MTCMOS Sequential Circuits

James Kao
Massachusetts Institute of Technology
Cambridge, MA 02139
jkao@mtl.mit.edu

Anantha Chandrakasan
Massachusetts Institute of Technology
Cambridge, MA 02139
anantha@mtl.mit.edu

Abstract

Multi-threshold CMOS is an increasingly popular circuit technique that enables high performance and low power operation, but requires sequential circuit structures that can retain state during standby modes. This work presents an in depth analysis of potential leakage paths, and yields several efficient MTCMOS flip flop implementations including a novel approach that utilizes a leakage feedback gate to enable state retention during standby modes in dynamic flip flops.

1. Background

In order to reduce overall power consumption, a well known technique is to scale supply voltages. However, to maintain performance, device threshold voltages must scale as well, which will cause subthreshold leakage currents to increase exponentially. Controlling subthreshold leakage has been explored significantly in the literature, especially in the context of reducing leakage currents in burst mode type circuits, where the system spends the majority of the time in an idle standby, or sleep, state where no computation is taking place. MTCMOS, or multi-threshold CMOS has been proposed as a very effective technique for reducing leakage currents during the standby by state by utilizing high V_t sleep devices to gate the power supplies of a low V_t logic block[1][2].

Although MTCMOS circuit techniques are effective for controlling leakage currents in combinational logic, a drawback is that it can cause internal nodes to float, and cannot be directly used in standard memory cells without corrupting stored data. As a result, several researchers have explored possible MTCMOS latch designs that can reduce leakage currents yet maintain state during the standby modes[3-7]. Some of this previous work did not fully characterize the sneak leakage paths that can arise in sequential circuits, and other techniques, such as the "balloon" flip flop require excessively complex timing requirements to store and retrieve memory states when entering and exiting the sleep state.

In this work, a more fundamental understanding of leakage currents in sequential MTCMOS circuits is explored, which leads to several possible master slave flip flop implementations that are streamlined and high performance, yet also eliminate all possible sneak leakage paths during the standby mode. A new type of flip flop utilizing a leakage feedback mechanism is also presented, which can be applied to dynamic flip flops that can actually retain state during the standby state.

2. MTCMOS Flip Flop Leakage Paths

An MTCMOS flip flop that can retain memory during the standby state is shown below in Fig. 1. This implementation is a straightforward extension of a standard master slave flip flop, where leakage paths are carefully eliminated.

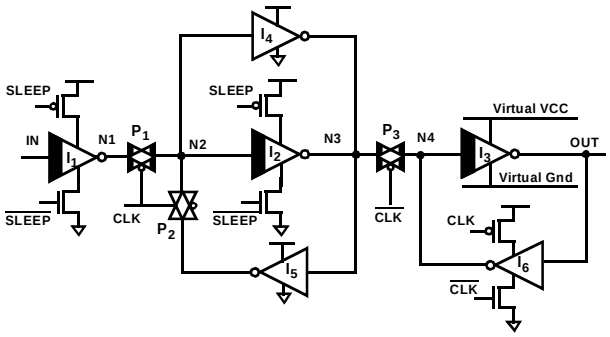

Figure 1. MTCMOS circuit structure (drives another MTCMOS block) with input buffer.

The basic latch structure used in this flip flop is similar to that of the MTCMOS latch first presented in [2], where a helper high V_t inverter is used to preserve state during the standby mode. However, here low V_t devices are used throughout the critical path while high V_t devices are used in peripheral circuitry that simply holds state. By making passgate P_2 high V_t, the performance of the master latch actually improves because during the transparent state, I_1 would not have to fight against I_5 through an off low V_t passgate. However, the drawback is that when CLK goes high, I_5 and P_2 need to be strong enough to correctly hold the state at node N_2 because N_2 and N_1 might be driven to opposite rails. Other than this sizing condition, the active operation of this flip flop is straightforward. During the standby state, all leakage current paths are also eliminated to minimize power dissipation.

