Optimization of Scannable Latches for Low Energy

Victor Zyuban, *Member, IEEE*

***Abstract*—This paper covers a range of issues in the design of latches and flip-flops for low-power applications. First it revisits, extends, and improves the energy-performance optimization methodology, attempting to make it more formal and comprehensive. The data-switching factor and the glitching activity are taken into consideration, using a formal analytical approach, then a notion of an energy-efficient family of configurations is introduced to make the comparison of different latch styles in the energy-performance space more fair. A recently proposed methodology for balancing hardware intensity in processor pipelines is applied to latch design to facilitate the selection of the objective function for tuning transistor sizes. The power dissipation of the clock distribution is taken into account, supported by simulations of extracted netlists for multibit datapath registers. Practical issues of building a low overhead scan mechanism are considered, and the power overhead of the scannable design is analyzed. A low-power level-sensitive scan mechanism is proposed, and results of a comparative study of scannable latches are shown. The applicability of the proposed scan mechanism to a wide variety of latches is demonstrated.**

***Index Terms*—Circuit tuning, curcuit power, energy, flip-flop, glitching, latch.**

I. Introduction

SINCE THE IMPORTANCE of designing low-power high-performance timing elements has been recognized, many low-power latch and flip-flop studies have been published [28], [10], [26], [2]. A number of papers have proposed new styles of latches and flip-flops that improve speed [8], [9], [13], [18], [19], [27], reduce energy [7], [14], [16], [17], [30], or demonstrate improvements in other important characteristics of timing elements, such as race immunity, noise immunity, and testability [12], [14]. Other works have compared various latch styles in the energy-performance design space and established a robust methodology for the comparison [6], [12], [15], [23], [24], [29].

This paper extends the existing power-performance optimization methodology in several aspects. First the methodology is formalized by using analytical formulas to take into account both the *data switching activity* and the *glitching factor*, based on [29]. A formal optimization of every latch style in the energy-performance space is performed before comparing different latch styles through constructing *energy-efficient families* of configurations for every latch. Issues of tuning transistor sizes in latches are considered, and efficient methods for building energy-efficient families are described.

In this work, we apply a recently introduced methodology [31] for balancing *hardware intensity* in processor pipelines to the latch design and derive relations which will help designers balance the tuning target in latches with that in logic, in an energy-efficient manner. Our analysis indicate that in many today's microprocessors latches are over designed, when compared to the rest of the hardware.

Breakdowns of the power for some of the commonly used latches show that the power of the local clock distribution network may constitute a significant portion of the power dissipated by a latch. Therefore, when comparing different latches for energy efficiency it is essential that the power dissipated in the local clock distribution network be taken into account, because different latches present different requirements as well as different amounts of capacitive load on the local clock distribution network. In this work we analyze power dissipated by the clock distribution network, and compare the single-phase clocking scheme with two-phase clocking.

One of the practical issues of building latches for testable designs is the integration of the scan mechanism into the latch. As shown in this paper, the power overhead of the scannable design can be significant, meanwhile the complexity of modern microprocessors has reached the point where saving power by implementing a nonscannable design is not viable. In this paper, we show a low-power overhead level-sensitive scan mechanism which can be used with a wide variety of recently reported edge-triggered and pulsed latches. We furthermore, compare it with other approaches for building scannable designs in terms of energy, based on simulations of extracted layouts, and give areas of the layouts drawn in a state of the art technology.

II. Optimization and Comparison Methodology

A. Performance Measurement

The methodology for comparing the performance of different latches used in this paper consists in evaluating the following metric [23], [24], [18], [27], based on simulation of the switching of the latch for varying values of the data setup time

$$\begin{aligned} D &= T_{\text{setup}} + D_{C\to Q} \\ &= \min[T_{D\text{-}to\text{-}C} + \max(D_{0\to 1}, D_{1\to 0})], \end{aligned} \tag{1}$$

where T_{setup} is the setup time, and $D_{C\to Q}$ 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, $T_{D\text{-}to\text{-}C}$, as shown in Fig. 1. For the design in Fig. 1, the minimum value for the sum is

Manuscript received March 5, 2002; revised August 8, 2002.

The author is with IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: zyuban@us.ibm.com).

Digital Object Identifier 10.1109/TVLSI.2003.814322

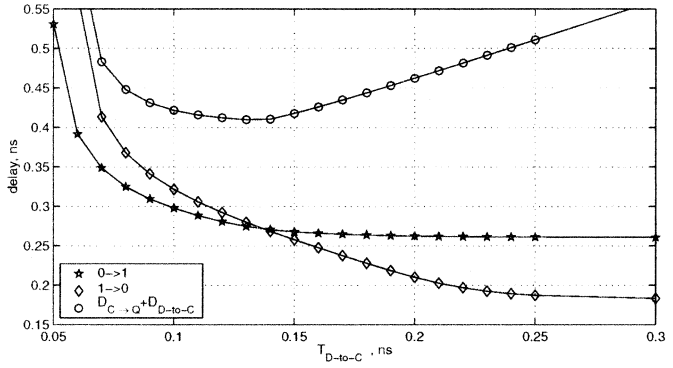

Fig. 1. Evaluation of the performance metric.

reached when the delay between the data and clock transitions is $T_{D\text{-}to\text{-}C} = 130$ ps, and the delay through the latch at this point is $\max(D_{0\to1}, D_{1\to0}) = 280$ ps. Thus, we put for this latch $T_{\text{setup}} + D_{C\to Q} = 130$ ps $+ 280$ ps $= 410$ ps. All delays are measured assuming a load of four minimum size inverters and a 50-track-long wire capacitance.

Another performance metric used in other works [14] consists in evaluating the sum

$$D = T_{\text{setup}} + (1+\delta)D_{C\to Q} \tag{2}$$

where $D_{C\to Q}$ is the nominal delay and T_{setup} is defined as data to clock offset that corresponds to a relative increase by δ in the $C \to Q$ delay from its nominal value. The selection of the value of δ is somewhat arbitrary. The value of $\delta = 0.05$ was used in [14]. Although, by construction of metrics (1) and (2), the latch insertion penalty D measured using (2) is always higher than that measured with (1), for all latches that were simulated in this work both metrics produce very close results. For example, for the curves in Fig. 1, the values of T_{setup} and $D_{C\to Q}$, measured according to (2) with $\delta = 0.05$ are 136 ps and 261 ps, respectively, which results in approximately the same value for D as measured with (1), $D = 410$ ps.

B. Energy 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 with respect to clock. The typical methodology used in most of the prior art is to estimate the power of a latch for two values of the switching activity at the data input: $\alpha = 0$ and $\alpha = 1$ [14], [18], and then calculate the average power as a linear combination of the power under these extreme cases, with the weights depending on the data switching factor α. Other works report energies dissipated for a limited set of input data patterns, assuming a free running clock [2], [7], [23], [24], The *spurious* switching activity, or *glitching* at the data input is typically either neglected [2], [24], [25], [27] or added afterwards [14].

Our analysis shows that in low-power designs, especially those with relatively shallow pipelines, the power dissipated in latches due to the spurious activity at the data input is *not* negligible. Table I shows the measured average number of spurious transitions per cycle at the outputs of several functional units. The second column shows the critical path delay, measured in FO4 units. The glitching activity was measured by simulating functional units for random input patterns with the switching factor $\alpha = 0.3$ applied to the inputs for 200 cycles. The total number of transitions was measured at the outputs of the functional units, then the simulation runs were repeated under the zero-delay model to measure the value of the "true" switching activity at the outputs. The difference between the two values gives the average number of spurious transitions per cycle.

