

# IBM Research Report

## A Low-Voltage Swing Latch for Reduced Power Dissipation in High-Frequency Microprocessors

**Pong-Fei Lu, Leon Sigal, Nianzheng Cao, Pieter Woltgens,  
R. Robertazzi, D. Heidel**

IBM Research Division  
Thomas J. Watson Research Center  
P.O. Box 218  
Yorktown Heights, NY 10598



Research Division

Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

# A Low-Voltage Swing Latch for Reduced Power Dissipation in High-Frequency Microprocessors

Pong-Fei Lu, Leon Sigal, Nianzheng Cao, Pieter Woltgens, R. Robertazzi, D. Heidel  
IBM T. J. Watson Research Center  
Yorktown Heights, NY 10598

## Abstract

We report in this paper a new low-swing latch (LSL) for low-power applications. Unlike the conventional transmission gate latch, the LSL allows reduced voltage on the clock inputs. Therefore the local clock buffer (LCB) can use reduced swing to save power while all other circuits are running at nominal voltage. We have implemented an accumulator loop experiment in an early version of IBM's 90 nm SOI technology [1] on a testchip. The experiment consists of an adder and a decremter surrounded by latches to mimic logic between pipeline stages. Side-by-side comparisons between the transmission gate latch and LSL are designed to illustrate the superior power-performance tradeoff of the LSL approach. Hardware measurements have shown 12% maximum AC power saving in 90nm technology.

## 1. Introduction

The power consumption of modern microprocessors design has become a major impediment to frequency scaling [2,3]. The clock power accounts for a big portion of the total power as the number of latches explodes in deep-pipelined designs [4]. Statistics compiled from past IBM microprocessors [5,6] show that the clock distribution, LCBs, and latches account for over 50% of the AC power on a chip. Yet, the data delay in the latch is only 10-20% of the cycle time. Thus, one way to achieve better power-frequency tradeoff is to reduce the supply voltage just on the clock nets, while using a latch topology that minimizes the delay penalty at low swing. However, the conventional latch (Figure 1a) has a clock splitter to drive the input transmission gates of the latch. Reduced voltage on the clocks will result in failure to turn off the P-FET when the clock is high, thus causing a totem pole current. Also, the output of the clock splitter becomes voltage-divided between the P- and N-FET, thus weakening the gate drive of the P-passgate. Therefore, if the clock swing were to be reduced, all circuitry would need to do the same in order to be functional. This mandatory uniform voltage scaling excludes the possibility of independent adjustment of the clocking power.

The proposed LSL topology, shown in Fig. 1b, replaces the transmission gate with a gated inverter, similar to the single-phase latch in [7]. (A different type of low-swing clocked latch design has been reported in [8].) The operation of the latch is as follows. Data are driven into the latch at the rising edge of the clock. In Figure 1b, when 'data' is a '1,' node 'a' is pulled down when C1 clock is switched from low to high, resulting in a pull-up of latch node, '11.' When 'data' is a '0,' node 'a' is a '1' at standby. As C1 clock is switched from low to high, node '11' will be pulled down. When the clock is low, the latch is opaque. Node 'a' may be tri-stated to '1' while node '11' tri-stated to '0.' The weak feedback PFET is used to improve noise immunity. The slave (L2) port operates in a similar fashion. (Note that there is an inherent skew in the delays when data is clock-launched, with data '0' being the faster edge.) Since C1/C2 clock pins only drive the gate of the NFETs, the latch can remain functional at reduced supply, with only minimal impact in delay due to the weakened stack strength. The intent here is to have two separate power supplies: a low-swing supply for LCB to distribute clocks over heavily loaded nets to the latches; and a regular supply for all logic circuitry and latches. The LSL is pin-compatible with the transmission-gate latch, and both are LSSD [9] scan testable with A/B clocks.

## 2. Experiment Design

We implemented an accumulator experiment in an early version of IBM's 90nm SOI technology [1]. The experiment (Fig. 2) consists of 4 major components: (1) a clock generator, (2) a block of glue logic for controlling the operation of the experiment, (3) the actual LSL experiment, and (4) a control experiment using the conventional latch. The clock generator can be driven by either an internal voltage controlled oscillator (VCO) or an external clock. The experiment is an accumulator loop: data is launched from the latch, going through two (identical) 65-bit adders in series, then wraps back to the latch. The second adder is configured as a decremter (i.e. the 'b' inputs are tied to ' $V_{dd}$ '). Its purpose is to lengthen the delay, so that the minimum cycle time fits in the frequency range (2-3 GHz) of interest. The latch data at the  $n^{\text{th}}$  cycle,  $A(n)$ , is:

$$A(n)=A(n-1)+B-1$$

where 'B' is a fixed number programmable through scan-only latches (see Fig. 2). Thus the latch content changes by fixed increment every cycle. The experiment will run through a fixed number of cycles programmable by the scan latches in the glue logic (not shown in Fig. 2), and at the end the latch data will be scanned out to compare with the expects from cycle simulations. The minimum cycle time should be the sum of latch latency and the (maximum) adder delays. The compare will eventually fail at high frequencies when the latch can no longer keep up with the change.

