

Clocking Strategies and Scannable Latches for Low Power Applications

V. Zyuban and D. Meltzer

IBM Research Division, T.J. Watson Research Center, Yorktown Heights, NY 10598
(D. Meltzer is now with Epson Research and Development, NY)

Abstract

This paper covers a range of issues in the design of clocking schemes for low-power applications. First we revisit, extend and improve the power-performance optimization methodology for latches, attempting to make it more formal and comprehensive. Data switching factor and the glitching activity are taken into consideration, using a formal analytical approach, then a notion of energy-efficient family of configurations is introduced to make the comparison of different latch styles in the power-performance space more fair, also the power of the clock distribution is taken into account. Practical issues of building a low overhead scan mechanism are considered, and the power overhead of the scannable design is analyzed. A low-power LSSD extension to single-phase latches is proposed, and results of a comparative study of LSSD-scannable latches are shown, supported by experimental data measured on a 0.18μ test chip.

Introduction

Since the importance of designing low-power high performance clocking schemes has been recognized, a number of low-power latch studies have been published [9, 12, 8, 11, 6, 2, 10, 5, 4]. Various latch styles have been compared in the power-performance design space, a number of useful criteria have been introduced, and several new low-power latches have been suggested.

This paper improves the existing power-performance optimization methodology in several aspects. First the methodology is formalized by using analytical formulas to take into account *both data switching activity* and the *glitching factor*, based on [12]. A formal optimization of every latch style in the power-performance space is performed before comparing different latch styles through constructing *energy-efficient families* of configurations for every latch. The importance of the latch scalability with respect to lowering supply voltage is emphasized by treating V_{dd} as a parameter, rather than a constant. Power of the clock distribution network is analyzed, and the single-phase clocking scheme is compared with two-phase clocking.

A number of practical issues of building latches have been missing in many academic studies. One of them is the testability issue, particularly, the scan mechanism. At the same time the power overhead of the scannable design can be very significant, and the complexity of modern designs has reached the point where saving power by implementing a non-scannable design is not viable. In this work we propose a low power overhead scan mechanism for single phase latches, and compare it with other approaches in terms of power.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED'01, August 6-7, 2001, Huntington Beach, California, USA.

Copyright 2001 ACM 1-58113-371-5/01/0008 ...\$5.00

Experimental verification of the ideas is becoming increasingly important as the technology moves to the deep sub-micrometer region, because it is getting more and more difficult to take into account in simulations all second order effects of the technology. In this paper we present experimental results for low-power latches built on a state-of-the art technology.

1 Optimization and comparison methodology

1.1 Performance Measurement

The state of the art methodology for comparing the performance of different latches consists in evaluating the following metric [8], based on simulation of the switching of the latch for varying values of the data setup time:

$$T_{setup} + D_{C \rightarrow Q} = \min[T_{D-to-C} + \max(D_{0 \rightarrow 1}, D_{1 \rightarrow 0})],$$

where T_{setup} is the setup time, and $D_{C \rightarrow Q}$ is a delay through the latch, measured from the appropriate transition at the clock input and the corresponding transition at the latch output. In the formula the max chooses the maximum delay between the positive and negative transitions, and the min chooses the smallest value of the sum for all values of the delay between the transitions at the data and clock inputs, T_{D-to-C} , as shown in Fig. 1. For this latch the minimum value for the sum is reached when the delay between the data and clock transitions is $T_{D-to-C} = 130ps$, and the delay through the latch at this point is $\max(D_{0 \rightarrow 1}, D_{1 \rightarrow 0}) = 280ps$. Thus, we put for this latch $T_{setup} + D_{C \rightarrow Q} = 130ps + 280ps = 410ps$. All delays are measured assuming a load of four minimum size inverters.




Figure 1: Evaluation of the performance metric.