TABLE I
AVERAGE NUMBER OF TRUE AND SPURIOUS TRANSITIONS PER CYCLE AT THE OUTPUTS OF FUNCTIONAL UNITS FOR RANDOM INPUT PATTERNS WITH $\alpha_{in} = 0.3$

unit	FO4 delay	true trans.	spurious trans.
saturated adder	17	0.41	0.28
int arithmetic	13	0.50	0.23
shifter	11	0.32	0.15
logic unit	7	0.49	0.12
multiplier tree	12	0.42	0.21
multiplier array	20	0.47	0.60

In microprocessors that extensively use the clock gating technique all transitions at the latch data input (true transition and glitches) while clock is gated can be modeled as spurious transition, which effectively increases the glitching factor [6]. Therefore, the effective glitching factor in clock-gated designs can be somewhat higher than the values reported in Table I.

Our simulations show that, depending on the latch circuitry, the data true switching factor, and time positions of spurious transitions at the data input, the presence of the glitching activity of 0.3 transitions per cycle increases the average energy dissipated by a latch by from 3% to 30%. For some latches the energy effect of glitching can be measured sufficiently accurately by simulating the energy dissipated by one glitch at the data input and then multiplying it by the average number of glitches per cycle [14], while for many latches such a simplification may produce an error of as high as 10%. Considering that many of competing latch styles proposed or compared in recent works differ by no more than 5% in energy dissipation, we believe, applying a more detailed analysis of energy dissipation in latches is justified.

A formal approach for analyzing energy dissipated in latches developed in [29] models the circuit as a directed graph, called the *state transition diagram* (STD) [11], [21], such that there is a one-to-one correspondence between edges in the graph and power-dissipating events in the circuit. An algorithm to construct the STD for a latch is described in [29]. As an example, Table II shows all reachable states for the sense amplifier latch [3], shown in Fig. 3(a), not distinguishing between states that only differ by voltage levels at nodes *G* and *H*. The eight columns in the table give the voltage levels at all nodes in the circuit for every state. The corresponding STD is shown in Fig. 2(a).

It is important to emphasize that we attribute the energy dissipated for charging/discharging the capacitances at the clock and data inputs to the latch itself, rather than to a fan-in gate or

TABLE II
REACHABLE STATES IN THE STD OF THE SA LATCH. "1" AND "0" DESIGNATE VOLTAGE LEVELS AT THE NODES OF THE CIRCUIT

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

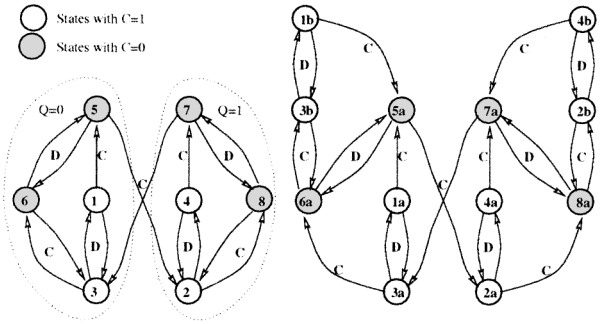

Fig. 2. State transition diagram of the (a) symmetric sense amplifier latch (simplified); and (b) proposed nonsymmetric SA latch.

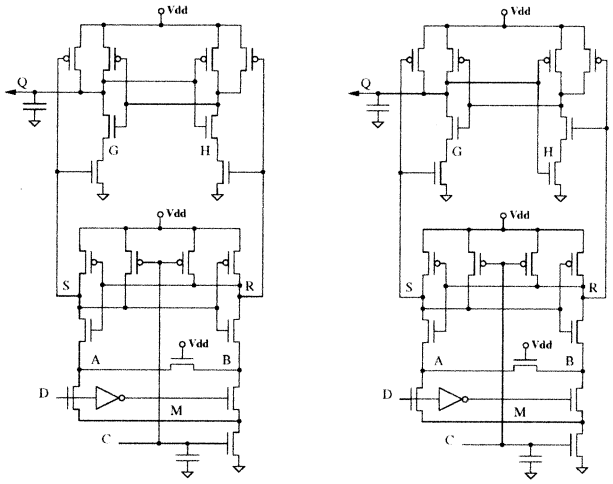

Fig. 3. Transistor diagrams. (a) Symmetric sense amplifier latch, (b) Proposed nonsymmetric sense amplifier latch.

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 energy dissipated for charging/discharging the load driven by a latch into latch energy. However, the output capacitance of the wiring *within* the latch layout *must* be included. Such a convention makes the energy comparison between different latches more fair, however this also makes the STD more complicated compared to those presented in other works [23].

Energy weights for every edge in the STD are calculated using a simulator or, for rough estimates, manually, using the formula $(1/2)\sum C_i V_{dd}\triangle V$, where $\sum C_i$ is the sum of capacitances at all nodes that have different voltage levels in the states connected by the edge. As an example, using Table II, the energy weight of path $\{s_7 \rightarrow s_3\}$ that corresponds to the latching of a new data (zero) at the rising clock edge is $E_{7\rightarrow 3} = (1/2)V_{dd}^2(C_C+C_Q+C_{Qb}+C_R)+(1/2)V_{dd}(V_{dd}-V_{TH})(C_A+C_B+C_M+C_H)$. To arrive at this expression, capacitances at all nodes in Table II that have different logic levels in columns three and seven were multiplied by $(1/2)V_{dd}$ and the corresponding voltage swing.

Then, based on the probabilistic analysis of the STD, presented in [29], 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

$$P_{\text{true}} = f(Q_0 + P_1 Q_1 + \alpha Q_2) \tag{3}$$

where

$$Q_0 = E_{00}, \qquad Q_1 = E_{11} - E_{00}$$
$$Q_2 = \tfrac{1}{2}\left(E_{01} + E_{10} - E_{00} - E_{11}\right).$$

In these formulas α is the switching activity as defined in [29], 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 STD in Fig. 2(a) 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 the energy weights of all edges in this path.

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 all derived 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 (3) is derived in [29] that is valid for all reasonable latches known to the author

$$P_{\text{total}} = \beta_0 P_{\text{true}} + f[\beta E_{\text{cycle}} + \beta^* E^*_{\text{cycle}} + (1-\beta_0)(Q'_0 + P_1 Q'_1 + \alpha Q'_2)] \tag{4}$$

where

$$Q'_0 = \sum_{p^1_{00}} E_{ij} - E_{\text{cycle}}, \qquad Q'_1 = \sum_{p^1_{11}} E_{ij} - \sum_{p^1_{00}} E_{ij}$$
$$Q'_2 = \tfrac{1}{2}\left(\sum_{p^1_{01}} E_{ij} + \sum_{p^1_{10}} E_{ij} - \sum_{p^1_{00}} E_{ij} - \sum_{p^1_{11}} E_{ij}\right).$$

Here, P_{true} is the average power in the absence of glitches calculated by (3), 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 exactly k spurious pulses occur during one clock cycle. In particular, β_0 is the probability that zero spurious pulses occur.

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 $\beta^* E^*_{\rm cycle}$, where β^* is the average number of spurious pulses per cycle occurring while clock is high. In these formulas p^k_{ij} denote a path in STD traversed when i was latched in the previous clock cycle and the true data value has changed to j, m, $n = \{0, 1\}$ (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. 2(a), $p^1_{00} = \{s_3 \rightarrow s_6 \rightarrow s_5 \rightarrow s_6 \rightarrow s_3\}$, and $p^1_{01} = \{s_3 \rightarrow s_6 \rightarrow s_5 \rightarrow s_6 \rightarrow s_5 \rightarrow s_2\}$. $E_{\rm cycle}$ is the energy dissipated by one spurious pulse, provided that at least one spurious pulse has occurred before it. For the STD in Fig. 2(a), $E_{\rm cycle}$ is the energy dissipated on the cycle $p_{\rm cycle} = \{s_6 \rightarrow s_5 \rightarrow s_6\}$, or $p_{\rm cycle} = \{s_8 \rightarrow s_7 \rightarrow s_8\}$, and $E^*_{\rm cycle}$ is the energy dissipated on the cycle $p^*_{\rm cycle} = \{s_3 \rightarrow s_1 \rightarrow s_3\}$, or $p^*_{\rm cycle} = \{s_2 \rightarrow s_4 \rightarrow s_2\}$.

