FA 8.5: Flow-Through Latch and Edge-Triggered Flip-flop Hybrid Elements

Hamid Partovi, Robert Burd, Udin Salim, Frederick Weber, Luigi DiGregorio, Donald Draper

NexGen, Inc., Milpitas, CA

This paper describes a hybrid latch-flipflop (HLFF) timing methodology aimed at a substantial reduction in latch latency and clock load. A common principle is employed to derive consistent latching structures for static logic, dynamic domino and self-resetting logic [1].

HLFF is similar to standard flip-flops in that it samples the data on one edge of the clock and thus eliminates a retardation of data flow on the opposite edge. It is similar to latches because it can provide a soft clock edge which allows for slack passing and minimizes the effects of clock skew on cycle time. At an operating frequency of 500MHz, and presenting half the capacitive load to the clock tree, its latency is about two-thirds the aggregate delays of a transparent low (TLL) and a transparent high (THL) latch. As a result, an improvement of at least 10% in cycle time, in addition to a reduction of about 30% in overall clock load, is achieved.

Flip-flops are commonly designed by cascading TLL and THL latches. To avoid an internal race, a delay element must be inserted between the master and the slave elements. HLFF, on the contrary, operates on a different principle. It is a latch with a brief transparency period. The duration of this period is determined by an integrated one-shot derived from the clock edge.

Figure 1 shows a variation of HLFF that can be modified to a storage element for dynamic circuits. Referring to Figure 1, prior to the rising edge of the clock, N1 and N4 are off while N3 and N6 (both gated by CKDB) in addition to P1 are on. As a result, node X is precharged to VDD and node Q (decoupled from X) holds the previous data. At the rising edge of the clock, N1 and N4 turn on while N3 and N6 stay on for a period determined by the inverter delay chain (I1-I3). It is in this period that the circuit is transparent and data at D can be sampled into the latch. Once CKDB transitions low, node X is decoupled from D and is either held at or begins to precharge to VDD by P3. At the falling edge of the clock, P1 fully precharges or holds X at VDD as long as the clock remains low. HLFF waveforms for data transitions of 1→0 and 0→1 are illustrated in Figure 2. The results were obtained in a 2.5V, 0.35μm technology at 2.3V, 85°C with typical devices.

The transparency period of HLFF also determines its hold time. While it is desirable to minimize hold time, the transparency period should be long enough to allow data to propagate to Q. As is seen from Figure 2, CKDB transitions low about 240ps after the rising edge of the clock, hence a hold time of 240ps. Further, both nodes X and Q evaluate 100ps prior to the falling edge of CKDB. This not only provides sufficient safety margin, but indicates a negative setup-time of the same magnitude. The flip-flop latency of 340ps is 100ps longer than its hold time; however, when variations in output load and clock skew were considered, it was required that two flip-flops be separated by at least three logic gates to eliminate race-through.

The negative setup-time of HLFF illustrates an attractive latch attribute known as the soft-clock edge. It allows a critical path to borrow time from the next stage. The soft edge may alternatively be thought of as the ability to overcome a loss of useful cycle time due to clock skew; with a negative setup-time of 100ps, this design tolerates a clock skew of up to 100ps with minimal impact in cycle time.

Level-sensitive latch pairs used in two high-performance processors (Figure 3) are used to assess the performance of HLFF [2, 3]. All were designed and optimized in the technology described earlier. Simulations were performed with equal output loads under two conditions: 1) equal aggregate clock load for latches and HLFF; 2) HLFF having half the clock load. Table 1 summarizes the aggregate latencies of HLFF and the latch pairs. It also includes the percentage gain in clock frequency for a 500MHz processor using latch pair type 2 when replaced by the other latch elements in the table. As can be seen from the table, HLFF designed at half the clock load (HLFF 1/2) can increase the operating frequency by 10% solely due to lower latency. The overall improvement is, however, larger due to three factors: 1) the reduction in clock load invariably reduces clock skew; 2) a percentage of clock skew is absorbed by the soft edge; and, 3) the retardation in data flow, as is the case with latch pairs, is eliminated at the falling edge of the clock.