One of the problems with sequential circuits that utilize feedback and parallel devices is that sneak leakage paths may exist. The flip flop of Fig. 1 is a good example of how distributed high V_t sleep transistors and dual polarity sleep devices are needed to eliminate sneak leakage currents during the standby condition. Sneak leakage paths can arise in MTCMOS circuits whenever the output of an MTCMOS gate is electrically connected to the output of a CMOS gate. In fact, the interfacing between MTCMOS type circuits and CMOS type circuits is what gives rise to potential leakage paths. For example, if a datapath block is implemented with only MTCMOS gates, then a single high V_t switch (either PMOS or NMOS) is sufficient to eliminate subthreshold leakage currents during the standby state because all current paths from V_{CC} to GND must pass through an off high V_t device. However, if CMOS gates and MTCMOS gates are combined together, sneak leakage paths can arise that bypass the off high V_t devices.

Fig. 2 shows two examples that illustrate how sneak paths may exist if only one polarity sleep device is employed and the output of a CMOS gate is connected (directly or through low V_t passgates) to the output of a MTCMOS gate.

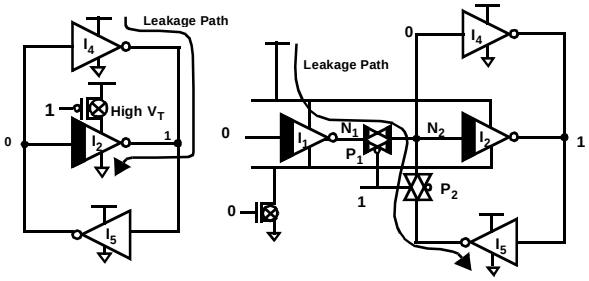

Figure 2. Potential sneak leakage paths with only one polarity sleep device (dual cases leak

In the first case, a direct leakage path exists that circumvents the off high V_t device completely. In the second scenario, a sneak path can also exist if a low V_t passgate is placed between an MTCMOS gate and a CMOS gate. If both ends of the passgate are driven to opposite polarities, then a large leakage current will exist. Even if the input to Fig. 2B were to go "high" instead, there would still be a potential sneak leak path through the off low V_t PMOS of I_1 in series with the turned-off (but still leaky) passgate P_1 to GND. These sneak leakage paths resulting from parallel gate constructions can easily be controlled by using both PMOS and NMOS polarity sleep transistors to ensure that both paths from V_{CC} and to GND are cutoff for MTCMOS circuits.

However, in some cases, simply using both polarity sleep devices is not enough. If high V_t sleep devices are shared among multiple MTCMOS gates, then there exists the possibility of sneak leakage paths that arise due to reverse conduction paths. These scenarios can arise when there are several MTCMOS - CMOS gate pairs that have a common output node, or are connected with low V_t passgates, as illustrated in Fig. 3.

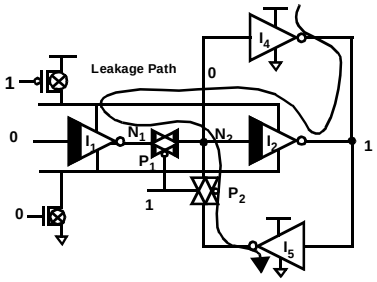

Figure 3. Reverse conduction sneak leakage path.

Such a sneak leakage path originates from V_{CC} of one of the CMOS gates, travels through the virtual power lines and exits to GND through another CMOS gate. The flow of current thus travels through an MTCMOS gate in a reverse conduction path (a PMOS current flowing from the device output towards virtual V_{CC}, or a NMOS device current flowing from virtual GND to the device output). In general, one can eliminate this type of sneak leakage path for a block with a common virtual power and virtual ground lines by ensuring that no more than one CMOS-MTCMOS gate pair has a common output node. Another way to eliminate this leakage path is to use separate local high V_t sleep devices for those MTCMOS gates with outputs that are electrically connected to CMOS outputs.

A final sneak current path can arise when an MTCMOS gate directly drives a CMOS gate. During the standby state, the output of the MTCMOS gate can float and as a result can cause short circuit currents to exist. For the flip flop of Fig. 1, nodes N_2 and N_3 are always driven during the standby state, so these nodes can directly drive CMOS gates I_4 and I_5. However, by construction node OUT can float so I_6 must be disconnected from the power supplies during the standby state and can only be interfaced to MTCMOS gates downstream.