1.2 Power Measurement

A significant obstacle in calculating power directly by simulation is that power dissipation is strongly pattern dependent. For example, in a latch the power depends on the average number of transitions at the data input, as well as their time positions. The state of the art

methodology used by other authors typically estimates the power of a latch for two values of the switching activity at the data input: $\alpha = 0$ and $\alpha = 1$, and then estimates the average power as a linear combination of the power under these extreme cases, with the weights depending on the data switching factor α . The *spurious* switching activity, or *glitching* at the data input is typically either neglected or added in an ad hoc manner.

In our study we used a more formal approach developed in [12] which models the circuit as a directed graph, called the *state transition diagram*, STD [3, 7], such that there is a one-to-one correspondence between edges in the graph and power-dissipating events in a circuit.

Table 1: Reachable states in the STD of the modified SA latch. '1' and '0' designate voltage levels at the nodes of the circuit.

node name	state											
	1a	1b	2a	2b	3a	3b	4a	4b	5a	6a	7a	8a
C	1	1	1	1	1	1	1	1	0	0	0	0
D	1	1	1	1	0	0	0	0	1	0	0	1
Q	0	0	1	1	0	0	1	1	0	0	1	1
S	1	1	0	0	1	1	0	0	1	1	1	1
R	0	0	1	1	0	0	1	1	1	1	1	1
A,B,M	0	0	0	0	0	0	0	0	1	1	1	1
G	0	0	1	0	0	0	1	0	0	0	0	0
H	0	1	0	0	0	1	0	0	1	1	0	0

To construct a STD for a latch we build a *state tree* using the state tree algorithm [12]. In the first row we draw one circle for each combination of voltage values at the input nodes. In the second row we specify all nodes whose voltages are uniquely specified given the voltage values at the nodes in the first row. Then the tree branches, each branch corresponding to the state of some node whose voltage value is so far independent of those at all nodes in this branch above it. In the next row, again, we write all nodes whose states are uniquely specified given the voltages at the nodes in the branch above them. This repeats until all nodes are specified. As an example, Table 1 shows all reachable states for the modified sense amplifier latch, shown in Fig. 3a. The modification of the latch consists in interchanging input to the NANDs in the second stage, so that the 'S' signal is connected to the lower input and the 'R' signal – to the upper input of the corresponding NAND gates. The result of this modification is that the capacitance charged/discharged on the slower $Q : 1 \rightarrow 0$ transition is minimized at the expense of somewhat higher capacitance charged/discharged on the other transition. Though a simple modification, it results in 5% delay reduction, compared to the configuration where both 'R' and 'S' signals are connected to the lower inputs of the NAND gates, and in 3% delay reduction compared to the configuration where both 'R' and 'S' signals are connected to the upper inputs of the NAND gates. The ability to do this sort of analysis easily is another reason for using the STD analysis.

After all states have been specified, we build a STD by starting with one state that is obviously reachable. For every state in the graph we find two other states reachable from it via edges corresponding to transitions at the clock and data inputs. This process is repeated until both edges leaving every node enter states that have already been counted. The STD for the modified SA latch is shown in Fig. 2a, with the states designated in Table 1.

To simplify the explanation in this section, we will merge the states in Fig. 2a that differ only by the voltage levels at nodes *G* and *H*. Then states with an index *a* in Fig. 2 merge with those that have index *b*, reducing the STD to a simpler one, shown in Fig. 2b. For our layout of the latch the capacitances at nodes *H* and *G* are about 8 times smaller than those at nodes *S* and *R*, making the simplification quite reasonable. Still, the methodology can be carried out




Figure 2: State transition diagram of the modified sense amplifier latch (a); simplified STD (b).

without any simplifying assumptions, and all results presented in this paper are based on the full analysis of the STD for every latch.

It is important to emphasize that we attribute the power dissipated for charging/discharging capacitances at the clock and data inputs to the latch itself, rather than to a fan-in gate or the clock distribution tree. It is important that the input capacitance of the wiring within the latch layout be included. Similarly, we do not include the power dissipated for charging/discharging the load driven by a latch into latch power. However, the output capacitance of the wiring within the latch layout *must* be included. Such a convention makes the power comparison between different latches more fair, however this also makes the STD more complicated compared to those published in other works.