The 'glue logic' starts the experiment at the falling edge of CLK<sub>G</sub> by sending out a 'start' pulse. The experiment then runs for a programmable number of 'warm-up' cycles to let the power supply stabilize. After these warm-up cycles, a 'reset' signal is issued to the low-swing experiment and the control experiment to initialize all 'A' latches to zero. In the meantime, the experiments continue to run for a programmable number of clock cycles, after which a 'stop' signal will be issued by the glue logic to stop the experiment. The experiments then return to the scan mode at standby, and the data are scanned out of the latches.

The only difference between the LSL experiment and the control experiment is the type of latch used. Functionally, however, both experiments are the same. Each experiment has its own LCB and LCB controller, with features to do various clock gating and clock stressing for testability. One LCB drives 32 latches; thus there are two LCBs in each experiment. To facilitate direct measurement of different power components, multiple  $V_{dd}$  domains are used.

## 3. Hardware Results

The measurements were performed by first determining the maximum frequency,  $F_{\text{max}}$ , at a fixed supply voltage. The accumulator functionality was verified at low frequencies using a complex add pattern over 50,000 cycles. The clock frequency was then gradually increased until the experiment failed. For power measurements, the clock frequency was fixed at  $F_{\text{max}}$  and the experiment was run for an extended time so that average currents could be measured. The same steps were repeated as the supply voltages were varied. The testing for the control and LSL experiments was done separately. The clock gating capability of the LCB/LCB controller was used to turn off the experiment not of interest for better noise isolation.

The system AC power consists of power dissipation from three components: LCBs, latches and adders. The three are independently measured via different power supply domains. For the control latch experiment, the three voltages are varied altogether. Fig. 3 shows the power consumption vs.  $F_{max}$  at different  $V_{dd}$ . It is seen that the LCB power and latch power account for roughly 45% of the total system power, and the adder power accounts for the remaining 55%; this result is close to what has been observed in high-frequency microprocessors. The latter indicates that our experiment is representative of a pipeline stage in real microprocessor designs. For the LSL experiment, the latch and adder  $V_{dd}$  supplies are fixed as in the control experiment, and the LCB supply is swept down from  $V_{dd}$  until the latch stops functioning. Figure 4 shows an example of  $V_{dd}$  equal to 1.0 V. As expected, the LSL latch slows down as the  $V_{LCB}$  is reduced. However,  $F_{max}$  is not significantly impacted until  $V_{LCB}$  reaches 0.65 V. Thus, there is potential to save power by reducing the local clock voltage swing without substantial impact on the frequency. Figure 5 shows the LCB power consumption at different supply voltages and corresponding  $F_{max}$  are presented. For the control experiment, the LCB power is measured with all components at the same voltage. For the LSL, each curve represents the LCB power measured at given system voltage  $V_{dd}$  (all components except LCB) as  $V_{LCB}$  is varied. The gradual drop of  $F_{max}$  with decreasing  $V_{LCB}$  in Figure 3 is reproduced for all  $V_{dd}$  values measured, demonstrating the potential for consistent power savings over a wide voltage range.

As discussed in Section 1, with the conventional transmission gate latch all  $V_{dd}$  must be adjusted all together to be functional; thus the clocking power cannot be reduced without significant degradation in  $F_{max}$ . In contrast, with the LSL a lower  $V_{LCB}$  will only impact the latch delay. The flatness of LCB power vs.  $F_{max}$  in Figures 4 and 5 suggests that the LSL approach can provide better system power-frequency tradeoff. Figure 6 shows the total system AC power consumption of the LSL and control experiments. The AC power is the sum of powers of LCBs, latches and adders at the operating frequency. For the control experiment (black dots), the system power is plotted as a function of the system voltage supply. For the LSL experiment, each curve represents the system power measured at given system voltage while  $V_{LCB}$  is varied. For nearly the whole measured range of system voltage, the LSL power curves are