In developing latch-flipflop libraries, it is desirable to design elements which incorporate logic functions. In addition, conditionally enabled flip-flops are useful as they provide the ability to deactivate functional blocks which are not used. Figure 4 is the circuit diagram for a conditional flip-flop implemented in HLFF.

Dynamic domino circuits are commonly used when it is not possible to meet timing requirements with static logic. A TLL must precede a dynamic block which is evaluated when clock is high. As an example, Figure 5 comprises a dynamic XOR function fed by a type 2 TLL. As is seen, the pulldown network must be conditioned to the clock because one of the latch outputs will be high during precharge.

Figure 6 illustrates the HLFF-based dual-rail storage element for dynamic domino circuits (DHLFF). It replaces the TLL preceding the dynamic blocks. Referring to the figure, both DHLFF outputs are pre-discharged to ground when the clock is low. At the rising edge of the clock and depending on the state of data at D, either QFH or QFL will be asserted. The outputs are held statically as long as the clock remains high. Unlike the TLL implementation, the XOR pulldown network of Figure 5 need not be conditioned to the clock as DHLFF outputs are both discharged when clock is low.

To evaluate DHLFF, its clock-to-output delay was compared with the setup requirement of TLL. The timing waveforms are included in Figure 7. A setup-time of 400ps was required for TLL such that its deselected output had a clock crossover below a threshold. The latency of DHLFF was only 280ps. Similar to its static counterpart, DHLFF latency is about two-thirds that of TLL.

A single-ended example for self-resetting circuits completes the suite of latching elements for different logic styles (Figure 8). The circuit samples the data on the rising edge of the clock and is reset through PR1 after a predetermined delay period which can be made as long as desired while keeping the hold time consistent with the other HLFF structures.

References:

[1] Chappell, T. I., et al., "A 2ns Cycle, 3.8ns Access 512kb CMOS ECL SRAM with a Fully Pipelined Architecture," J. Solid-State Circuits, Nov., 1991.

[2] Dobberpuhl, D. W., et al., "A 200MHz 64b Dual-Issue CMOS Microprocessor," J. Solid-State Circuits, Nov., 1992.

[3] Bowhill, W. J., et al., "A 300MHz 64b Quad-Issue CMOS RISC Microprocessor," ISSCC Digest of Technical Papers, pp. 182-183, Feb., 1995.
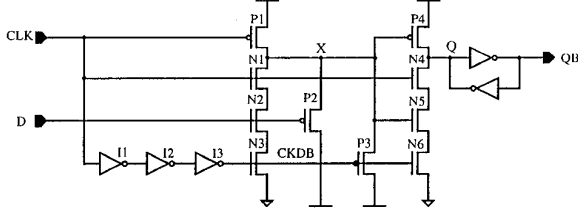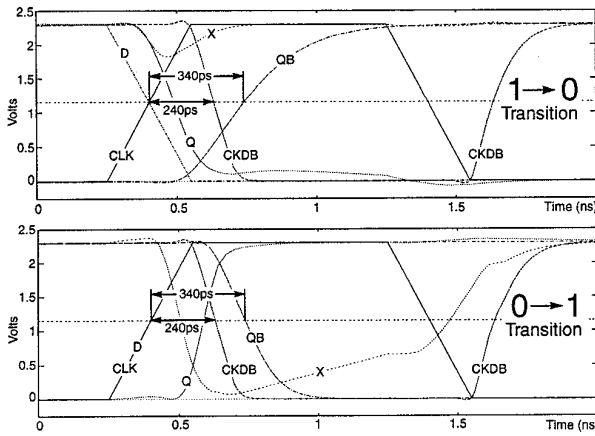

Figure 1: Basic HLFF circuit.