Thus, the average total latch power in (4) 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 some latches [in particular, for those represented by the STD in Fig. 2(a)], it cancels with $\beta_0 P_{\rm true}$, and the expression reduces to $P'_{\rm total} = P_{\rm true} + f\left[\beta E_{\rm cycle} + \beta^* E^*_{\rm cycle}\right]$. In this special case, the additional energy due to glitches at the data input can be measured by simulating the energy dissipated by one glitch and then multiplying it by the average number of glitches per cycle.

C. Proposed Nonsymmetric Sense Amplifier Latch

To demonstrate the usefulness of the STD technique, we show how a simple modification to the sense amplifier latch for single-rail designs may reduce the delay through the latch with *zero* increase in energy or area. The modification of the sense amplifier latch, shown in Fig. 3(b) 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 [path $\{7 \rightarrow 3\}$ in Fig. 2(a)] is minimized at the expense of somewhat higher capacitance charged/discharged on other transitions in the STD.

The state table of the modified nonsymmetric latch is shown in Table III, and the STD is shown in Fig. 2(b). According to Tables II and III the capacitance charged/discharge on the slower Q: $1 \rightarrow 0$ transition in the modified latch is reduced by the capacitance C_H at node H, compared to the original symmetric latch in Fig. 3(a). The capacitance charged/discharge on path $\{5 \rightarrow 2\}$ that corresponds to the faster Q: $0 \rightarrow 1$ transition in the modified latch is increased by C_H. As a result, the worst-case delay is reduced, at the expense of some increase in the delay of the faster transition, which improves the performance characteristics of the latch, according to the performance metric, defined in Section II-A. Since the total capacitance charged/discharged on any complete cycle in the STD in Fig. 2(b) remains unchanged, the described modification does not have any power overhead. We also verified that this modification does not incur any layout area overhead.

TABLE III
REACHABLE STATES IN THE STD OF THE NONSYMMETRIC 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

Simulations show that this simple modification results in 5% worst-case delay reduction, compared to the configuration where both "R" and "S" signals are connected to the lower inputs of the NAND gates, Fig. 3(a), and in 3% delay reduction compared to the configuration where both "R" and "S" signals are connected to the upper inputs [7] of the NAND gates.

D. Tuning Transistor Sizes

Before comparing it with other latches, every latch must be optimized for an energy-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 [6].

To avoid this uncertainty, *energy-efficient family* of configurations should build for every latch, which is a set of configurations, each of which results in the highest performance among all configurations dissipating the same power, or the one that dissipates 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. As a demonstration, Fig. 4 shows an example of building an energy-efficient family for a low-power version of the sense amplifier latch built in an experimental 0.18 μm technology.

It was shown in [20] and [31] that the energy-efficient curve in energy-versus-delay coordinates can be approximated by relation

$$(E - E_0)(D - D_0) = \kappa E_0 D_0, \tag{5}$$

where κ is a technology-dependent constant, whose value was empirically measured to be in the range from 0.2 to 0.4 for a 0.13 μm bulk technology, depending on the circuit type, and D_0 and E_0 are the asymptotes of delay and energy that can be approached through tuning the circuit.

In this work we built energy-efficient families by optimizing every latch to minimize the cost function

$$F_c = \gamma(E/E_0)^2 + (1-\gamma)(D/D_0)^2, \qquad 0 \le \gamma \le 1 \tag{6}$$

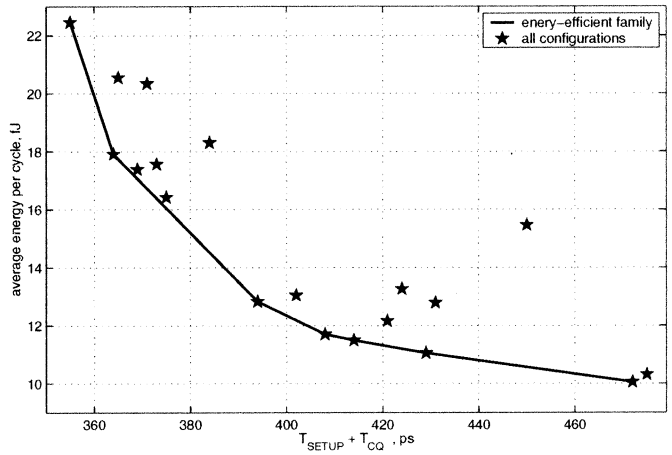

Fig. 4. Building energy-efficient family for a latch.

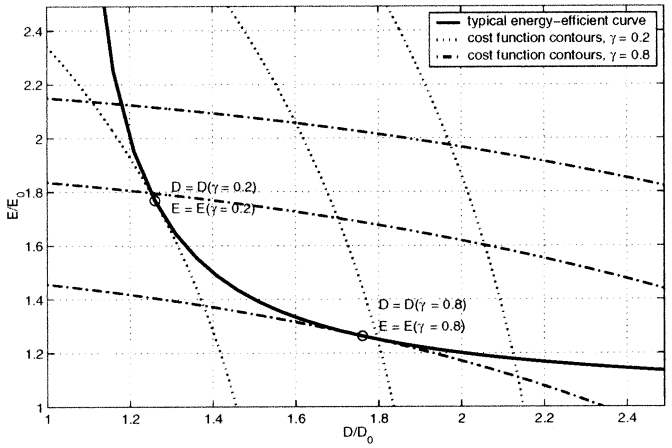

Fig. 5. Typical energy-efficient curve and constant cost function contours for $\gamma = 0.2$ and $\gamma = 0.8$.

where D is the sum of the setup time and the delay through the latch, as defined in Section II-A, and E is the average energy dissipated by the latch in one clock cycle, determined according to Section II-B for $\alpha = 0.3$ and $\beta = 0.16$, D_0 and E_0 are estimates of the corresponding lower bounds.

Fig. 5 gives a graphical interpretation of the process of building the energy-efficiency family using cost function (6). The solid line plots a typical energy-efficient curve for a latch. Dotted lines show several contours of the cost function (6), for two values of optimization parameter γ. Point (D, E) at which the energy-efficient curve touches the lowest of the contours ($F_c = A$ with the smallest value of A) corresponds to the energy-efficient implementation for this value of γ.

Many types of functions can be used as a cost function [20]. We used (6) because it resulted in a consistent convergence of the circuit tuner. Sometimes in theoretical studies the following expression [1], [5], [20], [22], [30], [31] is used for the cost function:

$$F_c = (E/E_0)(D/D_0)^{\eta} \tag{7}$$

where η is an optimization parameter, $\eta \geq 0$. The convergence of the tuning process using an objective function of this form requires that the curvature of the energy-delay curve satisfies $\partial^2 E/\partial D^2 > (1/E)(\partial E/\partial D)^2 - (1/D)(\partial E/\partial D)$. This constraint is typically satisfied in CMOS circuits which have a nonzero lower bounds on delay and energy, D_0 and E_0. Then the tangent to the energy efficient curve $\partial E/\partial D$ at point (D, E) is related to the optimization parameter η in (7) as

$$\frac{\partial E}{\partial D} = -\frac{\partial F_c}{\partial D} \Big/ \frac{\partial F_c}{\partial E} = -\eta \frac{E}{D} \tag{8}$$

or

$$\eta = -\frac{D \partial E}{E \partial D}. \tag{9}$$

Thus, η is the ratio of the relative increase in energy to the corresponding relative gain in delay achieved through transistor tuning for designs on the energy-efficient family. Simply put, it is the value of % power per % performance for an energy-efficient design. Because of this property optimization parameter η in (7) was called a *hardware intensity* in [31]. In other words, hardware intensity is a measure of how aggressively a hardware block has to be tuned to meet the delay requirement.