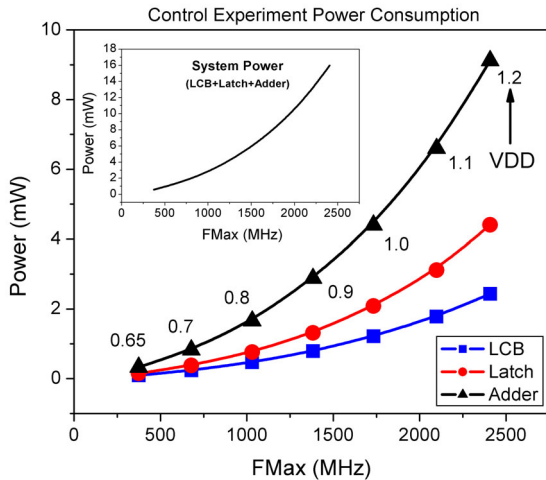


Fig. 3: AC Power vs.  $F_{max}$  for control experiment

below the control experiment power curve, confirming its superior power-frequency tradeoff. For example, at the nominal  $V_{dd}$  of 1.1 V for the 90nm technology, a maximum power saving of 12% can be obtained at  $\sim 2$  GHz using the LSL with  $V_{LCB}$  at 1.0 V and other voltages at 1.1 V. The latter can also be interpreted as the following. Part of the powers saved from running the LSL with lower  $V_{LCB}$  can be re-allocated to latches and adders for speed-up. The latter overcompensates the slowdown at the LSL input NFET device, resulting in higher  $F_{max}$  for the same power, or lower power for the same  $F_{max}$ . Note that in Figure 6, the relative power saving is observed to be reduced at lower  $V_{dd}$  and  $F_{max}$ . The latter can be attributed to the fast increase in the channel resistance of the clock input NFETs in the LSL as the supply voltage approaches the threshold voltage.

#### 4. Conclusion

In modern microprocessors design, there is an imbalance of large clock power (50%) and relatively small latch delay (10-20%) component in a pipeline. Thus there is a potential opportunity for better power vs. performance tradeoff using a low-voltage swing latch. In this paper, we have proposed a new latch circuit, and demonstrated through hardware results that significant power saving can be gained.

**Acknowledgements:** We would like to thank SRDC Technology Group in Fishkill, NY for processing the testchip.

#### References

1. M. Khare et al, IEDM Technical Digest, p. 407, 2002.
2. Patrick P. Gelsinger, Technical Digest of ISSCC 2001, p. 22.
3. M. J. Flynn et al., IEEE Micro, 19(4): 11-12, July/Aug 1999.
4. V. Srinivasan et al., 35<sup>th</sup> IEEE/ACM Int. Symp. On Microarchitecture 2002.
5. C. Anderson et al., Technical Digest of ISSCC 2001, p. 232.
6. B. Curran, Technical Digest of ISSCC 2001, p. 238.
7. D. Dobberpuhl et al, JSSC, pp.1555-1567, 1992.
8. R. K. Krishnamurthy et al, 2002 VLSI Circuit Symp. Digest, p. 128.
9. E. B. Eichelberger, U.S. Patent 3,761,695, 1973.

