make sure that the rise and fall times for the control voltages are slow enough to reduce capacitive coupling. We have simulated the MTCMOS and MVCMOS ring oscillators with only NMOS gating transistors in sleep mode (see figure 4) and there is no reason to reduce the negative control voltage to less than $-0.5V$ because at some point the subthreshold current becomes negligible and the leakage of the reverse biased drain junctions becomes dominant. This should be perfectly safe with no oxide breakdown for all current technologies. The sleep voltage can be either generated off-chip or on-chip and the edges of the signal should be very slow to avoid latchup. A possible on-chip circuit is the classical substrate pump used in memory design [4].

4 Sizing the Gating Transistors

The "virtual" power rails for MVCMOS (as for MTCMOS) have a much higher impedance than the true power rails and will unavoidably "bounce" which means that:

- The speed of the logic gates is reduced as the effective supply voltage is reduced by the bounce magnitude.
- The noise margins are reduced which becomes a problem especially when communicating with gates with different "virtual" rails.

Figure 5 shows the ring oscillator delay at $V_{dd} = 0.5V$ for MVCMOS with both PMOS and NMOS, only NMOS or only PMOS gating transistors. Using both transistors is the slowest solution and, although the PMOS is sized three times larger than the NMOS, the circuit using only a gating NMOS is the fastest. Using just a gating NMOS has the lowest area overhead and is also the fastest but has a virtual ground which will bounce. Designers generally prefer having a noisy supply and maintain a "clean" ground which can be obtained using a gating PMOS even if it is a slower solution. A tool was proposed for sizing MTCMOS gating transistors for speed [5], but we believe that a more critical issue is sizing for noise margins. Figure 6 shows the variation of ground bounce with the size of the sleep transistors. The ground bounce is a noise problem because many logic gates share a *centralized* gating transistor, hence the same virtual ground, which is a characteristic carried over from the MTCMOS design style.

Instead of centralized gating transistors there is also the option of using *distributed* gating transistors, one for each logic gate. Such a distributed solution will completely eliminate the ground bounce noise problem but may result in a slower circuit as can be seen in figure 7. The layout with distributed sleep transistors can be done very efficiently using *diffusion sharing* with the logic transistors. In this case the layout for the pull-down network (PDN) for a gated inverter looks exactly like the PDN of a standard CMOS NAND gate and represents a small area increase. Distributing the gating transistors with every gate does not incur a large penalty which suggests the following design strategy:

- All the gates that are not on the critical path should have *distributed* gating transistors in order to eliminate any noise problems due to ground bounce.

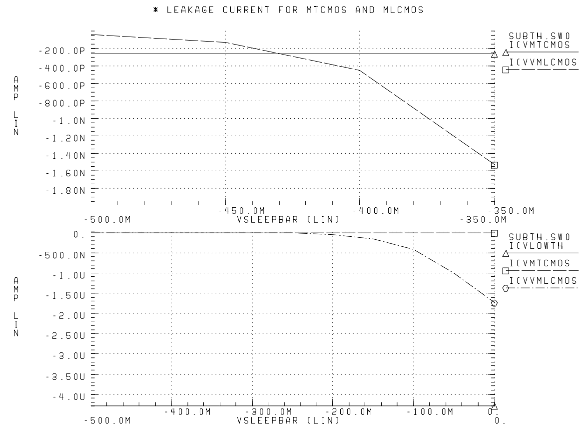

Figure 4: Subthreshold current for MTCMOS with SLEEP at $0V$ and MVCMOS with SLEEP between $0V$ and $-0.5V$.

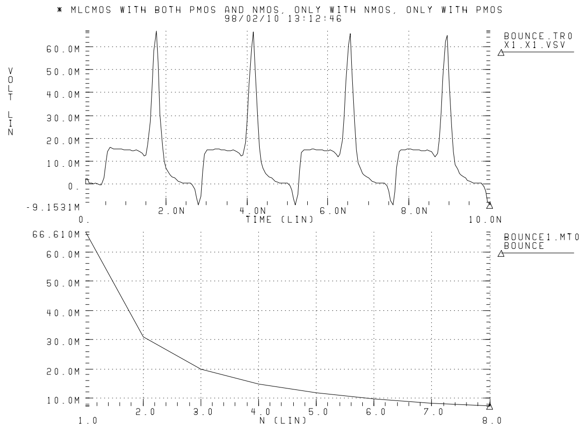

Figure 6: A. Ground bounce for MVCMOS ring oscillator, B. Bounce as a function of the size of the gating transistor.