Energy weights for every edge in the STD are calculated using an analog simulator or, for rough estimates, manually, using the formula $\frac{1}{2} \sum C_i V_{dd} \Delta V$, where $\sum C_i$ is the sum of capacitances at all nodes that have different voltage levels in the states connected by the edge. Then, based on the probabilistic analysis of the STD, presented in [12], analytical formulas are derived for the power of a latch that express latch power in terms of true and spurious switching activities at the data input.

$$P_{\text{true}} = f(Q_0 + P_1 Q_1 + \alpha Q_2) \quad (1)$$

where

$$Q_0 = E_{00}, \quad Q_1 = E_{11} - E_{00},$$

$$Q_2 = \frac{1}{2} (E_{01} + E_{10} - E_{00} - E_{11})$$

In these formulas α is the switching activity as defined in [12], P_1 is the probability of latching '1', E_{mn} are energy weights of paths p_{mn} , $m, n = \{0, 1\}$ in the graph that are traversed when m was latched in the previous clock cycle and the data has changed to n (or has not changed if $m = n$). For the simplified STD in Fig. 2b some of the paths are $p_{01} = \{s_3 \rightarrow s_6 \rightarrow s_5 \rightarrow s_2\}$, $p_{10} = \{s_2 \rightarrow s_8 \rightarrow s_7 \rightarrow s_3\}$, $p_{00} = \{s_3 \rightarrow s_6 \rightarrow s_3\}$, and $p_{11} = \{s_2 \rightarrow s_8 \rightarrow s_2\}$. The path energy weight is obtained by summing energy weights of all edges along it.

For accurate estimates, we simulate the latch, using an input pattern that causes the latch to go through every single edge in the STD, and energy weights are measured by the simulator. It turns out that every term in power formulas includes only energies dissipated on complete cycles in the STD, which allows us to measure these energies by integrating the current of the power supply, in case the simulator does not support the measurements of the instantaneous power.

In the presence of the spurious activity at the data input a formula similar to (1) is derived in [12] that is valid for all reasonable




Figure 3: Transistor diagrams: **a** – modified sense amplifier latch, **b** – true single phase RAM latch.

latches known to the authors:

$$P_{\text{total}} = \beta_0 P_{\text{true}} + f \left[\beta E_{\text{cycle}} + \beta^* E_{\text{cycle}}^* + (1 - \beta_0)(Q'_0 + P_1 Q'_1 + \alpha Q'_2) \right] \quad (2)$$

where $Q'_0 = \sum_{p_{00}^1} E_{ij} - E_{\text{cycle}}$, $Q'_1 = \sum_{p_{11}^1} E_{ij} - \sum_{p_{00}^1} E_{ij}$, $Q'_2 = \frac{1}{2} \left(\sum_{p_{01}^1} E_{ij} + \sum_{p_{10}^1} E_{ij} - \sum_{p_{00}^1} E_{ij} - \sum_{p_{11}^1} E_{ij} \right)$. Here P_{true} is the average power in the absence of glitches, calculated by (1), β the average number of spurious pulses during one clock cycle calculated as $\beta = \sum \beta_k k$, where β_k is the probability that k spurious pulses occur during one clock cycle. In many latches spurious pulses occurring when clock is high dissipate more (or less) energy than those occurring when clock is low. This is accounted by the term $\beta^* E_{\text{cycle}}^*$, where β^* is the average number of spurious pulses per cycle occurring while clock is high. In the above formulas p_{ij}^k denote a path in STD traversed when i was latched in the previous clock cycle and the true data value has changed to j (or has not changed if $i = j$) and when k spurious pulses occurred in this clock cycle. In the above formula the summation of the energy weights is taken along such paths. For example, in Fig. 2b, $p_{00}^1 = \{s_3 \rightarrow s_6 \rightarrow s_5 \rightarrow s_6 \rightarrow s_3\}$, and $p_{01}^1 = \{s_3 \rightarrow s_6 \rightarrow s_5 \rightarrow s_6 \rightarrow s_5 \rightarrow s_2\}$. E_{cycle} is the energy dissipated by one spurious pulse, provided that at least one spurious pulse has occurred before it. For the STD in Fig. 2b, E_{cycle} is the energy dissipated on the cycle $p_{\text{cycle}} = \{s_6 \rightarrow s_5 \rightarrow s_6\}$, or $p_{\text{cycle}} = \{s_8 \rightarrow s_7 \rightarrow s_8\}$, and E_{cycle}^* is the energy dissipated on the cycle $p_{\text{cycle}}^* = \{s_3 \rightarrow s_1 \rightarrow s_3\}$, or $p_{\text{cycle}}^* = \{s_2 \rightarrow s_4 \rightarrow s_2\}$. There is one more cycle in the STD in Fig. 2a, $p_{\text{cycle}} = \{s_{3b} \rightarrow s_{1b} \rightarrow s_{3b}\}$, or $p_{\text{cycle}}^* = \{s_{2b} \rightarrow s_{4b} \rightarrow s_{2b}\}$, but it has the same energy.

