

Boost Logic : A High Speed Energy Recovery Circuit Family

Visvesh S. Sathe, Conrad H. Ziesler and Marios C. Papaefthymiou

Electrical Engineering and Computer Science

University of Michigan

Ann Arbor, MI 48109, USA

email{vssathe,cziesler,marios}@eecs.umich.edu

Abstract

In this paper, we propose Boost Logic, a logic family which relies on voltage scaling, gate overdrive and energy recovery techniques to achieve high energy efficiency at frequencies in the GHz range. The key feature of our design is the use of an energy recovering “boost” stage to provide an efficient gate overdrive to a highly voltage scaled logic at near threshold supply voltage. We have evaluated our logic family using post-layout simulation of an 8-bit pipelined array multiplier in a $0.13\mu\text{m}$ CMOS process with $V_{th}=340\text{mV}$. At 1.6GHz and a 1.3V supply voltage, the Boost multiplier dissipates 8.11pJ per computation. Comparing results from post-layout simulations of boost and voltage-scaled conventional multipliers, our proposed logic achieves 68% energy savings with respect to static CMOS. Using low V_{th} devices, Boost Logic has been verified to operate at 2GHz with a 1.25V voltage supply and 8.5pJ energy dissipation per cycle.

1 Introduction

Power minimization has become one of the primary concerns in VLSI design. Several conventional techniques are utilized to curb dynamic and leakage power in conventional CMOS circuits. One of the most effective methods is pipelining and subsequent voltage scaling to minimize energy at a given operating frequency. At high frequencies of operation, however, the energy and delay overhead of pipeline registers becomes significant and results in a degradation of system efficiency.

Energy recovery circuits offer an alternative approach to the reduction of dynamic energy dissipation. Several energy recovery logic styles have been proposed [1, 2, 3, 4, 5]. Over a range of relatively low operating frequencies (a few hundred megahertz), these energy recovery techniques have been shown to achieve the same performance at lower energy dissipation when compared to voltage scaled CMOS. Achieving energy savings over CMOS at higher operating frequencies has remained elusive, however.

Although performance limits of energy recovery circuits are fundamentally determined by the need for gradually transitioning power clocks, prevalent operating frequencies in energy recovery circuits are more a consequence of design than any such fundamental constraint. Some of the main factors that lead to lower speeds in energy recovery circuits are the use of diode-connected transistors [6, 7], the use of pMOS devices in evaluation trees [8, 9], and the excessive time required to resolve the complementary outputs of the dual-rail gates during evaluation [2, 4].

In this paper, we propose a novel dynamic logic family called Boost Logic. This family is a fine-grained, two-phase hybrid logic that consists of conventional switching and energy recovery stages and can achieve significant energy savings over voltage-scaled CMOS across a range of frequencies much higher than currently demonstrated in energy recovery literature. A unique feature of Boost Logic gates that enables high throughput operation is the “boost” stage at the output of the gate. The boost stage serves to provide a greater gate overdrive for the evaluation trees of fanout gates, thereby reducing the delay in the aggressively voltage-scaled logic evaluation stage. Thus, the boost stage achieves lower energy dissipation without incurring the same performance degradation experienced in conventional voltage-scaled designs.

Figure 1(a) illustrates the concept behind Boost Logic. Each Boost Logic gate consists of 2 parts: A conventionally-switching logical evaluation stage “Logic” and a charge recovering “Boost” stage. Boost Logic employs a convention-