If the optimization function (7) is used for building an energy-efficient family, the convergence of the tuning process is not as good as with function (6), because the contours of function (7) have a shape (curvature) close to that of the typical energy-efficient curve (5). However, for any energy-efficient curve there is a one-to-one correspondence between optimization parameter γ in (6) and the energy efficiency η in (7)

$$\eta = \frac{1-\gamma}{\gamma} \frac{(D/D_0)^2}{(E/E_0)^2}. \tag{10}$$

This allows us to use for optimizing latches a cost function that results in a consistent convergence of a tuner, and still use the theoretical expressions for hardware intensity derived in [31].

Typically, the cycle time requirement can be met at different combinations of hardware intensity η and power supply voltage V_{dd}. In [31], a condition for the optimal balance between V_{dd} and η was derived, such that for a given critical path delay requirement, the energy reaches its minimum over the two-dimensional space (η, V_{dd}). For a pipeline stage that consists of M blocks that are designed and optimized independently, the optimal balance between the power supply voltage and the hardware intensity is achieved when

$$\eta_i = \frac{u_i}{w_i} \theta(V_{dd}), \qquad 1 \leq i \leq M \tag{11}$$

where η_i is the hardware intensity in block i, u_i is the delay weight of block i, $u_i = D_i/D$, and w_i is the corresponding energy weight, $w_i = E_i/E$, calculated taking into account the activity factors. Here D is the clock period and E is the average total energy of the pipeline stage. The voltage intensityθ is defined in [31] as a ratio of the dimensionless derivatives of the delay and energy with respect to the power supply voltage

$$E_v = \frac{v}{E} \frac{\partial E}{\partial v}, \qquad D_v = -\frac{v}{D} \frac{\partial D}{\partial v}, \qquad \theta = \frac{E_v}{D_v}. \tag{12}$$

Voltage intensity function $\theta(V_{dd})$ depends on many factors. For a 0.13 μm bulk technology, $\theta(V_{dd})$ was measured to grow ap-

proximately linearly with V_{dd}, from a value of 0.5 at the low range of V_{dd} to a value of 3 at the upper limit of V_{dd}.

Let us apply (11) to a pipeline stage that consists of a block of latches and a cloud of logic which are typically designed and tuned independently. Assume that $\theta = 2$ for the chosen design point, and assume that the latch delay constitutes 20% of the cycle time and the delay through the logic is 80% of the cycle time. Furthermore, assume that latches are responsible for 30% of the total power of the pipeline stage, not including the power of the clock distribution. Then, using (11), the optimum hardware intensity for latches is $\eta_{\text{latch}} = (0.2/0.3)2.0 = 1.3$, and that for the logic is $\eta_{\text{logic}} = (0.8/0.7)2.0 = 2.3$. Thus, for the assumptions above, latches should be optimized less aggressively than logic. With the help of (11) the actual value of optimal hardware intensity for latches and logic can be determined for any microprocessor.

III. Clock Distribution Power

When comparing different latches for energy efficiency it is essential that the power dissipated in the local clock distribution network be taken into account, because different latches present different requirements as well as different amounts of capacitive load on the local clock-distribution network. Moreover, some latch styles require two clock phases, others use only one phase, and pulsed latches require a pulse generation circuitry. The power overhead of the clock gating mechanism also needs to be taken into account. The power dissipated in upper levels of the clock distribution network largely depends on the chip size, the total number of latches and their locations, and much less on the latch design. Therefore, the power of the global clock-distribution network can be analyzed independently.

To determine the portion of the total power that is dissipated in the local clock-distribution network we developed layouts for 16, 32, and 40 bit datapath registers built of single-phase nonsymmetric sense amplifier latches, Fig. 3(b) and two-phase master-slave C^{2}MOS latches, Fig. 8, both designed in a 0.13 μm technology, with a 12 track bit step. The latches were optimized for low values of hardware intensity $\eta = 1$, which, according to (11) corresponds to the energy-efficient operation point at $V_{dd} = 0.9$ V [31]. The designs included the local clock buffers and the clock gating mechanism which consists of a retiming latch and a NAND gate. For two-phase latches the local clock buffer includes a clock splitter that generates two nonoverlapping phases of clock.

We ran power measurements on the extracted netlists, according to the methodology described in the previous section, for the data switching factor $\alpha = 0.3$, and the glitching activity $\beta = 0.16$. Figs. 6 and 7 show the power breakdown for the two-phase C^{2}MOS and single-phase sense amplifier latches. The total power of the multibit registers which includes the power dissipated in the local clock buffer and clock gating mechanism was divided by the number of bits in those registers. The results in Figs. 6 and 7 correspond to the 16-bit registers, the corresponding results for wider register are similar, except that the power overhead of the retiming latch is shared between a larger number of individual latches.

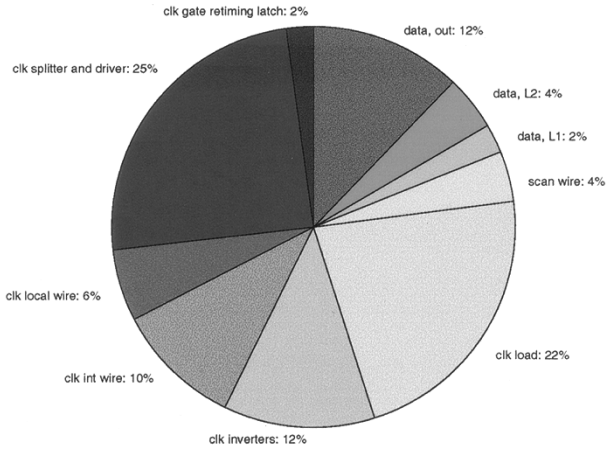

Fig. 6. Power breakdown for the two-phase C^{2}MOS datapath latch.

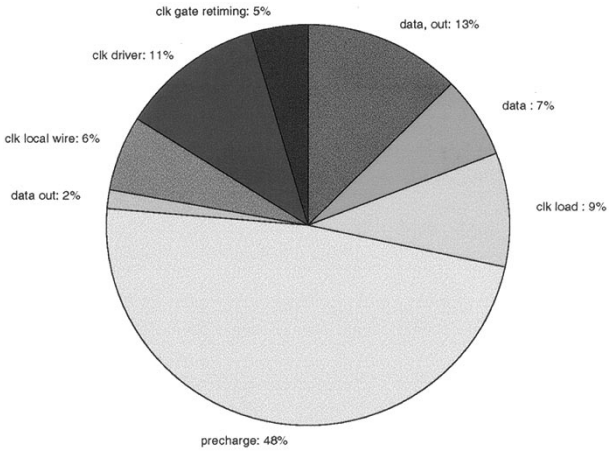

Fig. 7. Power breakdown for a SA datapath latch.

Results in Figs. 6 and 7 show that the power dissipated in the local clock distribution network is significant, and it is different for different latch styles, constituting 31% of the total power of the two-phase C^{2}MOS latch and 17% of the total power of the single-phase sense amplifier latch. The power portion associated with the clock distribution is typically even higher in autoplaced or ASIC designs, because of longer local-clock wires. It can also be more significant in high-frequency designs, where the objective of minimizing the clock skew is achieved by reducing the transition times in the clock distribution network, which requires using larger transistor sizes in clock buffers.