The flip flop of Fig. 1 utilizes local sleep devices of both polarities to effectively eliminate sneak leakage paths. Although sleep transistor area can be large because sleep devices cannot be shared among multiple blocks (like in combinational MTCMOS circuits), the penalty is not too severe because having local control of sleep devices makes it easier to size the sleep devices and also decouples noise from different switching blocks from sensitive storage nodes. However, if area is of premium importance, one can modify the architecture of the flip flop of Fig. 1 by simply disconnecting I_4 and I_5 from the internal node N_3, which disconnects CMOS outputs from MTCMOS output nodes, while still providing a latch recirculation path during the opaque state. In this case, local high V_t sleep devices are needed for I_1, but a shared virtual V_{CC} or virtual GND line with a common sleep transistor can be used for I_2 and I_3. In fact, only a single polarity shared sleep device is needed to eliminate leakage currents. In a large register for example, sharing a common NMOS sleep transistor among several flip flop and logic blocks can result in large area savings at the expense of more complicated sleep transistor sizing methodologies[2].

3. Leakage Feedback Gate

The most straightforward way to provide state retention in sequential circuits it to utilize high V_t parallel gates to recirculate data during the standby mode. However, the use of parallel high V_t inverters is detrimental to performance because these gates provide extra load to the critical path, but do not appreciably improve current drive. An alternative approach to maintain state during the standby mode is to utilize a leakage feedback gate, as shown below.

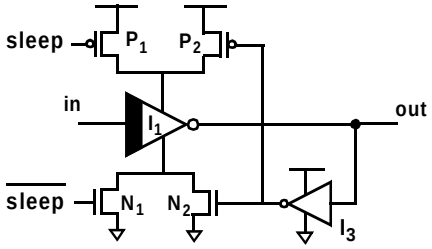

Figure 4. Leakage feedback gate

The important characteristic of this type of gate is that depending on the state of the latest output, one but not both, helper sleep devices (P_2 or N_2) is turned on. During the standby state, both high V_t sleep devices P_1 and N_1 are turned off, but only one of the helper sleep devices will be turned off. The helper sleep device that is kept on will simply correspond to the one that continues to drive the output signal to the appropriate rail. As a result, during the standby state the output node will still be driven to one rail or the other, yet because any path from power to ground must encounter a strongly turned off high V_t

device, the leakage currents will be reduced by several orders of magnitude. This provides a mechanism where an MTCMOS gate can be put in a low leakage state, yet still actively drive its outputs.

Even if the input signal to a leakage feedback gate does transition or float after the gate is placed in the standby mode, the output voltage will still be held to the same logic value (through a leakage path). This occurs because the output of the leakage feedback gate will be determined by the relative strengths of the leakage current through an "on" high V_t device in series with an "off" low V_t device pulling against the leakage current of an "off" high V_t device in series with an "on" low V_t device. As long as the high threshold voltages are larger than the low threshold voltages, then the path with the "on" high V_t device and "off" low V_t device will dominate and thus continue to pull the output voltage to the same rail as before. Because the functionality of this leakage feedback selectively enables either the virtual power or virtual ground lines depending on the last data present, leakage feedback gates must utilize local high V_t sleep devices. However, utilizing local sleep devices may already be desirable for sensitive circuits like flip flop and latches that should be decoupled from the switching activities of neighboring gates.

An immediate use of the leakage feedback gate is that it can be used as an interface circuit between MTCMOS and CMOS logic blocks. For example the last stage of any MTCMOS block can be implemented as a leakage feedback gate such that during the standby mode the output is still driven. As a result, this stage can safely drive a standard CMOS gate without creating short circuit currents due to floating inputs.

4. Leakage Feedback Static Flip Flop

Another use of the leakage feedback gate is to modify the static MTCMOS flip flop to eliminate the need for the parallel inverter to recirculate data. The leakage feedback structure can be used instead, which does not slow down the critical path because no extra capacitances are introduced on internal nodes as seen in Fig. 5. The addition of the helper sleep devices only add load to the outputs of I_4 and I_5 which are not part of the critical path, so the speed of the MTCMOS flip flop is not compromised.

