

A Double-Pulsed Set-Conditional-Reset Flip-Flop

Albert Ma and Krste Asanović
MIT Laboratory for Computer Science
200 Technology Square
Cambridge, MA 02139
{ama,krste}@lcs.mit.edu

Abstract

A new flip-flop design using a double-pulsed static latch is presented. The flip-flop has only a single stage of logic in the critical path and as a result is up to three times faster than the fastest previously known flip-flops, while consuming approximately the same energy as the lowest-power flip-flops. The flip-flop has asymmetric timing properties which make it a good match to skewed static logic styles. A novel dual-pulse generator further reduces power requirements.

Introduction

Flip-flops are critical timing elements in digital circuits and have a large impact on circuit speed and power consumption. Consequently, extensive research has been performed to develop fast and low-power flip-flops [3, 2, 8, 1]. The primary measure of performance of a flip-flop is the minimum D-to-Q delay [8], as this determines how much impact the flip-flop has on cycle time. Recently, pulsed latch structures have emerged as the fastest known flip-flop structures [3, 2]. By reducing the transparency period of a latch to a narrow window, the latch can operate as a flip-flop with the additional advantage of allowing limited time-borrowing across cycle boundaries to reduce sensitivity to clock skew and jitter. These structures have the disadvantage of large positive hold times which complicates timing verification. The pulse generators can also consume considerable energy as pulses must be generated locally to avoid pulse distortion. Nonetheless, because of their performance advantages, these pulsed latch structures have been used in several commercial high-performance microprocessors [4, 5]. Apart from raw performance and energy consumption, other attributes are used to evaluate flip-flop structures including robustness, compatibility with high-performance logic families, and ability to embed logic into the flip-flop.

In this work we introduce a new flip-flop structure, the double-pulsed set-conditional-reset flip-flop (DPSCRFF), which is up to three times faster than the fastest previously known flip-flops while consuming the same power as the lowest-power flip-flops. The DPSCRFF is a single-ended static flip-flop design with a single logic stage which can include arbitrary logic functionality. The DPSCRFF is compatible with static or dynamic logic, and in particular can directly drive following dynamic logic.

DPSCRFF Design

Fig. 1 shows the design of the DPSCRFF. The DPSCRFF is composed of two pieces: a static set-reset latch and pulse-generator.

*This work was partly funded by DARPA PAC/C award F30602-00-2-0562 and by NSF CAREER award CCR-0093354.

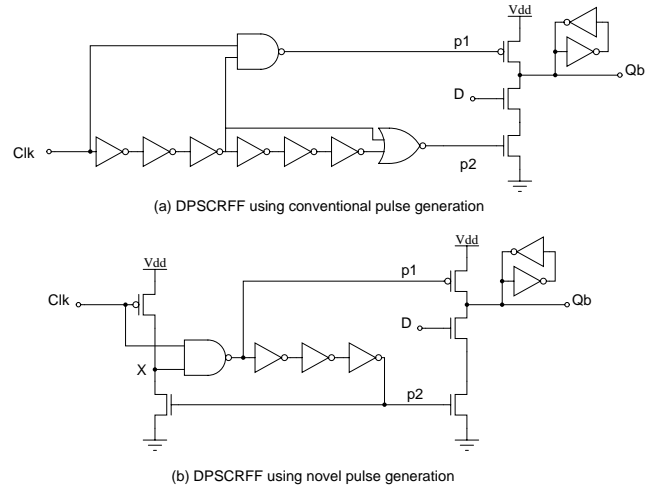


Figure 1: DPSCRFF




Figure 2: DPSCRFF operation

Fig. 2 shows the operation of the static latch. The latch requires two clock pulses, p1 and p2, which are generated from the active clock edge. The first pulse presets the output node high using the p-type pull-up. The second pulse conditionally resets the output node, based on the value of the data input. The precharge causes a glitch at the output node whenever the output is supposed to remain low, which is further discussed below. An additional inverter can be added to the output stage to isolate the storage node from the output load.

The path from input to output is only a single stage of logic which is the key to the design's high-performance. In addition, arbitrary logic can be embedded into the pull-down stack, similar to a domino pull-down tree, as shown in Fig. 3. Another advantage is that the data input sees only a single transistor load which reduces required input drive and energy consumption.




Figure 3: A 2-input mux embedded in the DPSCRFF

Double Pulse Generator

The two pulses are generated by a local pulse generator to avoid pulse distortions from additional pulse buffers and wiring. The pulse generator can be shared by a few neighboring flip-flops to reduce pulse generator area and energy overheads. The width of the pulses is controlled by the inverter delay chain. The inverters in the chain can be skewed to control the lengths of p1 and p2. The width of p2 determines the transparency window of the latch. To reduce setup and hold time requirements, p2 should be made as small as possible. However, if p2 is too short, the circuit will not function. Detailed simulation at all process corners and careful control of clock pulse loading will ensure proper functionality.

