

Suggested Layout, SP25.7

SP25.7: Skew-Tolerant Domino Circuits

David Harris*, Mark A. Horowitz

Stanford University, Stanford, CA

* also with Intel Corporation, Santa Clara, CA

As cycle time of chips shrinks and die size grows, clock skew measured as a fraction of the cycle time is increasing. Traditional domino circuits shown in Figure 1 are especially sensitive because skew must be budgeted in both half-cycles. The problem with such domino pipelines is that evaluation starts (indicated by the heavy dashed line) when the clock connected to the first gate in the half-cycle rises but the output needs to be valid before the clock on the output latch falls. In the worst case, the evaluate clock is late and the latch clock is early, decreasing time for logic. Many designers realize that some of the overhead can be reduced by using differential domino (also called dual rail) designs. An SR latch or “pipeline latch” [1] at the end of dual-rail circuits lessens sensitivity to the falling edge. Self-timed techniques [2] eliminate clocks and clock skew, but raise new issues of control overhead, timing assumption verification, and testability. This paper describes a methodology which boosts operating frequency by tolerating clock skew, eliminating latches from the critical path, and better balancing logic between phases of the pipeline.

The key observation is that logic and clock waveforms may be designed such that the domino gate is in evaluation whenever the inputs become valid, even under worst clock skew. Figure 2 shows a 2-phase skew-tolerant domino pipeline. Non-monotonic inputs must become valid before the earliest a skewed clock might rise, yet worst case timing assumes that the domino phase actually begins evaluation at the latest possible skewed clock. We eliminate latches from the critical path by guaranteeing that the domino gates in the subsequent clock phase can evaluate before the domino gates in the current phase fully precharge [2]. Thus, state is stored on dynamic nodes. Replacing the keeper of the first domino gate in each cycle with a cross-coupled inverter allows lossless stop-clock operation, as shown in Figure 3.

In an N-phase clocking system with period T, symmetry arguments dictate all phases should be identical except for an offset of T/N between phases. Each phase is high for an evaluation period t_e and low for a precharge period t_p . The nominal logic delay in each phase is T/N. Figure 4 illustrates pipeline timing requirements. The precharge period must be long enough to guarantee that a domino gate precharges and that the output of the subsequent static gate falls low before the subsequent domino gate re-enters evaluation, even under worst-case clock skew between ϕ_{1a} and ϕ_{1b} :

$$t_p \geq t_{prech} + t_{skew} \quad (\text{EQ 1})$$

The evaluation period must be long enough that the last domino gate in a phase does not precharge until the subsequent domino gate properly evaluates under worst skew between ϕ_{1b} and ϕ_{2a} :

$$t_e \geq T/N + t_{skew} + t_{hold} \quad (\text{EQ 2})$$

This required overlap of evaluation phases is t_{hold} . The maxi-

imum tolerable skew is therefore a function of the cycle time, overhead, and number of clock phases:

$$t_{skew-max} = [T(N-1)/N - t_{hold} - t_{prech}]/2 \quad (\text{EQ 3})$$

Assuming an aggressive cycle time T of 16 fanout-of-4 (FO4) inverter delays, a hold time of 0 (simulations show a small negative hold time under reasonable cell library restrictions), and required precharge time of 4 FO4 delays, a 2-phase system can tolerate a skew of 2 FO4 delays (with $t_e=10$), while a 4-phase system can tolerate a skew of 4 (with $t_e=8$).

Logic seldom can be partitioned to exactly fill a clock phase, so allowing time borrowing between phases simplifies design and eliminates wasted time at the end of phases. Furthermore, time borrowing automatically averages out delay variations along a path caused by process variation and modeling inaccuracies. The amount of time that a phase of logic can borrow is the skew tolerance minus the worst case skew, where t_{skewG} is the maximum skew between clock phases anywhere on the die and t_{skewL} is the maximum skew of a single phase within a local clock domain:

$$t_{borrow} = T(N-1)/N - t_{hold} - t_{prech} - t_{skewG} - t_{skewL} \quad (\text{EQ 4})$$

With the same delay assumptions as above and skew budgets of 1 FO4 delay locally and 2 globally, a 2-phase system can borrow up to 1 FO4 delay, while a 4-phase system can borrow up to 5 (in both cases setting $t_e=11$).

Two phase skew-tolerant domino stretches traditional domino clocks to considerably improve performance by eliminating latch delays and skew budget from the critical path and by borrowing time to balance pipe stages [3]. Four phase skew-tolerant domino allows much greater skew tolerance and/or time borrowing at the expense of quadrature clock phase generation (with open or closed loop techniques). Other numbers of phases require more difficult clock generation.

To evaluate the performance benefits of skew-tolerant domino, we compared two 64-bit adder self-bypass paths, one using traditional domino with latches and the other using 4-phase skew-tolerant domino, as shown in Figure 5. The paths were simulated in a 0.6 μm 3-metal process assuming a microarchitecture and floorplan similar to the dual integer ALUs of the DEC Alpha 21164 [4]¹. The adder employs two levels of carry selection implemented in dual rail domino logic. With no skew, a traditional design has a latency of 13.0 FO4 delays, but a cycle time of 16.6 due to an 8.3 delay first half-cycle. The skew-tolerant design has a latency of 11.9 FO4 delays because latches were eliminated. Cycle time is also 11.9 because time borrowing is used to balance logic among the phases. Introducing a local skew of 1 FO4 delay does not affect the skew-tolerant design, but increases the traditional latency to 15.0 because skew must be budgeted in both phases.

¹ The 21164 overlaps clocks to eliminate a latch from the ALU.

Figure 1: Traditional domino pipeline

Figure 3: Cross-coupled inverter allows static operation

Figure 2: Skew-tolerant domino pipeline

Acknowledgments

This work was partially funded by a National Science Foundation fellowship.

References

- [1] Heikes, C., "A 4.5 mm² Multiplier Array for a 200MFLOP Pipelined Coprocessor," ISSCC Digest of Technical Papers, pp. 290-291, Feb. 1994.
- [2] Williams, T., "Self-timed rings and their application to division," Ph.D. dissertation, Stanford University, Stanford, CA, May 1991.
- [3] Intel Corporation, *Opportunistic Time-Borrowing Domino Logic*, US Patent #5,517,136, May 14, 1996.
- [4] Bowhill, W., et. al., "A 300 MHz 64b Quad-Issue CMOS RISC Microprocessor," ISSCC Digest of Technical Papers, pp. 182-183, Feb. 1995.

Figure 4: Worst case timing requirements

Figure 5: Adder self-bypass comparison