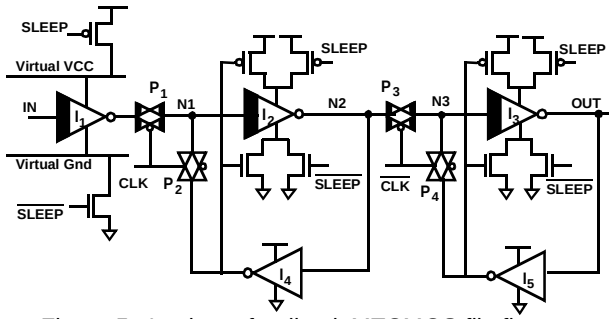

Figure 5. Leakage feedback MTCMOS flip flop (static)

During the active operation, the leakage feedback static flip flop operates like an ordinary master slave flip flop with performance superior to that of the original flip flop of Fig. 1. During the standby state, the leakage feedback gate stores the flip flop state in the master stage while clock is high. Furthermore the slave stage is configured such that the output node does not float during standby mode so that it can interface directly to CMOS devices. These two added functionalities can be implemented using leakage feedback gates without any loss in performance since no extra loading is introduced on the critical path. Since both nodes N_1 and N_3 are actively driven during the standby state, the leakage feedback gates actively hold their outputs as well, and thus functions exactly like the previous flip flop using high V_t recirculation paths.

5. Leakage Feedback Dynamic Flip Flop

One especially useful application of the leakage feedback gate is to implement dynamic flip flops that retain state during the standby mode yet still have very fast circuit performance during the active state.

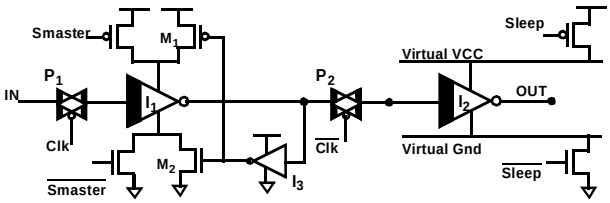

Figure 6. Leakage feedback dynamic flip flop

During the active switching mode, the leakage feedback dynamic flip flop operates like a conventional one, where the state of the master and slave latches are simply stored on the dynamic nodes at the inputs of I_1 and I_2. Because leakage currents will be large when using low V_t devices, the clock period must be made fast enough such that the node voltages can be stable over the clock periods of interest. By utilizing leakage feedback gates in a dynamic flip flop however, it will be possible to retain data during the standby state when CLK is high. This enables a timing methodology where standby gating can be performed so that a block can be stalled in time, and yet can be woken up again to complete the computation. Clock gating alone would not be sufficient because during the active state the flip flop's dynamic nodes would simply leak away, and data would be corrupted. The leakage feedback gate provides an added functionality to architectures with dynamic flip flops not available before. Conventional dynamic flip flops are incapable of maintaining data during the standby modes, so previous architectures that provide for block shut down periods must explicitly provide peripheral circuitry to retain state.

When the dynamic flip flop is placed in a sleep condition however, the leakage feedback mechanism retains the data even if the internal dynamic nodes change. However, because the input voltage to the leakage feedback gate can float, the timing requirement coming out of the sleep state requires that the phase of the master latch sleep signals, S_{master} and $\overline{S}_{master}$, lag the main sleep signals, Sleep and $\overline{Sleep}$, by one half cycle. This is because when transitioning from the sleep state to the active mode state, one cannot immediately turn on the master latch high V_t sleep transistors because that might accidentally cause the data to flip state. As a result, the master slave can exit the sleep condition only after the clock goes low and the slave stage latches the stored data.

6. Simulations

Simulations were performed on the various MTCMOS flip flop architectures in a 0.16u technology with high V_t approximately .15V and low V_t approximately .05V (defined the 10 nA @ 10um point). Fig. 7 illustrates how the leakage feedback gate of Fig. 4 can hold data during the standby state even when the

input varies during the sleep state.

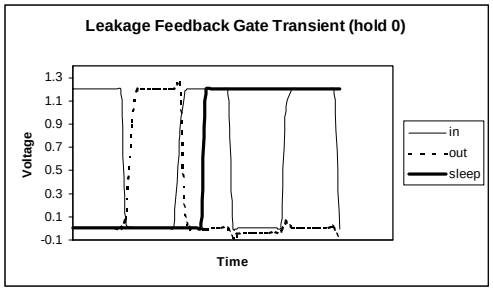

Figure 7. Leakage feedback gate transient behavior showing output tracking clock until sleep mode activates.