The conventional way to generate a pair of pulses uses an inverter delay chain as in Fig. 1(b). This design has a large number of intermediate nodes and thus dissipates a significant amount of energy. Our alternative design reduces the number of intermediate nodes by using an inverter delay chain both to generate p2 and to turn off p1. As shown in Fig. 2, intermediate node X is precharged high during the low phase of the global clock. When the clock rises, p1 falls. After some delay, p2 rises. This causes node X to discharge, causing p1 to rise. After some delay, p2 falls. Note that node X floats in the low state until the global clock goes low. This can be a concern if the global clock is held high for a long time. In this design, p1 and p2 overlap by some amount. This causes some overlap current in the latch when the data input is high. However, the extra energy dissipation caused by the overlap current is not too great, and is much less than the energy savings from using this pulse generator design. It is possible to design the pulse generator to separate the pulses, but the energy cost to separate the pulses with a longer inverter chain is larger than that of the overlap current.

DPSCRFF Timing Analysis

The DPSCRFF has asymmetric timing properties, but this can be exploited in skewed static logic and dynamic domino logic styles. A low input propagates through the flip-flop in negative time as the output is preset at the start of p1. A low input must be setup by the start of the second sampling pulse p2, and the hold time lasts for the duration of p2. A high input, however, can arrive later during the transparency period p2. The hold time of the high input just has to be large enough to switch the state of the static latch. The high value will still be correctly registered at the end of p2 even if the high value drops low again during p2.

Fig. 4 shows two DPSCRFFs connected as a shift register to illustrate hold time violations. Consider the state just before a clock edge, when the first DPSCRFF had a reset value on Qb. This will be propagating through the combinational logic to the input of the second DPSCRFF. At the clock edge, pulse p1 is generated and




Figure 4: DPSCRFF shift register

the first DPSCRFF will begin propagating a preset value from its output before the second DPSCRFF has sampled its input using pulse p2, potentially causing a hold time violation. A conservative approach would be to require sufficient logic levels between DPSCRFFs such that the preset value initiated by pulse p1 could not arrive at the second DPSCRFF until the end of pulse p2. A more aggressive approach takes advantage of the asymmetry of the sampling input. If there are an odd number of inverting logic levels between the two DPSCRFFs, then the high-going preset value from the first DPSCRFF eventually propagates into a low-going value at the input to the second DPSCRFF. This low-going value will not cause a hold-time violation even if it arrives before the end of p2, provided that the previous input was high long enough to flip the latch state. In our technology, we found that five levels of FO4-loaded inverters between DPSCRFFs were sufficient to ensure no hold-time violations across PVT corners with ample margin (three levels just failed in one process corner).

This DPSCRFF does not allow arbitrary time borrowing across the transparency window as with other pulsed latches. Time borrowing is only possible for late arriving high inputs, e.g., from a preceding domino logic stage or a preceding skewed static logic stage.

The output of the DPSCRFF has a glitch in the case where the output Qb is to stay low, i.e., the input remained high. The precharge pulse p1 first forces the output high before the data input resets the output. This glitch can cause additional power dissipation in downstream logic. There is a tradeoff between the additional power dissipation caused by the glitch, and the possible power savings the glitch provides by enabling the use of highly skewed static logic. This is similar to the energy tradeoffs of precharged domino logic versus static logic.

Interfacing to Domino Logic

Fig. 5 shows a DPSCRFF interfacing to domino logic at its input and output. By adding an output inverter, the DPSCRFF can be treated as another domino logic stage. The monotonic rising output of a preceding domino gate can arrive late into the p2 sampling period of the DPSCRFF, reducing effective setup time. The pulsed preset value on the output of a DPSCRFF also simplifies driving a following domino gate. The following domino gate does not have to wait until the worst case Clk-to-Q of the flip-flop to enter evaluate, as the DPSCRFF will first set its output inverter low then give a monotonic rising output in the same way as a domino gate.

Evaluation Methodology

The DPSCRFF, along with other previously published designs (Fig. 6), were simulated using HSpice from schematic netlists annotated with accurate source/drain parasitic diode parameters using a TSMC 0.25 μm process. Fig. 7 shows the testbench used for the evaluation. The testbench is based on that in [7]. However, we