Thus, the average total latch power in (2) is a sum of the 'true' portion multiplied by the probability that no spurious pulses occur during one clock cycle, β_0 , and the 'spurious' portion which depends on three parameters: β and β^* – the average number of spurious pulses per cycle, when clock is low and high, respectively, and β_0 . The term $f(1 - \beta_0)(Q'_0 + P_1 Q'_1 + \alpha Q'_2)$ accounts for the difference between the energies dissipated by the first and subsequent spurious pulses. For many latches (which is the case for the STD in Fig. 2b), it cancels with $\beta_0 P_{\text{true}}$, and the expression reduces to $P'_{\text{total}} = P_{\text{true}} + f \left[\beta E_{\text{cycle}} + \beta^* E_{\text{cycle}}^* \right]$.

1.3 Tuning Transistor Sizes

Before comparing it with other latches, every latch must be optimized for a power-performance metric, in order to make sure that the best configurations of every latch are compared. However, it turns out that it is virtually impossible to come up with a single metric that would be fair for every latch – some latches are more suitable for high-speed designs, others – for low-power, but slower designs.

To avoid this uncertainty, we built for every latch an *energy-efficient family* of configurations, which is a family of configurations, obtained by optimizing a latch to minimize the cost function $\gamma(E/E_0)^2 + (1 - \gamma)(D/D_0)^2$ for all values of the optimization parameter γ in the range $0 \leq \gamma \leq 1$. Here, D is the sum of the setup time and the delay through the latch, as defined in section 1.1, and E is the average energy dissipated by the latch in one clock cycle, determined according to section 1.2. This cost function was used because it resulted in a consistent convergence of the circuit tuner.

Thus, every configuration in the *energy-efficient family* is the one that results in the highest performance among all configurations dissipating the same power, or the one that dissipated the least power among all configurations that deliver the same performance. If plotted in the power-versus-performance coordinates, *energy-efficient* configurations form a *convex hull* of all possible configurations of a given latch, Fig. 4.




Figure 4: Building energy-efficient family for a latch.

Having built *energy-efficient families* for every latch allows us to compare different latches over the whole energy-performance design tradeoff space, rather than comparing particular configurations of every latch.

2 Clock Distribution Power

When comparing different latches for energy efficiency it is essential that the power dissipated in the clocking tree be taken into account, because different latches present different requirements as well as different amounts of capacitive load on the clock distribution network. Moreover, some latch styles require two clock phases, while others use only one phase.

To evaluate the effect of the power of the clock distribution tree on the latch energy efficiency, we simulated a clock distribution tree for a 32-bit datapath latch, with a 12 track bit step, using a 0.18μ technology with Vdd set to 1V. When calculating the capacitive load presented by a latch to the clock distribution network, it is important to include the capacitance of the clock wiring inside the latch cell. We found that for latch design using very small transistor sizes, the internal wiring may represent from 5% to 20% of the total capacitive load on the clock.