Fig. 8 shows simulations comparing the total delay (T_{setup} + T_{CQ}) of the flip flops of Fig. 1, Fig. 5, and Fig. 6 as function of the sleep transistor W/L ratio expressed as a percentage of the total width. As can be seen, the leakage feedback static flip flop shows a slight improvement over the MTCMOS flip flop of Fig. 1 because of reduced loading. This benefit is compounded by the fact that the leakage feedback flip flop can drive an ordinary CMOS output whereas the standard flip flop would have to modified and take another performance hit to provide this functionality. The delay through the leakage feedback dynamic flip flop is shown to be the fastest of all three circuits, as expected since loading of internal loads is minimized. Again, this dynamic flip flop can still retain state during the standby modes, which yields significant improvement over standard dynamic flip flops.

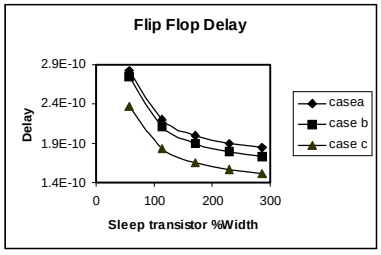

Figure 8. Flip flop delays (T_{setup} + T_{CQ}) for a) MTCMOS static FF, b) leakage feedback static FF, and c) leakage feedback dynamic FF.

Finally, Fig. 9 shows how leakage currents are reduced during the sleep condition. All three flip flop implementations are shown to significantly reduce leakage currents during standby operation, although the exact amount of leakage reduction is dependent on choice of technology and the selection of high V_t and low V_t levels. To the first order, the sleep condition leakage currents corresponds to the leakage of an all high V_t implementation.

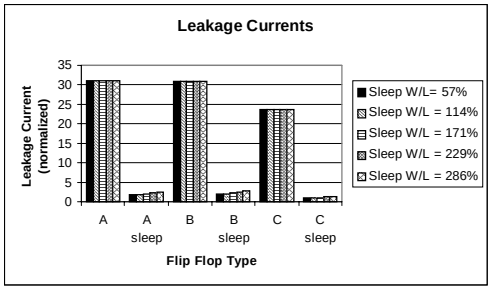

Figure 9. Active and standby leakage currents for each flip flop for varying W/L

7. Conclusion

This paper illustrates several MTCMOS flip flop implementations that can provide high speed active mode operation, yet retain state and exhibit low leakage currents during the standby mode. By thoroughly analyzing all possible leakage paths, it is possible to ensure that no sneak leakage paths exists, and also gives a framework for how MTCMOS and CMOS logic blocks can be combined. Finally, a unique leakage feedback gate was introduced, which can ensure low leakage conditions during the standby mode, yet still have actively driven outputs, regardless of the input behavior. This type of gate lends quite nicely to the development of a unique dynamic flip flop that can actually retain state during the standby mode.

8. References

[1] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, J. Yamada, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," IEEE JSSC, vol. 30, no. 8, pp. 847-854, August 1995.

[2] J. Kao, A. Chandrakasan, "Dual-Threshold Voltage Techniques for Low Power Digital Circuits," IEEE JSSC, vol.35, no. 7, pp. 1009-1018, July 2000.

[3] S. Shigematsu, S. Mutoh, Y. Matsuya, Y. Tanabe, J. Yamada, "A 1-V High-Speed MTCMOS Circuit Scheme for Power-Down Application Circuits," IEEE JSSC, vol. 32, no. 6, pp. 861-869, June 1997.

[4] H. Akamatsu, T. Iwata, H. Yamamoto, T. Hirata, H. Yamauchi, H. Kotani, A. Matsuzawa, "A Low Power Data Holding Circuit with an Intermittent Power Supply scheme for sub-1V MT-CMOS LSIs," 1996 Symposium on VLSI Circuits Digest of Technical Papers, pp. 14-15, 1996.

[5] K. Kumagai, H. Iwaki, H. Yoshida, H. Suzuki, T. Yamada, S. Kurosawa, "A Novel Powering-down Scheme for Low Vt CMOS Circuits," 1998 Symposium on VLSI Circuits Digest of Technical Papers, pp. 44-45, 1998.

[6] H. Makino, Y. Tsujihashi, K. Nii, C. Morishima, Y. Hayakawa, T. Shimizu, A Arakwa, "An Auto-Backgate-Controlled MT-CMOS Circuit," 1998 Symposium on VLSI Circuits Digest of Technical Papers, pp. 42-43, 1998.