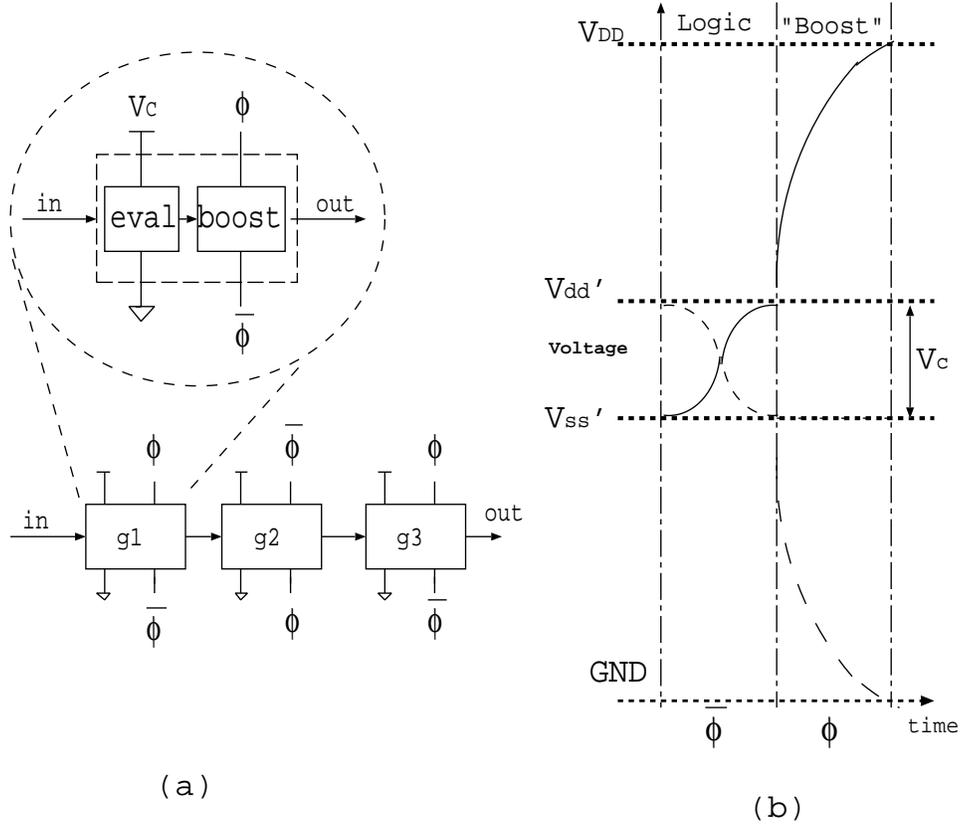


Figure 1: Boost Logic (a) Cascade and (b) Operation

ally switching logic stage which provides it with greater voltage scalability as compared to fully energy recovering circuits. This conventional logic operates at an ultra-low DC voltage supply. An efficient amplifying stage (“Boost”) is then applied at the output of the logic stage to boost the voltage level of the logic “1” node from V_{dd}' to the nominal voltage V_{dd} and from V_{ss}' to GND as shown in Figure 1(b). The value of V_c is approximately V_{th} . The logic and boost stages of a Boost Logic gate operate on complementary phases of the clock.

In Boost Logic, both dynamic and leakage power in the evaluation stage are greatly reduced as a result of the low supply voltage. Despite this scaled voltage, the evaluate stage is able to function in the gigahertz range due to the gate overdrive of $V_g' = (V_{dd} - V_{th})/2$ provided to its n-type evaluation trees by the boost stages of its fanin gates.

The idea of providing greater gate overdrive has been previously proposed in [3, 10], where bootstrapping was used to that end. Such techniques, however, lack the robustness offered by the boost stage and are limited in the amount of gate overdrive that can be achieved. More recently, LVS logic [11] has been proposed, where sense amplifiers operating at a higher supply voltage are used to provide gate overdrive.

The dynamic energy consumed by a Boost Logic gate with a voltage supply of V_c for one transition is:

$$E = \frac{1}{2} \cdot CV_c^2 + E_{boost}, \quad (1)$$

where E_{boost} is the energy dissipated in the boost stage, C is the switching capacitance and V_c is the voltage swing of the capacitance. Although the boost stage provides significant advantages by reducing the energy dissipated in its logic stage and increasing its speed, it is vital that the power dissipation of the boost converter itself does not nullify these advantages. By using an efficient high-speed energy recovering circuit to perform the operation of the boost stage, the latter is implemented with a low energy overhead.

We have performed several simulation experiments to verify and characterize the performance and energy dissipation of Boost Logic. Since Boost Logic gates are driven by complementary power-clocks, we also characterized the

robustness of standard Boost Logic gates to clock skew.