Figure 2: Waveforms for HLFF (Tsu = 0ns).
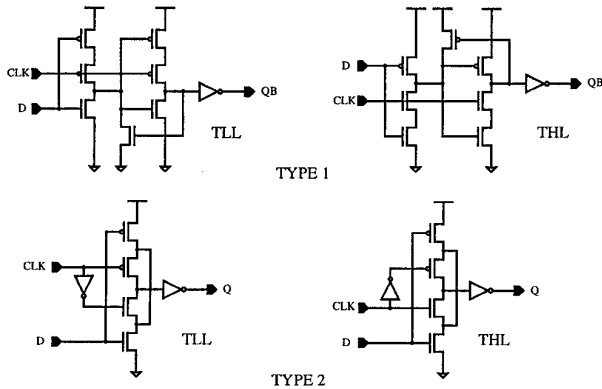

TYPE 1

TYPE 2

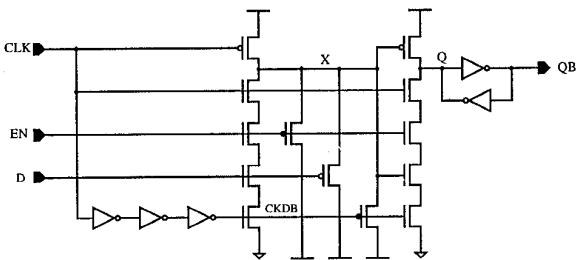Figure 3: Level-sensitive latch pairs.
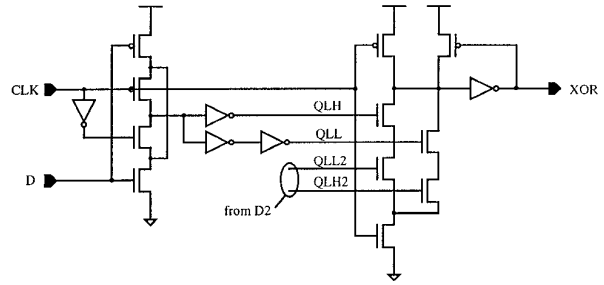

Figure 4: HLFF incorporating ENABLE function.

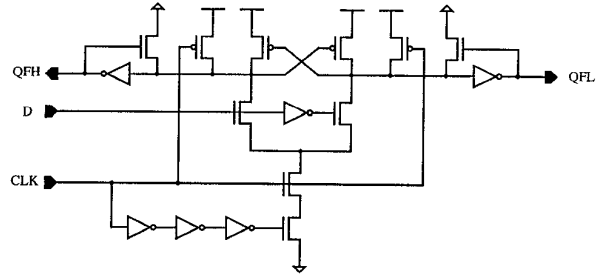Figure 5: Dynamic XOR circuit driven by TLL type 2.
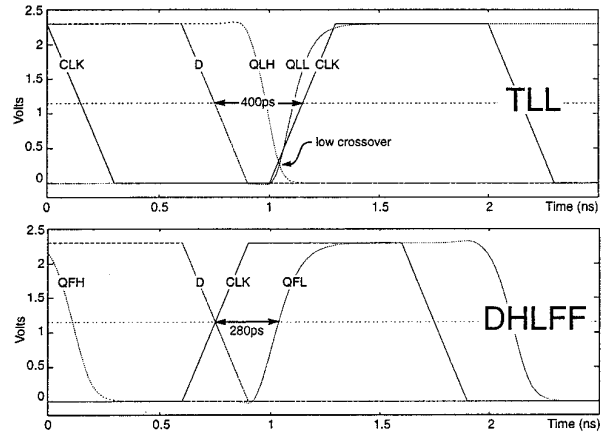

Figure 6: Dual-rail dynamic HLFF.

Figure 7: Waveforms of TLL and DHLFF.

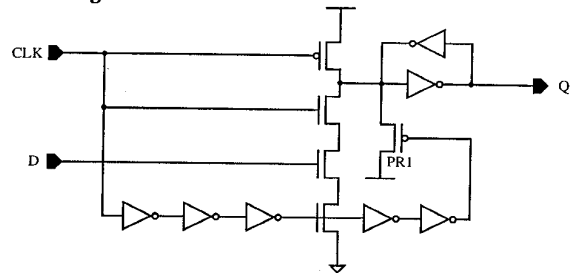Figure 8: Single-ended dynamic HLFF for self-resetting logic.

	HLFF	HLFF 1/2	Latch pair type 2	Latch pair type 1
Latency(ps)	270	340	530	750
Gain(%)	14	10	---	-10

Table 1: Latch latency and gain in operating frequency.