Another observation from data in Fig. 6 is that for calculating the load presented by a latch to the local clock-distribution network, it is important not to neglect the capacitance of the clock wiring inside the latch cell (clock internal wire in Fig. 6). We found that for latch designs using very small transistor sizes, the internal wiring of the clock may represent from 5% to 20% of the total capacitive load on the clock.

Finally, our 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 at least a 15% power advantage over a latch that requires two phases. As mentioned above, the power overhead of generating and distributing the second clock phase may be even higher (as high as 30%) in high-frequency designs, or ASIC designs.

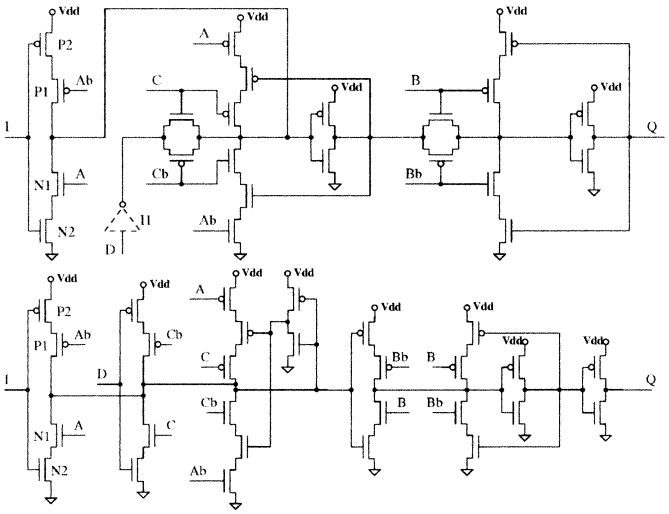

Fig. 8. LSSD transmission gate latch (above) and C^{2}MOS latch (below).

IV. Scannable Latches

The integration and complexity of modern systems has grown to the point where saving power by building a nonscannable design is no longer an option. The power overhead of the scannable design may be significant. For example, the study in [4] has reported a 54% increase in power of an LSSD standard cell design over the identical nonscannable design. In this section we analyze the power overhead of existing approaches to building scannable latches and describe a new, low power level sensitive scan mechanism that can be applied to a variety of single phase and two-phase latches, including edge-triggered latches, pulsed latches [8], [16], [23], [24], [27], and dual edge triggered flip-flops [15], [17], [27].

There exist two major approaches to building scannable designs: edge-triggered and level sensitive, level sensitive scannable design (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 [4]. For this reason LSSD is the scan mechanism of our choice.

The standard way of implementing an LSSD master-slave latch is shown in Fig. 8 for the transmission-gate latch, and C^{2}MOS latch. In these latches, the power overhead of the scan mechanism in the normal operation mode consists of the power of charging/discharging the drain capacitances of transistors $\mathbf{N1}$ and $\mathbf{P1}$ (which are cutoff in the normal operation mode), the gate capacitances of transistors $\mathbf{N2}$ and $\mathbf{P2}$ in the next latch of the scan chain, and the capacitance of the wire that connects latches in the scan chain, $C_{\text{scan}} = 2C_d + 2C_g + C_{\text{scan wire}}$.

The length of the scan wire can be quite small (10 μ to 20 μm) between adjacent latches in custom datapath registers, however it is much longer (several hundred μm) for wires that connects the scan output of the last latch in a register with the scan input of the first latch in the next register in the scan chain. Moreover, extra buffering may be required to decouple the capacitance of the long scan wire from the datapath. Also, the scan wire capacitances are much higher in ASIC designs. We will use a somewhat optimistic value of 40 μm for the average length of the scan wire, which translates into approximately 8 fF of capacitance in a 0.13 μm technology.

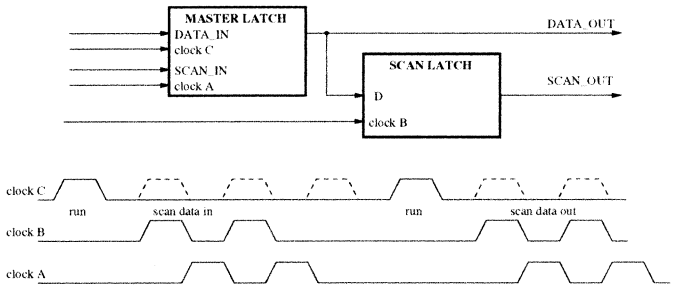

Fig. 9. Low-power level-sensitive scan mechanism.

The extra capacitance of the scan mechanism C_{scan} is charged/discharged during the normal operation mode every time the output of the latch changes, so that the energy overhead of the scan mechanism in the normal operation mode is $E_{\text{scan}} = (1/2)\alpha V_{dd}^2 C_{\text{scan}}$ per latch. For the listed assumptions C_{scan} is approximately 12 fF for a low-power custom design, but may be higher for the ASIC design flow. For a "typical" low-power microprocessor with 30 000 latches, running at $f = 400$ MHz at $V_{dd} = 1.2$ V, with average $\alpha = 0.3$ the power overhead of the scan mechanism is approximately 31 mW, or even higher, which may not be negligible in state of the art low-power designs.

Even if the power overhead of the conventional scan mechanism is tolerable, the scannable latches in Fig. 8 require two phases of clock, $\mathbf{C}$ and $\mathbf{B}$, in the normal operation mode, which, according to Section III, increases the total power of the clocking system in a low-power microprocessor by 15% to 30%.

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 nonoverlapping clock phases, as required by the LSSD standard. Also, to reduce the power overhead of the scan, it is desirable to separate the scan output of the latch from the data output, so that the wire connecting latches in the scan chain does not toggle in the normal operation mode.

Fig. 9 shows the proposed scan mechanism that has this property. The master latch in Fig. 9 can be any type of a single phase latch, or a two-phase latch, for example, edge-triggered latch, pulsed latch [23], [27], or dual edge triggered flip-flop [17], [15].

The scan latch is a low area, slow, level-sensitive latch, controlled by clock B. The output of the scan latch is the scan output of the entire flip-flop. It is connected to the scan input of another latch in the scan chain. During normal operation mode, clock A and clock B are kept at the low level, and the flip-flop works as a conventional latch, whereas scan latch is in the nontransparent state, so that the scan output does not toggle, and the internal capacitances inside the scan latch do not toggle either. This reduces the power dissipation in the normal operation mode. During the scan mode, clock C is kept at the low level, and the flip-flop works as a master-slave latch, controlled by nonoverlapping clocks A and B, providing a robust, level-sensitive scan operation.