An 8-bit Boost array multiplier with BIST was designed in an industrial $0.13\mu\text{m}$ process. To compare the performance of Boost Logic with other design styles, we also implemented a pipelined, voltage-scaled CMOS array multiplier. Industrial synthesis and place and route tools were used to design a static CMOS multiplier, pipelined and voltage scaled so as to achieve minimum energy dissipation at 1.6GHz. Energy comparisons between the two multipliers were made at the frequency of 1.6GHz. All simulations were performed on post-layout designs with extracted parasitics. In simulations, Boost Logic achieved energy savings of 68% over its pipelined static counterpart.

Boost Logic performance is enhanced considerably with the use of low V_{th} devices in the logic stage. The use of these devices provides more slack for the logic evaluation stage by improving the transistor drive strength. Given the low supply voltage that the logic stage operates under, leakage power resulting from the sub-threshold leakage component in the logic stage is insignificant. Using low V_{th} devices offers an additional advantage of extending the time allotted for logical evaluation in each cycle. In simulations of our Boost multiplier with device threshold voltages of $V_{th}=200\text{mV}$, a further 29% energy savings was achieved over the nominal V_{th} boost design at 1.6GHz.

The remainder of the paper is organized as follows: In Section 2, we present Boost Logic and discuss its structure. We also discuss the efficiency of the boost stage which plays a pivotal role in the efficient operation of Boost Logic. Results obtained from numerous simulations such as energy-performance characteristics of Boost gates and the benefit derived from low V_{th} design are discussed in Section 3. In that section we also present the 8-bit carry-save array multiplier and compare its energy and throughput to a voltage-scaled pipelined CMOS implementation. Conclusions are given in Section 4.

2 Energy Recovering Boost Logic

In this section, we first analyze the structure and operation of Boost Logic. We subsequently consider the energy and delay equations that apply to Boost Logic and show how Boost Logic achieves high throughput with significant energy savings.

2.1 Boost Logic structure

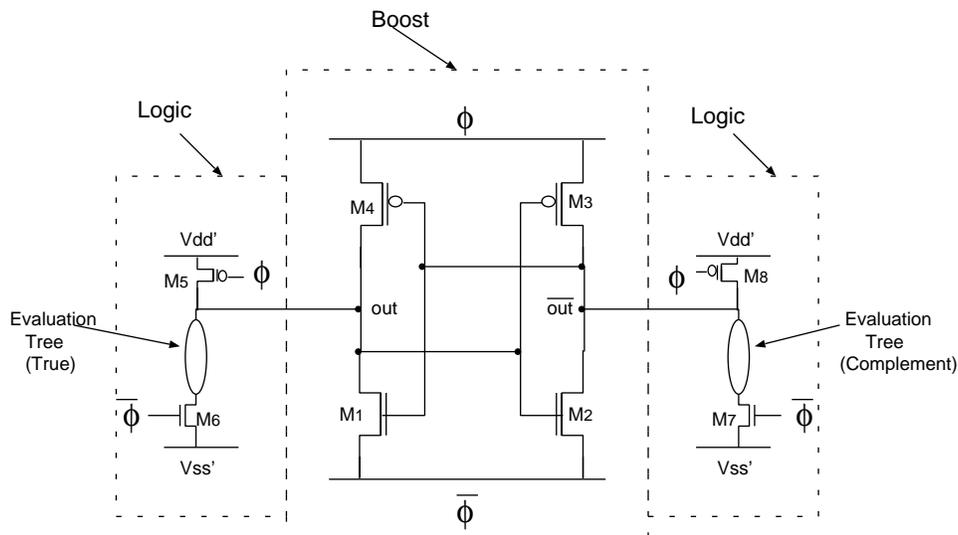


Figure 2: Boost Logic

Figure 2 shows a typical Boost Logic logic gate. Boost Logic is a two-phase, dual-rail, partially energy recovering logic. The operation of a Boost gate can be divided into two parts—logical evaluation (“Logic”) and boost conversion