The simulated circuitry included a clock splitter that generates two non-overlapping clock phases and a distribution network feeding every latch. Transistor sizes in the clock drivers were set to the minimal sizes that are needed to guarantee that the slope at every node in the clock distribution network is no more than 100ps.

The simulation results showed that the whole simulated clock distribution network dissipates 240fJ per clock cycle, of which 100fJ is dissipated for driving the wire capacitance and latch input capacitance, and 140fJ is dissipated in clock drivers. Divided by 32 latches in the simulated structure, this yields 7.5fJ per latch per clock cycle. Note that the power dissipated for driving the capacitance at the clock input of the latch is counted as power dissipated in the latch, rather than in the distribution tree. Taking this into account yields an estimate of 6.5fJ per latch per clock cycle overhead for distributing one clock phase. The overhead for distributing two clock phases is 13fJ per latch per clock cycle, which is comparable to the energy dissipated within the latch itself, if the latter is built using very small transistor sizes.

This analysis indicates that in a low power design which uses very small transistors, a latch that can work with a single phase of the clock has a 30% power advantage over a latch that requires two phases for robust operation.

3 Scannable latches

The integration and complexity of modern systems has grown to the point where saving power by building a non-scannable design is no longer an option. The power overhead of the scannable design may be very significant. For example, the study in [1] has reported a 54% increase in power of an LSSD standard cell design over the identical non-scannable design. However, the effect of scannable design on power has not received sufficient attention in recent works on low-power latches. In this work we try to fill this gap by analyzing the power overhead of existing approaches to building scannable latches and proposing a new, low power overhead LSSD compatible extension to edge-triggered latches.

There exist two major approaches to building scannable designs: edge-triggered and level sensitive, LSSD scan. Because the LSSD scan is race free, it is more robust than the edge-triggered scan, and it preserves the integrity of the scan chain even in the presence of significant clock skews [1]. For this reason LSSD is the scan mechanism of our choice.

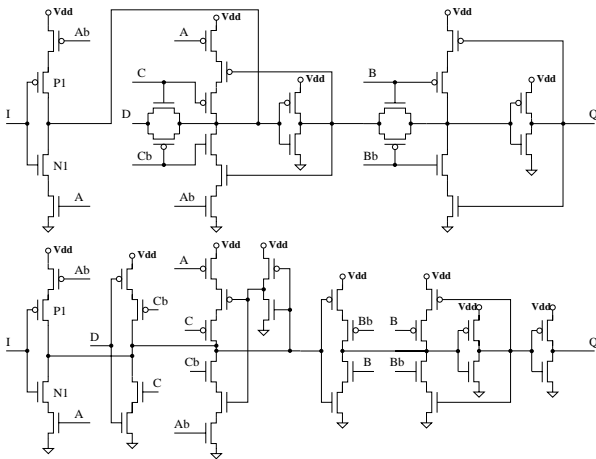


Figure 5: LSSD transmission gate latch (above) and NORA latch (below).

The standard way of implementing an LSSD master-slave latch is shown in Fig. 5 for the transmission-gate latch (called PowerPC latch in [8]), and NORA latch (called C²MOS latch in [8]). In these latches the power overhead of the scan is quite small – only drain capacitances of transistors **N1** and **P1** (which are cut off) are charged/discharged during the normal operation mode. However, these latches require two phases of clock, **C** and **B**, in the normal operation mode, which, according to Section 2, increases the total power of the clocking system by 30%.

In order to avoid the power penalty of the second clock phase, the latch should operate with a single clock phase during the normal mode, and during the scan mode it should operate as a master-slave latch with two non-overlapping clock phases, as required by the LSSD standard. Fig. 6a shows the proposed LSSD extension to the sense amplifier latch that has this property.




Figure 6: Scannable sense amplifier latch: **a** – proposed LSSD scannable latch, **b** – prior art [6] edge-triggered scannable SA latch.