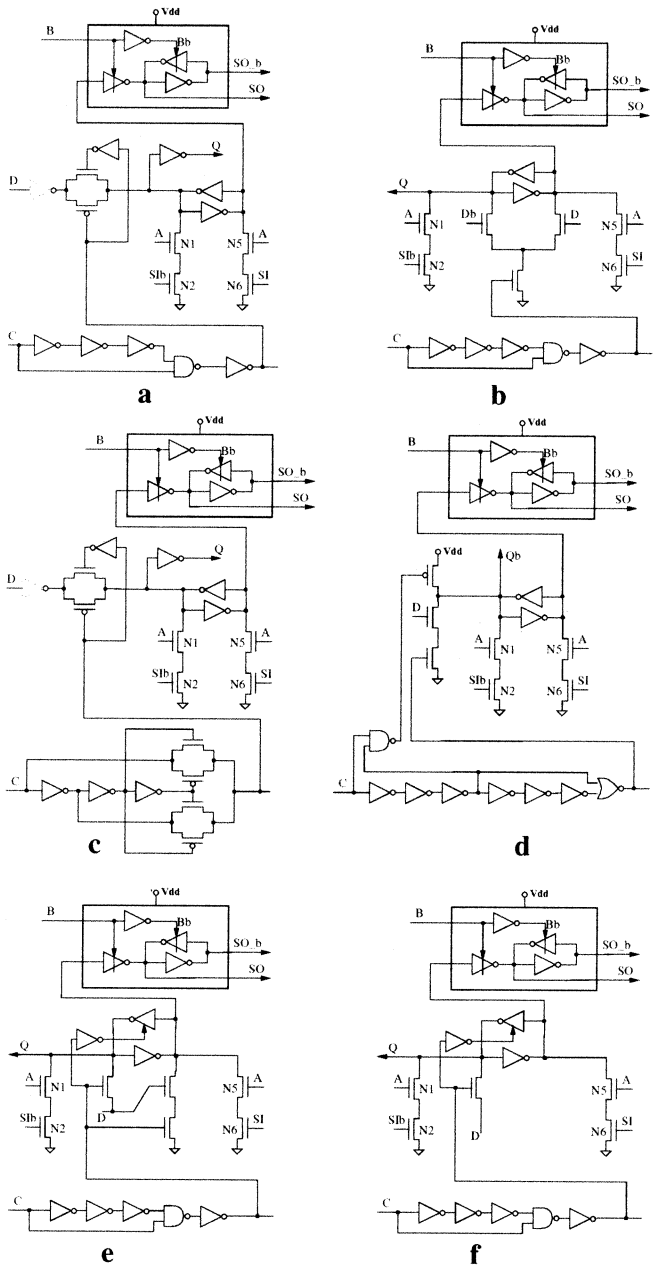

Fig. 10. Scannable latches. (a) ep-SFF. (b) SSASPL. (c) ep-DSFF. (d) DPSCRFF. (e) PPCF. (f) N-CF.

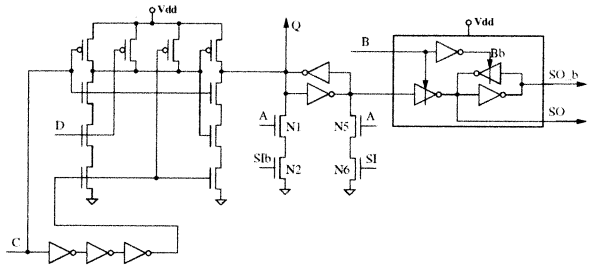

Fig. 11. Scannable HLFF latch.

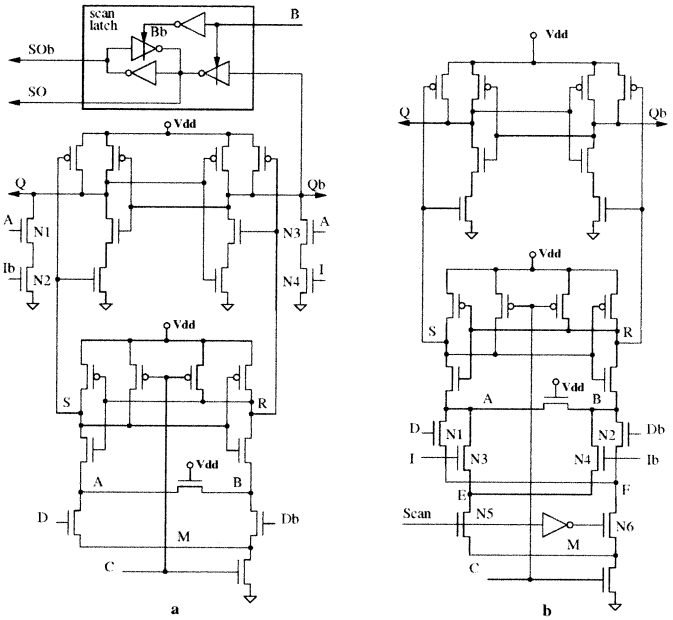

Fig. 12. Scannable sense amplifier latch. (a) Proposed LSSD scannable latch. (b) Prior art edge-triggered scannable SA latch.

Figs. 10, 11, and 12(a) show implementation examples of the above scan mechanism, applied to some of the recently published latches [13], [15], [17], [23], [27].

The low-energy overhead of this scan mechanism is achieved by mixing in the scan-in data at the second stage of the latch. The scan-in data signal, $\mathbf{I}$ is written in through transistors $\mathbf{N1}$ and $\mathbf{N2}$, or $\mathbf{N5}$ and $\mathbf{N6}$. High level of clock $\mathbf{A}$ enables the scan-in write operation. The "scan" latch is a level sensitive latch controlled by clock $\mathbf{B}$. During the scan mode clock $\mathbf{C}$ is kept at the low level, and the master stage of the latch and the "scan" latch work as a master-slave latch, controlled by clocks $\mathbf{A}$ and $\mathbf{B}$, as required by LSSD.

During the normal operation mode clocks $\mathbf{A}$ and $\mathbf{B}$ are kept at the low level, and the latch operates as a conventional latch. The energy overhead of the proposed scan extension is reduced to the drain capacitance of two minimum-sized transistors $\mathbf{N1}$ and $\mathbf{N5}$, connected to the complimentary nodes of cross-coupled inverters, and gates of two minimum-size transistors in the scan latch, plus some overhead due to internal wiring and increase in area, $C_{\text{scan}} = 2C_d + 2C_g + C_{\text{area incr}}$.

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 energy overhead of the proposed scan mechanism is

$$\triangle E_1 = \tfrac{1}{2}\,\alpha C_{\text{scan}} V_{dd}^2$$

where α is the "true" switching activity at the data input.

A prior art edge-triggered scannable version of the SA latch [18] is shown in Fig. 12(b). During the normal mode of operation the input signal **Scan** is low, and the SA current flows through transistors $\mathbf{N1}$ or $\mathbf{N2}$, controlled by the input data signals $\mathbf{D}$ and $\mathbf{Db}$. During the scan mode the signal $\mathbf{Scan}$ is high, and the SA current flows through transistors $\mathbf{N3}$ or $\mathbf{N4}$, controlled by the scan-in signals $\mathbf{I}$ and $\mathbf{Ib}$. This implementation of the scan-in capability has a very high-power overhead, because it significantly increases the capacitance at the bottom part of the latch (nodes $\mathbf{A}$, $\mathbf{B}$, $\mathbf{E}$, $\mathbf{F}$, and $\mathbf{M}$). Since these nodes are

TABLE IV
ENERGY, AREA, AND PERFORMANCE OVERHEAD OF SCANNABLE SA LATCH

design	area μ^2	energy overhead formula	C value	value for $V_{dd} = 1$v, $\alpha = 0.3, \beta = 0.08$	D overhead
prior art	30.72	$V_{dd}(V_{dd} - V_T)\triangle C_2 + \frac{1}{2}\alpha V_{dd}^2 C_{scan\ out}$	10.4 fF and 2*8.5fF	7.6fJ + 2.6fJ = 10.2fJ	28ns
mux-based	40.32	$\frac{1}{2}(\alpha + 2\beta)V_{dd}^2 C_{mux} + \frac{1}{2}\alpha V_{dd}^2 C_{scan\ out}$	6 fF and 9fF	1.4 fJ + 1.4 fJ = 2.8fJ	100ns
proposed	40.32	$\frac{1}{2}\alpha V_{dd}^2 C_{scan}$	4.1 fF	0.62 fJ	25ns

charged and discharged every clock cycle, independent of the switching activity, the increase in energy dissipation equals

$$\triangle E_2 = V_{dd}(V_{dd} - V_T)\triangle C_2 + \tfrac{1}{2}\,\alpha V_{dd}^2 C_{\text{scan out}}$$

where $\triangle C_2$ is the increase of the capacitance at nodes $\mathbf{A}$, $\mathbf{B}$, $\mathbf{E}$, $\mathbf{F}$ and $\mathbf{M}$ in Fig. 12(b), and $C_{scan\ out}$ is the sum of the gate capacitances of transistors $\mathbf{N3}$ and $\mathbf{N4}$ in the next latch of the scan chain, and the capacitance of the wire that connects latches in the scan chain, $C_{\text{scan out}} = 2C_g + 2C_{\text{scan wire}}$.