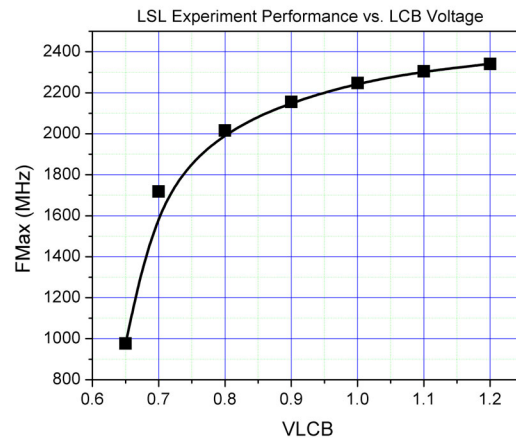


Fig. 4: Frequency vs. LCB voltage with  $V_{dd}=1.0$  V for the LSL experiment

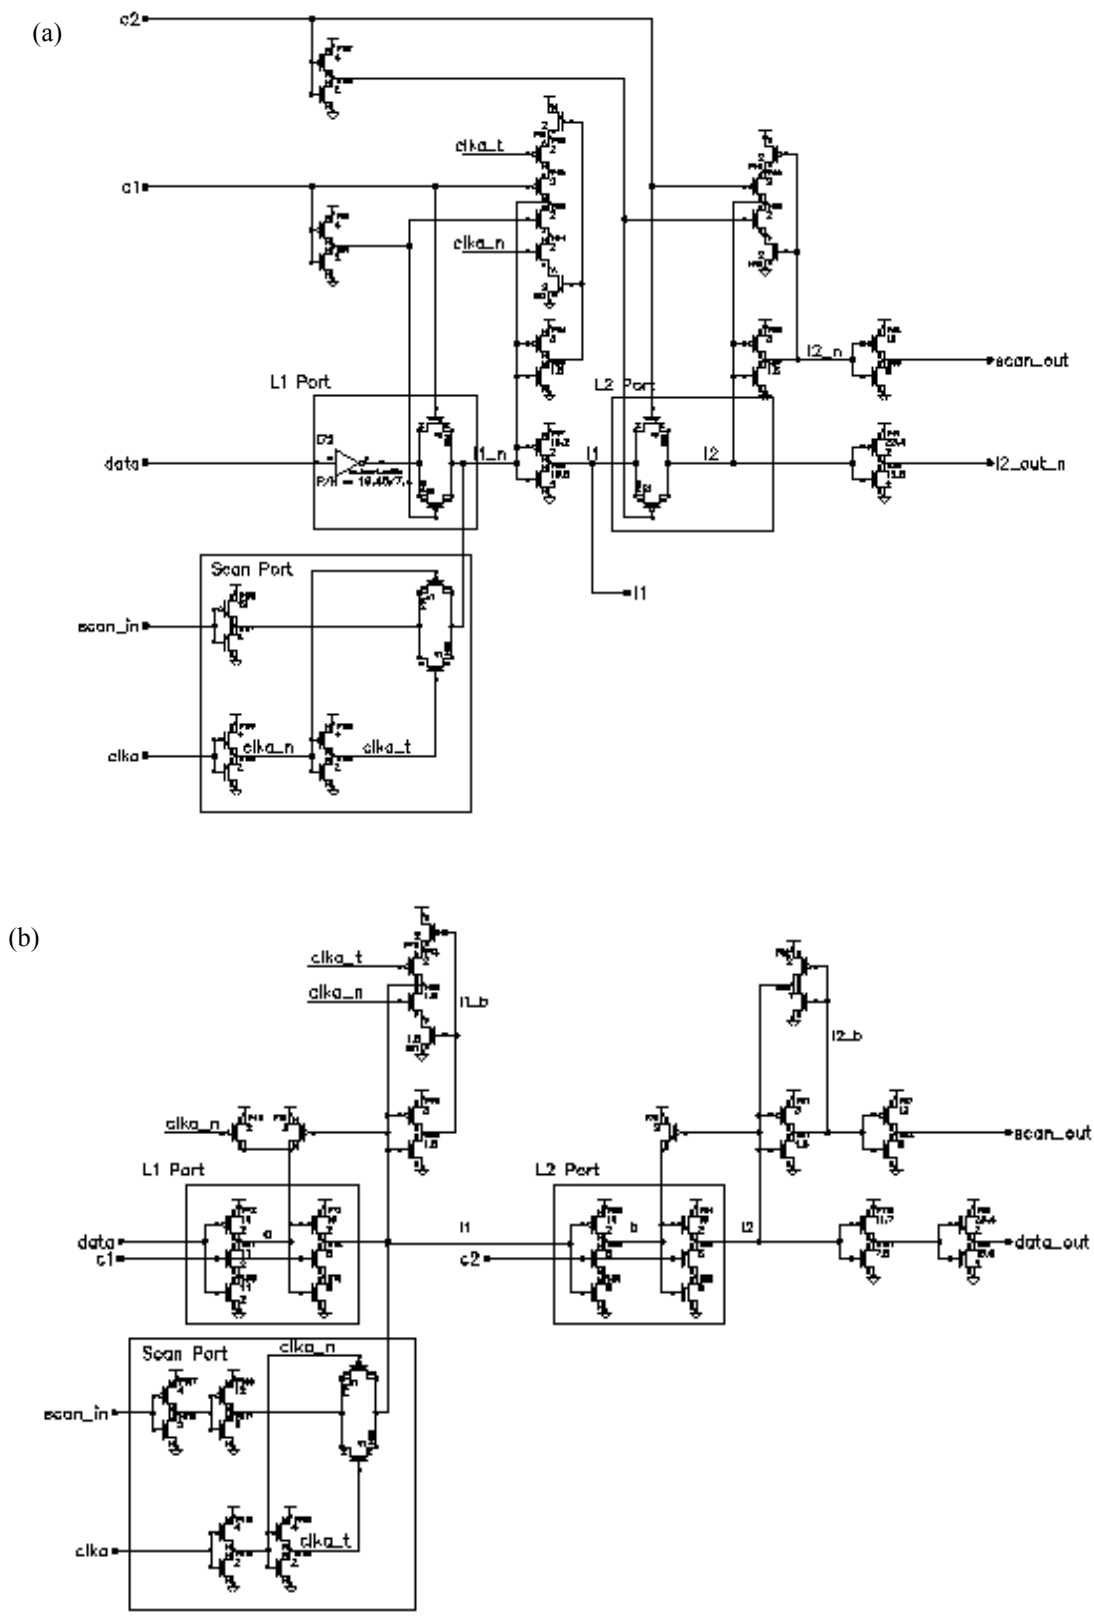
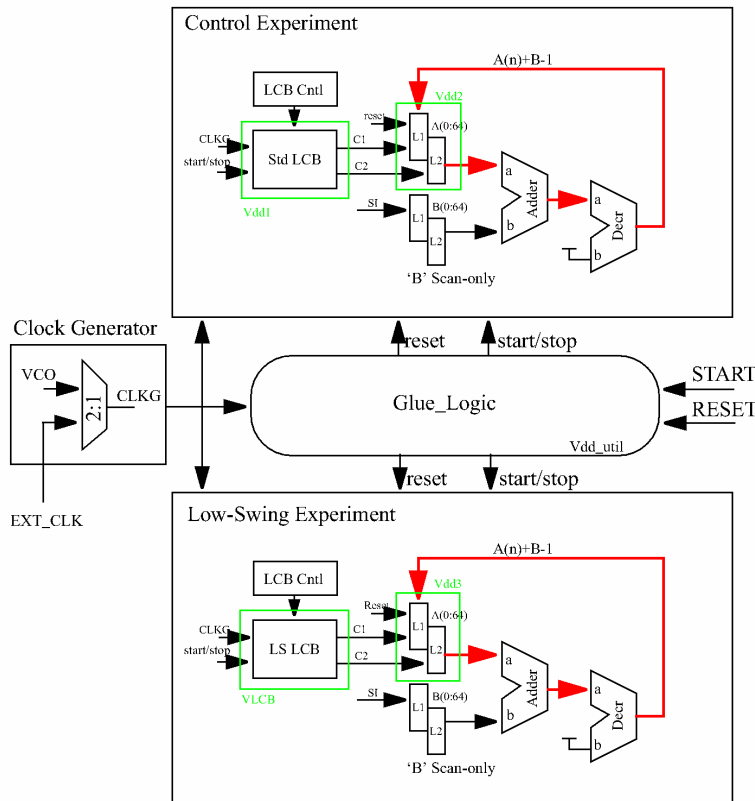