The result is achieved by mixing in the scan-in data at the second stage of the latch, R-S stage. The scan-in data signal, **I** is written to the R-S stage of the latch through transistors **N1** and **N2**, or **N2** and **N4**. High level of clock **A** enables the scan-in write operation. The 'scan' latch in Fig. 6a is a level sensitive latch controlled by clock **B**. During the scan mode the clock **C** is kept at the low level, and the R-S stage of the SA latch and the 'scan' latch work as a master-slave latch, controlled by clocks **A** and **B**, as required by LSSD.

During the normal operation mode clocks **A** and **B** are kept at the low level, and the latch operates as the conventional SA latch. The power overhead of the proposed scan extension is reduced to the drain capacitance of two minimum-sized transistors **N1** and **N3**, connected to the output nodes **Q** and **Qb**. This extra capacitance is charged or discharged at most once per clock cycle, and is not affected by spurious transitions at the data input. Thus, the power overhead of the scan extension is

$$\Delta P_1 = \frac{1}{2} f \times V_{dd}^2 \times C_1 \times \alpha,$$

where C_1 is the drain capacitance of transistors **N1** and **N3** in Fig. 6a, and α is the 'true' switching activity at the data input.

A prior art edge-triggered scannable version of the SA latch [6] is shown in Fig. 6b. During the normal mode of operation the input

Table 2: Power overhead of adding scan the SA latch.

approach	energy overhead formula	C value	value for $V_{dd} = 1\text{V}$, $\alpha = 0.3, \beta = 0.3$
prior art	$V_{dd}(V_{dd} - V_T)\Delta C_2$	14.7 fF	11.0 fJ
mux-based	$\frac{1}{2}(\alpha + \beta)V_{dd}^2 C_{mux}$	15.2 fF	4.6 fJ
proposed	$\frac{1}{2}\alpha V_{dd}^2 C_1$	2.2 fF	0.3 fJ

signal **Scan** is low, and the SA current flows through transistors **N1** or **N2**, controlled by the input data signals **D** and **Db**. During the scan mode the signal **Scan** is high, and the SA current flows through transistors **N3** or **N4**, controlled by the scan-in signals **I** and **Ib**. This implementation of the scan-in capability has a very high power overhead, because it significantly increases capacitance at the bottom part of the latch (nodes **A**, **B**, **E**, **F** and **M**). Since these nodes are charged and discharged every clock cycle, independent of the switching activity, the increase in power dissipation equals

$$\Delta P_2 = f \times V_{dd}(V_{dd} - V_T) \times \Delta C_2,$$

where ΔC_2 is the increase of the capacitance at nodes **A**, **B**, **E**, **F** and **M** in Fig. 6b.

An alternative implementation of the scan capability by means of multiplexing the input and scan-in data degrades the performance of the latch by increasing the setup time, moreover, it leads to an increase of the the power dissipation which is proportional to the sum of the input data switching activity and glitching factor,

$$\Delta P_3 = \frac{1}{2}f \times V_{dd}^2 \times C_{mux} \times (\alpha + \beta),$$

where C_{mux} is the capacitance of the multiplexor at the input, α is the input data switching activity, and β is the glitching factor at the data input.

Based on capacitance values for a 0.13μ technology, Table 2 estimates the power overhead of adding the scan feature to the SA latch using the two prior art approaches and the proposed approach. The fourth column gives the energy overhead estimates for typical values of the data switching activity and the glitching factor. The Table shows that the proposed LSSD extension reduces the energy overhead of the scannable latch 12 times (which could be even more in high glitching nodes), compared to the input multiplexed design, and 37 times, compared to the prior art design in Fig. 6b. Under the same conditions the full sense amplifier latch in Fig. 6a dissipates about 8.5fJ per clock cycle. Thus, using the proposed approach results in more than 50% power savings in the scannable SA latch, and about 30% savings in the total latch power, including the clock distribution tree.