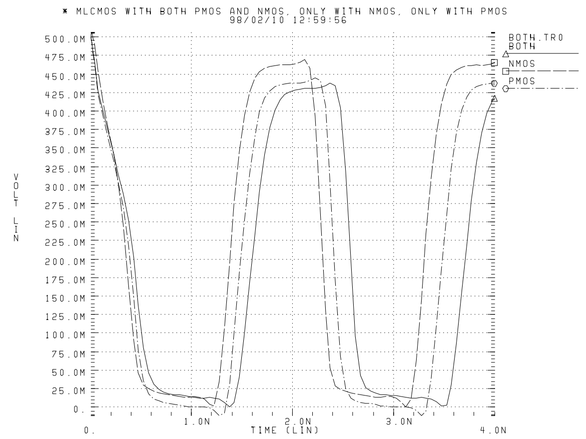

Figure 5: MVCMOS delay with one or two gating transistors at $V_{dd} = 0.5V$: A. both, B. PMOS only, C. NMOS only.

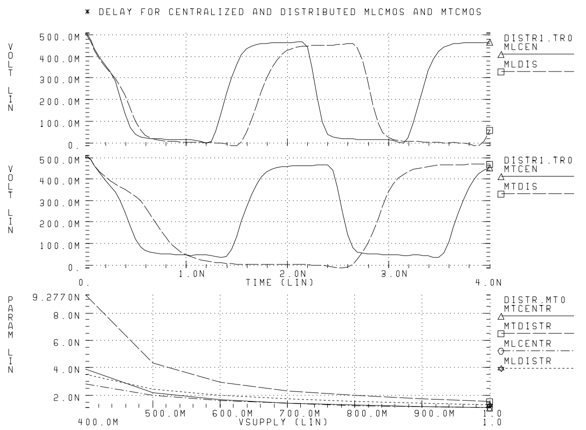

Figure 7: Delay for centralized vs. distributed sleep transistors for MTCMOS and MVCMOS: A. Time representation, B. As a function of supply voltage.

- Only gates that are critical should use common gating transistors for faster operation. A thorough analysis is required for noise and performance issues but since the number of transistors is small this should be easier.
- Gates that switch at the same time do not benefit (speed-wise) from sharing gating transistors, gates that switch at different moments do.

5 Conclusion

We have presented Multi-Voltage CMOS (MVCMOS), a design methodology for very low supply voltages with good performance and low leakage during sleep. The area overhead is smaller than for MTCMOS and the ground bounce problems can be eased by distributing the low threshold gating transistors. As suggested by a reviewer the gating transistors could be overdriven both "on" and "off" which requires even less width for the same current and reduces even more the area penalty for gating. More work is needed for analyzing dynamic and sequential MVCMOS circuits.

Acknowledgements
This work was supported in part by NSF CAREER grant MIP-9703440. I thank Sreekanth Nallagatla for helping with some of the initial simulations.

References

[1] U. Berkeley. BSIM3v3.1 SPICE MOS device models, 1997. http://www-device.EECS.Berkeley.EDU/ bsim3/.
[2] J. B. Burr and A. M. Peterson. Ultra low power CMOS technology. In *NASA VLSI Design Symposium*, 1991.
[3] A. P. Chandrakasan, S. Sheng, and R. W. Brodersen. Low-power CMOS digital design. *Proceedings of the IEEE*, 81, 1992.
[4] K. Itoh, K. Sasaki, and Y. Nakagome. Trends in low-power ram circuit technologies. *Proceedings of the IEEE*, 83(4):524–543, Apr. 1995.
[5] J. Kao, A. Chandrakasan, and D. Antoniadis. Transistor sizing issues and tool for multi-threshold CMOS technology. In *Design Automation Conference*, Anaheim, CA, June 1997.
[6] D. Liu and C. Svensson. Trading speed for low power by choice of supply and threshold voltages. *IEEE Journal of Solid-State Circuits*, 28(1):10–17, Jan. 1993.
[7] Mutoh, S. et al. A 1V power supply high-speed digital circuit technology with multithreshold voltage CMOS. *IEEE Journal of Solid-State Circuits*, 30(8):847–854, Aug. 1995.
[8] S. Shigematsu et al. A 1-V high-speed MTCMOS circuit scheme for power-down application circuits. *IEEE Journal of Solid-State Circuits*, 32(6):861–869, June 1997.