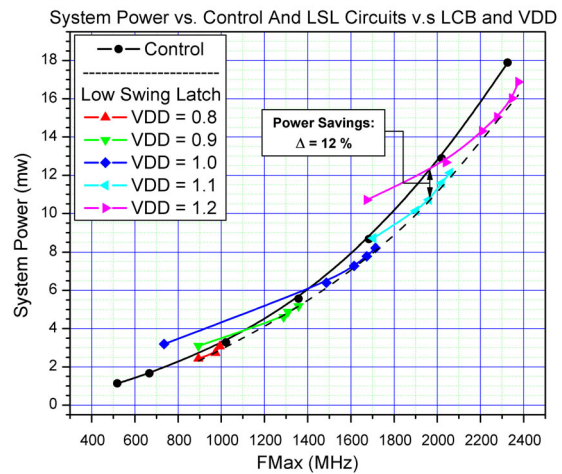
Fig 1: (a) Conventional transmission gate master-slave latch, and (b) the proposed low-voltage swing master-slave latch



**Fig. 2:** Experimental setup: A ‘glue logic’ block controls the side-by-side experiments. It can start/stop the experiments as well as reset the latches. The global clock can either be generated by an internal VCO or be fed from external clock. The LCB/LCB controller have capabilities of clock gating to gate off the clocks, as well as clock stressing to enhance testability. There are two adders in series in the accumulator loop. The second adder is configured as a decremter by tying the ‘B’ inputs to a ‘1.’ The loop performs a function of  $A(n)=A(n-1)+B-1$ , where constant B is programmable through scan-only latches. Five  $V_{dd}$  domains are used to facilitate power measurements: four for the latches and LCBs of each experiment, and one for the adders and miscellaneous circuitry (clocks and ‘Glue\_logic’).



**Fig. 5:** LCB power vs. frequency. For the control experiment, all voltages are varied together. For LSL, the  $V_{dd}$  is fixed as in the control experiment while the LCB voltage is swept down from  $V_{dd}$  to 0.6V.



**Fig. 6:** System AC power vs.  $F_{max}$ . For the control experiment, all voltages are varied together. For LSL, the  $V_{dd}$  is fixed as in the control experiment while the LCB voltage is swept down from nominal  $V_{dd}$  to 0.6 V.