In terms of the effect on the latch performance, the proposed scan extension has approximately the same decrease in performance as the prior art approach in Fig. 6b, and a significantly smaller decrease in performance than the multiplexor-based approach. The proposed LSSD extension can be used with many other single phase latches, including those described in [11].

4 Comparative study

We have done a comparative study of a large number of different latch styles to identify the ones that are most suitable for a low power design. Since the power supply reduction is essential for reducing power, we primarily focused on static and semi-static latches because of their higher noise margin. Also, since V_{dd} reduction plays such an important role, we were particularly interested in those latches whose performance degrades the least as V_{dd} is reduced.

In this paper we show results only for four latch styles: LSSD scannable NORA and transmission gate latches, Fig. 5, proposed LSSD sense amplifier latch, Fig 6a and semi-static true single phase RAM latch, Fig 3b, derived from [11]. The optimization described in Section 1 was applied to every latch. The optimization parameter γ was changed in the range from 0.1 to 0.9 to generate the *energy efficient* curve for each latch for $V_{dd}=0.9\text{V}$. A bulk technology was used with a 0.13μ feature size. The power and performance of the tuned latch were measured as described in Section 1. Then all *energy efficient* configurations of every latch were simulated for lower values of V_{dd} , $V_{dd}=0.8\text{V}$ and $V_{dd}=0.7\text{V}$. No additional tuning was done, however. The results are shown in Fig. 7, for the activity factor of 0.3 transitions per cycle, and the spurious activity of 0.15 glitches per cycle.




Figure 7: Average energy per cycle versus performance. Switching factor $\alpha = 0.3$, glitching activity $\beta = 0.15$ ($\beta_1 = 0.1, \beta_2 = 0.02, \beta_3 = 0.005$). Solid lines connect points of *energy efficient* configurations for every value of V_{dd} .

Extensive use of clock gating effectively increases the switching factor and the glitching activity. Fig. 8 plots the average energy per cycle of the same latches for higher value of the switching and glitching activities. Low power consumption of the LSSD sense amplifier latch even in the presence of significant glitching activity, as well as its ability to operate with reduced swing signals make it a very good candidate for low power designs.

5 Experimental data

A test site was constructed in an experimental 0.18μ CMOS process to investigate the ability of single clock sense amplifier style latches built of very small width devices to capture data with poor slew rates and low V_{dd} . The true/complement input sense amplifier style latch was modified for single ended input in two versions as shown in Fig. 9. The first version used a gate input with an added inverter of minimum size devices to drive the opposite gate. The second version used a mixed gate plus source input. The input was connected to four 1.6mm wires with 16 tristate drivers distributed along each wire. The driver data and selects are controlled by a scan latch chain and the output of the sense amplifier latch is observed at the pads. The experiment measured the maximum frequency at which the latch could capture alternating ones and zeroes inserted at the end of the long wire versus V_{dd} . Because the absolute frequency is proportional to both the large wire delay and the setup+hold time of the latch, the results of the experiment are presented in relative terms. Figure 10 shows the results over the voltage range equal to 0.55V to 1.5V . Both circuits operated




Figure 8: Average energy per cycle versus performance. Switching factor $\alpha = 0.5$, glitching activity $\beta = 0.55$ ($\beta_1 = 0.2, \beta_2 = 0.1, \beta_3 = 0.05$). Solid lines connect points of energy efficient configurations for every value of Vdd.




Figure 9: Latches on the test chip: **a** – gate input SA latch, **b** – mixed input SA latch.

over the full voltage range, including the extreme low voltage corresponding to $V_{TP} + V_{TN}$. The conventional gate input showed better setup+hold time as well as lower energy.