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

$$\triangle E_3 = \tfrac{1}{2}\,V_{dd}^2 C_{\text{mux}}(\alpha + 2\beta) + \tfrac{1}{2}\,\alpha V_{dd}^2 C_{\text{scan out}}$$

where C_{mux} is the capacitance of the multiplexor at the input, α is the input data switching activity, β is the glitching factor at the data input (one glitch represents two spurious transition), and $C_{\text{scan out}}$ is the sum of the input capacitance of the multiplexor in front of the next latch in the scan chain, and the capacitance of the wire that connects latches in the scan chain, $C_{\text{scan out}} = C_{\text{in mux}} + C_{\text{scan wire}}$.

To determine the exact energy overheads of the described scan mechanisms we developed layouts in a 0.13 μ technology for three versions of the scannable sense amplifier latch: scannable latch with the prior art scan mechanism [18] in Fig. 12(b), multiplexor-based design, described above, and scannable latch with the new scan mechanism, Fig. 12(a). The layouts for all latches were designed to fit into 12 tracks. The second column in Table IV shows the layout areas of these designs. Notice that the prior art design [18] has a 25% area advantage over the two other designs.

The next three columns in Table IV estimate the energy overhead of adding the scan feature to the SA latch using the three approaches. The fourth column estimates the capacitance values in the energy overhead formulas, based on wire capacitance for a 0.13 μm technology and simulations of the extracted netlists. The fifth column gives the energy overhead estimates for typical values of the data switching activity and the glitching factor, from Table I. Notice that $\beta = 0.08$ corresponds to the average of $2\beta = 0.16$ spurious transitions per cycle.

The results in Table IV show that the proposed scan mechanism reduces the energy overhead of the scannable latch 4.5 times, compared to the input multiplexed design, and 11 times, compared to the prior art design in Fig. 12(b). This advantage can be even more significant in high glitching nodes, or in clock-gated designs. Under the same conditions the full sense amplifier latch in Fig. 12(a) dissipates about 11 fJ per clock cycle. Thus, using the proposed approach results in more than 30% energy savings in the scannable SA latch.

In terms of the effect on the latch performance, the proposed scan mechanism has a 25 ps delay overhead in a latch with the lowest power level. This result is based on the simulations of the extracted netlist, and it includes the effect of increase in the area. The delay overhead is approximately the same as the delay overhead of the prior art approach in Fig. 12(b) which is 28 ps, and is significantly smaller than the setup time overhead of the multiplexor-based approach which is close to 100 ps for the same conditions. The latch insertion delay overheads are summarized in the last column of Table IV.

V. COMPARATIVE STUDY

We have done a comparative study of a number of different latch styles to identify the ones that are most suitable for low-power designs. Since the power-supply reduction is essential for reducing energy, we primarily focused on static and semistatic latches because of their higher noise margin. Also, since Vdd reduction plays such an important role, we were particularly interested in those latches whose performance degrades the least as Vdd is reduced [2].

In this paper, we show results only for four latch styles: LSSD scannable C^2MOS latch; LSSD transmission gate latch, Fig. 8 without the input data inverter which is needed to protect the transmission gate input, but can be omitted under certain conditions in custom datapaths; scannable sense amplifier latch, Fig. 12(a), and semistatic true single phase RAM latch [30] derived from [28]. The optimization described in Section II was applied to every latch: the sum of the setup time and delay was used as a measure of performance, D, and the expression (6) was used as a tuning cost function. The optimization parameter was changed from 0.1 to 0.9 to generate the *energy efficient* curve for each latch for $Vdd = 0.9$ V. A bulk technology was used with a 0.13 μm feature size. Then all *energy efficient* configurations of every latch were simulated for lower values of Vdd, $Vdd = 0.8$ V and $Vdd = 0.7$ V. No additional tuning was done however. The results are shown in Fig. 13, for the activity factor of 0.3 transitions per cycle, and the spurious activity of 0.15 glitches per cycle.

Extensive use of clock gating effectively increases the switching factor and the glitching activity. Fig. 14 plots the average energy per cycle of the same latches for higher value of the switching and glitching activities. Low-energy consumption of the scannable sense amplifier latch even in the presence of significant glitching activity, as well as its ability to operate

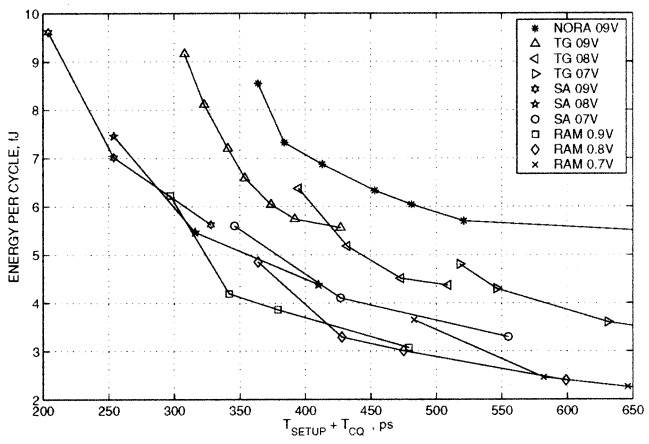

Fig. 13. Average energy per cycle versus performance. Switching factor $\alpha = 0.3$, glitching activity $\beta = 0.15$ ($\beta_1 = 0.1$, $\beta_2 = 0.02$, $\beta_3 = 0.005$). Solid lines connect points of *energy efficient* configurations for every value of Vdd.

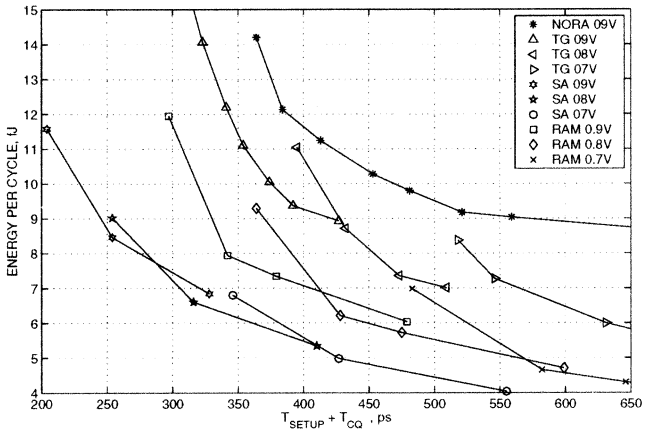

Fig. 14. Average energy per cycle versus performance. Switching factor $\alpha = 0.5$, glitching activity $\beta = 0.55$ ($\beta_1 = 0.2$, $\beta_2 = 0.1$, $\beta_3 = 0.05$). Solid lines connect points of *energy efficient* configurations for every value of Vdd.

with reduced swing signals make it a good candidate for low-power designs.