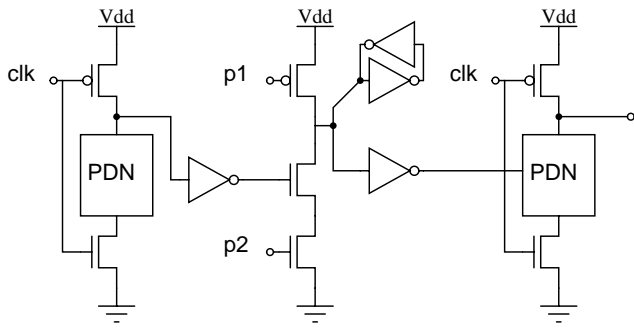


Figure 5: DPSCRFF with domino logic




Figure 7: Testbench setup

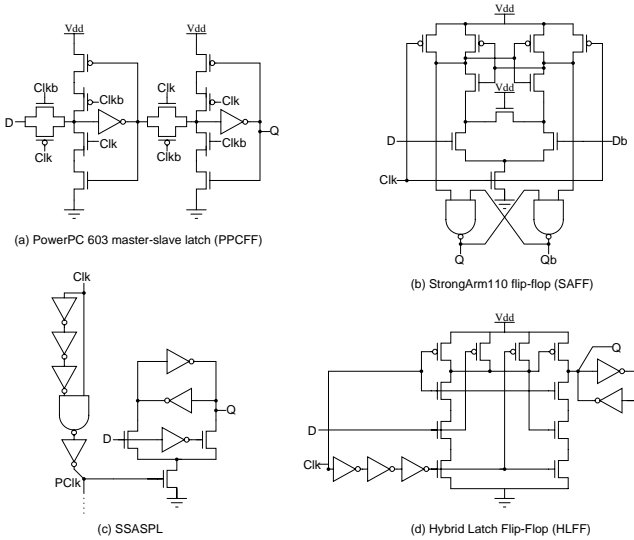


Figure 6: Flip-flops for comparison

those more balanced 2/1 inverters instead of minimum sized inverters in the data and clock buffers. As in [7, 8], we subtracted out the energy dissipated in charging and discharging the output load capacitors. In addition, as in [8], we also subtracted out the energy dissipated in the input buffers.

The relative ranking of flip-flops depends on the loading conditions assumed [6]. For this evaluation, we chose a load of (7.2 fF) which corresponds to four minimum sized inverters in this technology. This represents a typical light load in a datapath structure [7]. To drive higher loads, it is likely that additional levels of output buffering should be used [6].

The pulse generators of the DPSCRFF and the SSASPL were connected to four of the flip-flops, and the energy cost of the pulse generation is considered to be amortized between them.

The transistor sizes in the designs were each optimized for several design points. This optimization was performed using data inputs that were stable well before and after the arrival of the clock. The clock was ungated and the data alternated on every cycle. Clk-to-Q delay and energy were measured. Afterward the minimum D-to-Q delays were found by optimizing the data input arrival times. The minimum D-to-Q delay is the best metric in measuring the performance of timing-elements as it takes into account the relationship between input arrival time and Clk-to-Q delay [8].




Figure 8: Energy versus delay

Results

Fig. 8 show the results. The rising and falling delays for the DPSCRFF have been separated out since they differ significantly. The rising delays are negative since the output precharges before the input is required to arrive. The flip-flops were optimized for the worst-case positive delay, which in some cases increases the negative delays.

As described above, the negative delay can be used to improve performance or to lower power if skewed logic circuits are used. As can be seen, the fastest DPSCRFF at 54 ps is significantly faster than the next fastest flops (HLFF and SSASPL) at roughly 150 ps. The lowest-power DPSCRFF at 141 fJ is comparable to the lowest-power flop (PPCFF) at 130 fJ. However, it has a propagation delay of only 167 ps compared to 342 ps.

Fig. 9 show how the energy dissipation varies with different clock and data input patterns for the different flip-flops. Note that the flip-flops shown in this figure have widely varying propagation delays as shown by the labels in the axes. When the data is held low while the clock continues to run, the energy dissipation of the DPSCRFF is reduced. However, if the clock is running and the data is held high, the DPSCRFF actually dissipates more power than for the full activity waveforms because of its output glitches. When the clock is held stable, no internal nodes change state and only the single data input gate toggles. The DPSCRFF therefore has low energy when the local clock is gated.




Figure 9: Energy Dissipation across different input waveforms

Conclusion

The DPSCRFF has the smallest D-to-Q delay of published flip-flop designs, with comparable energy to the lowest-power flip-flop designs. When the clock is gated, the DPSCRFF has the lowest possible data input loading (a single transistor gate). The asymmetric propagation delay enables the use of highly-skewed static logic to reduce cycle time and energy. The glitching present at the output may cause additional energy dissipation in downstream logic dependent on signal statistics.

References

- [1] B. Nikolić *et al.* Improved sense-amplifier-based flip-flop: Design and measurements. *IEEE Journal of Solid-State Circuits*, 35(6):876–884, June 2000.
- [2] F. Klass *et al.* A new family of semidynamic and dynamic flip-flops with embedded logic for high-performance processors. *IEEE JSSC*, 34(5):712–, May 1999.
- [3] H. Partovi *et al.* Flow-through latch and edge-triggered flip-flop hybrid elements. *Digest ISSCC*, pages 138–139, February 1996.
- [4] M. Golden *et al.* A seventh-generation x86 microprocessor. *IEEE JSSC*, 34(11):1465–1477, November 1999.
- [5] R. Heald *et al.* A third-generation SPARC V9 64-b microprocessor. *IEEE JSSC*, 35(11):1526–1538, November 2000.
- [6] S. Heo and K. Asanović. Load-sensitive flip-flop characterization. In *IEEE Workshop on VLSI*, Orlando, FL, April 2001.
- [7] S. Heo, R. Krashinsky, and K. Asanović. Activity-sensitive flip-flop and latch selection for reduced energy. In *19th Conference on Advanced Research in VLSI*, Salt Lake City, UT, March 2001.
- [8] V. Stojanović and V. Oklobdžija. Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems. *IEEE Journal Solid-State Circuits*, 34(4):536–548, April 1999.