6 Conclusions

Power-performance optimization methodology for latches was extended to formally parameterize latch power in terms of switching factor and glitching activity. The concept of energy efficient family of configurations was introduced and used for formal comparison of different latch styles in the power-performance space. Clock distribution power was found to be a significant component of the total power of the clocking system in low-power designs, and latches using a single clock phase were found to dissipate 30% less power than those requiring two phases of clock. Practical issues of building a low power overhead scan mechanism were considered, and a low-power LSSD extension to single-phase latches was proposed and demonstrated to significantly reduce the power overhead of LSSD design. Results of a comparative study of LSSD latches are shown, and the modified sense amplifier latch with the proposed LSSD extension was found to be a very strong candidate for low-power designs. Results are supported by experimental data measured on a 0.18μ test chip which showed a robust operation of the sense amplifier latch built of very small width devices over the full voltage range. Gate-input version of the sense amplifier latch

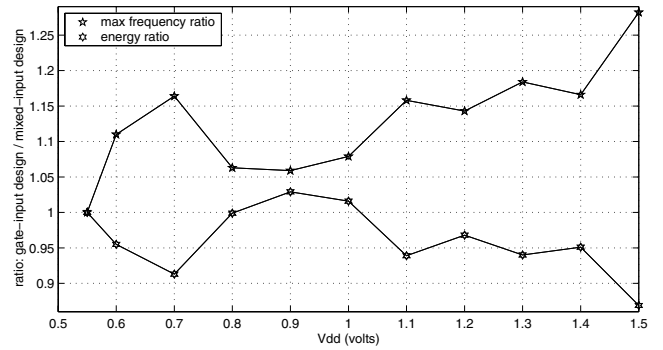


Figure 10: Latch comparison experiment: gate-input (Fig. 9a) versus mixed-input designs (Fig. 9b).

was found to be both faster and lower-power than the mixed-input variation.

Acknowledgment

The authors would like to thank S. Kosonocky, K. Chin, A. Haen, D. Knebel and W. Hwang for useful discussions; M. Immediato for measuring experimental data; G. Gristede for tool support; and K. Warren and J. Moreno for management support.

References

- [1] S. Faris. Circuit design for full scan ATPG. In *Forth Annual IEEE International ASIC Conference*, pages 6.1–6.4, 1991.
- [2] T. Lang, E. Musoll, and J. Cortadella. Individual flip-flops with gated clocks for low power datapaths. *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, 44(6):507–516, June 1997.
- [3] J.Y. Lin, T.C. Liu, and W.Z. Shen. A cell-based power estimation in CMOS combinational circuits. *ICCAD*, pages 304–309, 1994.
- [4] N. Nedovic and V. Oklobdzija. Dynamic flip-flop with improved power. In *Proceedings of the International Conference on Computer Design*, 2000.
- [5] N. Nedovic and V. Oklobdzija. Hybrid latch flip-flop with improved power efficiency. In *Proceedings of the Symposium on Integrated Circuits and Systems Design*, 2000.
- [6] B. Nikolic et al. Improved sense-amplifier-based flip-flop: Design and measurements. *IEEE Journal of Solid-State Circuits*, 35(6):876–883, June 2000.
- [7] J. H. Satyanarayana and K. K. Parhi. Heat: Hierarchical energy analysis tool. *Proceedings of the 33rd Design Automation Conference*, pages 9–14, June 1996.
- [8] V. Stojanovic and V. Oklobdzija. Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems. *IEEE Journal of Solid-State Circuits*, 34(4):536–548, April 1999.
- [9] V. Stojanovic, V. Oklobdzija, and R. Bajwa. A unified approach in the analysis of latches and flip-flops for low-power systems. In *Proceedings of the International Symposium on Low Power Electronics and Design*, pages 227–232, August 1998.
- [10] C. Svensson and J. Yuan. Latches and flip-flops for low power systems. In A. Chandrakasan and R. Brodersen, editors, *Low Power CMOS Design*, pages 233–238. IEEE Press, 1998.
- [11] J. Yuan and C. Svensson. New single-clock CMOS latches and flipflops with improved speed and power savings. *IEEE Journal of Solid-State Circuits*, 32(1):62–69, January 1997.
- [12] V. Zyuban and P. Kogge. Application of STD to latch-power estimation. *IEEE Transactions on VLSI Systems*, 7(1), March 1999.