VI. Conclusions

The energy-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 used for formal comparison of different latch styles in the energy-performance space. Issues of tuning transistor sizes in latches were considered, along with methods for building energy-efficient families. Theoretical analysis for balancing hardware intensity was applied to the latch design, which indicated that the optimal value of hardware intensity parameter in latches is typically less than 2. Thus, it turns out that in many today's microprocessors latches are over-designed, compared to the rest of the hardware. Power breakdowns for some of commonly used latches was analyzed, and local clock-distribution power was found to be a significant component of the total power of the clocking system in low-power designs, which demonstrates the power advantage of latches that operate with a single phase of clock during the normal operation mode. Practical issues of building a low-energy overhead scan mechanism were considered, and a recently proposed low-power scan mechanism was analyzed and demonstrated to significantly reduce the energy overhead of scannable design. Results of a comparative study of scannable latches are shown, and the proposed nonsymmetric sense amplifier latch with the proposed scan extension was found to be a very strong candidate for low-power designs.

Acknowledgment

The author would like to thank D. Meltzer and S. Kosonocky for useful discussions; G. Gristede and A. Haen for the design flow support; and K. Warren and J. Moreno for the management support.

References

[1] T. Burd and R. Brodersen, "Energy efficient CMOS microprocessor design," in *Proc. 28th Annu. Hawaii Int. Conf. Syst Sciences*, 1995, pp. 288–297.

[2] H. Dao, K. Nowka, and V. Oklobdzija, "Analysis of clocked timing elements for dynamic voltage scaling effects over process parameter variation," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 2001, pp. 56–59.

[3] D. Dobberpuhl, "The design of a high performance low power microprocessor," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 1996, pp. 11–16.

[4] S. Faris, "Circuit design for full scan ATPG," in *Proc 4th Annu. IEEE Int. ASIC Conf.*, 1991, pp. 6.1–6.4.

[5] R. Gonzalez and M. Horowitz, "Energy dissipation in general purpose microprocessors," *IEEE J. Solid-State Circuits*, vol. 31, no. 9, pp. 1277–1283, Sept. 1996.

[6] S. Heo, R. Krashinsky, and K. Asanovic, "Activity-sensitive flip-flop and latch selection for reduced energy," in *Proc. 19th Conf. Advanced Research VLSI*, Mar. 2001, pp. 59–74.

[7] H. Kawahuchi and T. Sakurai, "A reduced clock-swing flip-flop (RCSFF) for 63% power reduction," *IEEE J. Solid-State Circuits*, vol. 33, pp. 807–811, May 1999.

[8] F. Klass *et al.*, "A new family of semidynamic and dynamic flop-flops with embedded logic for high-performance processors," *IEEE J.Solid-State Circuits*, vol. 34, pp. 712–716, May 1999.

[9] U. Ko and P. Balsara, "High-performance elergy-efficient D flip-flop circuits," *IEEE Trans. VLSI Syst.* , vol. 8, pp. 94–98, Feb. 2000.

[10] T. Lang, E. Musoll, and J. Cortadella, "Individual flip-flops with gated clocks for low power datapaths," *IEEE Trans Circuits Syst. II*, vol. 44, pp. 507–516, June 1997.

[11] J. Y. Lin, T. C. Liu, and W. Z. Shen, "A cell-based power estimation in CMOS combinational circuits," in *Proc. Int. Conf. Computer-Aided (ICCAD)*, 1994, pp. 304–309.

[12] R. Llopis and M. Sachdev, "Low power, testable dual edge triggered flip-flops," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 1996, pp. 329–334.

[13] A. Ma and K. Asanovic, "A double-pulsed set-conditional-reset flip-flop," MIT LCS, Tech. Rep. 844, May 2002.

[14] D. Markovic, B. Nikolic, and R. Brodersen, "Analysis and design of low-energy flip-flops," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 2001, pp. 52–55.

[15] N. Nedovic, M. Aleksic, and V. Oklobdzija, "Timing characterization of dual-edge triggered flip-flops," in *Proc. Int. Conf. Computer Design (ICCD)*, Aug. 2001, pp. 538–541.

[16] N. Nedovic and V. Oklobdzija, "Dynamic flip-flop with improved power," in *Proc. Int. Conf. Computer Design*, Sept. 2000, pp. 323–326.

[17] ——, "Hybrid latch flip-flop with improved power efficiency," in *Proc. Symp. Integrated Circuits and Systems Design*, 2000, pp. 211–215.

[18] B. Nikolic *et al.*, "Improved sense-amplifier-based flip-flop: Design and measurements," *IEEE J. Solid-State Circuits*, vol. 35, pp. 876–883, June 2000.

[19] H. Partovi *et al.*, "Flow-through latch and edge-triggered flip-flop hybrid elements," in *Proc. Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 1996, pp. 138–139.

[20] P. Penzes and A. Martin, "Energy-delay efficiency of VLSI computations," in *Proc. Great Lakes Symp. VLSI*, Apr. 2002, pp. 104–107.

[21] J. H. Satyanarayana and K. K. Parhi, "Heat: Hierarchical energy analysis tool," in *Proc. 33rd Design Automation Conf.*, June 1996, pp. 9–14.
[22] M. Stan, "Low-power CMOS with subvolt supply voltages," *IEEE Trans. VLSI Syst.*, vol. 9, pp. 394–400, Apr. 2001.
[23] V. Stojanovic and V. Oklobdzija, "Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems," *IEEE J. Solid-State Circuits*, vol. 34, pp. 536–548, Apr. 1999.
[24] V. Stojanovic, V. Oklobdzija, and R. Bajwa, "A unified approach in the analysis of latches and flip-flops for low-power systems," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 1998, pp. 227–232.
[25] A. G. M. Strollo, E. Napoli, and D. De Caro, "New clock-gating techniques for low-power flip-flops," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 2000, pp. 114–119.
[26] C. Svensson and J. Yuan, "Latches and flip-flops for low power systems," in *Low Power CMOS Design*, A. Chandrakasan and R. Brodersen, Eds. Piscataway, NJ: IEEE Press, 1998, pp. 233–238.
[27] J. Tschanz *et al.*, "Comparative delay and energy of single edge-triggered and dual edge-triggered pulsed flip-flops for high-performance microprocessors," in *IEEE Symp. Low-Power Electronics and Design*, Aug. 2001, pp. 147–152.
[28] J. Yuan and C. Svensson, "New single-clock CMOS latches and flipflops with improved speed and power savings," *IEEE J. Solid-State Circuits*, vol. 32, pp. 62–69, Jan. 1997.
[29] V. Zyuban and P. Kogge, "Application of STD to latch-power estimation," *IEEE Trans. VLSI Syst.*, vol. 7, pp. 111–115, Mar. 1999.
[30] V. Zyuban and D. Meltzer, "Clocking strategies and scannable latches for low power applications," in *Proc. IEEE Symp. Low-Power Electronics and Design*, Aug. 2001, pp. 346–351.
[31] V. Zyuban and P. Strenski, "Unified methodology for resolving power-performance tradeoffs at the microarchitectural and circuit levels," in *Proc. Int. Symp. Low-Power Electronics and Design*, Aug. 2002, pp. 166–171.

Victor Zyuban (S'97–M'00) received the B.S. and M.S. degrees from the Moscow Institute of Physics and Technology, Russia, in 1993 and 1995, respectively, and the Ph.D. degree in computer science and engineering from the University of Notre Dame, IN, in 2000.

From 1995 to 1996, he was with the Moscow Center for SPARC Technologies, Russia. Currently, he is a Research Staff Member with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, where from 2002 to 2003, he worked on a low-power digital signal processing (DSP) research project in which he was involved in ISA definition, microarchitecture, and physical design, and lead the development of a semicustom eLite core test chip; and is now working on a high-performance general-purpose microprocessor core. His research interests include high-frequency, low-power circuitry, microarchitecture and methodologies for energy-efficient design.