(“Boost”). The logic stage comprises a dual-rail pseudo nMOS evaluation tree. The design of the logic stage differs from conventional pseudo nMOS evaluation in that the weak pMOS pull-up and the footer transistor both turn on only during the evaluation of the logic stage. At other times, they are off, isolating the output node from the conventional voltage supply rails. The pseudo nMOS-like gate is chosen to reduce the loading on the gate thereby improving performance. To improve the robustness of the design, a clock-gated CMOS logic stage can be used instead of the pseudo nMOS evaluation tree. The power supply rails are at voltages:

$$V_{dd}' = \frac{1}{2} \cdot (V_{dd} + V_{th}), \quad (2)$$

$$V_{ss}' = \frac{1}{2} \cdot (V_{dd} - V_{th}). \quad (3)$$

This choice of voltage values is motivated by the operation of the boost stage which will be discussed in greater detail in Subsection 2.2. The potential difference between the voltage supply rails in the logic stage is therefore $V_c = V_{th}$. The boost stage, which is essentially an energy recovering sense amplifier, resembles back-to-back CMOS inverters. The only difference is that the V_{dd} and Gnd rails are replaced by ϕ and $\bar{\phi}$.

Boost Logic utilizes a dual-rail gate structure to ensure that the capacitance presented to the power-clock by the gate is balanced and data-independent, reducing clock jitter. The use of the pseudo nMOS-type evaluation tree reduces the input loading of the gate at the expense of short-circuit dissipation in the gate. The delay penalty due to the header and footer can be reduced by sizing up transistors M_5 , M_6 , M_7 , and M_8 . Since the gate inputs to these transistors are resonant clocks, wider transistors result in significantly lower energy penalties compared to a conventional clock. To reduce the susceptibility of gate performance to process variation, a complementary pMOS evaluation tree can be used in series with M_5 and M_8 .

2.2 Operation

Figure 3 illustrates the operation of a Boost inverter. The complementary clock waveform $\bar{\phi}$ is not shown in the figure but is exactly in anti-phase with ϕ . By design, the logic and boost stages evaluate at mutually exclusive intervals. As such, when the logic stage evaluates, the boost stage does not drive the outputs and vice-versa. Consider the operation of the gate whose waveforms are shown in Figure 3. When the logic stage evaluates (ϕ falls and $\bar{\phi}$ rises), the header transistors M_5 and M_8 and footer transistors M_6 and M_7 turn on. As out evaluates high, the header transistor M_5 pulls the output node to V_{dd}' . The complementary output discharges through the evaluation tree to nearly V_{ss}' . At this time, the energy recovering sense amplifier is in pre-charge with $\phi = 0$ and $\bar{\phi} = V_{dd}$. In this state, it is easily verified that as long as the outputs stay within the conventional supply rails, none of the transistors in the sense amplifier are turned on, and no crowbar current flows in the Boost converter. As ϕ begins to rise past V_{ss}' (or 750mV in Figure 3), the logic stage is deactivated, disconnecting from V_{dd}' and V_{ss}' . As ϕ continues to rise past V_{dd}' , the boost conversion begins to operate. Since out is at V_{dd}' and \overline{out} at nearly V_{ss}' , transistors M_2 and M_4 turn on as $\bar{\phi}$ (ϕ) goes past V_{ss}' (V_{dd}'), causing \overline{out} (out) to subsequently follow $\bar{\phi}$ (ϕ). During boost conversion, as the voltage difference between out and \overline{out} increases, transistors M_2 and M_4 turn more strongly on, reducing the voltage difference across the current-carrying transistors further. Finally, the nodes out and \overline{out} reach the rails ϕ and $\bar{\phi}$, respectively. These outputs will drive the next gate during its logical evaluation stage.

As ϕ and $\bar{\phi}$ transition once again, entering the next logic phase, the outputs track the corresponding complementary clocks once again through the same transistors M_2 and M_4 . As the voltage difference between out and \overline{out} approaches V_{th} , conduction in any of the four transistors in the sense amplifier stops and the logic stage once again begins to evaluate.

Figure 3 shows Boost Logic operating with sinusoidal power-clocks. While sinusoidal power-clocks are natural to resonant clock generation [9, 12], Boost Logic also operates correctly with trapezoidal clocks such as the resonant clock generated by the rotary clock [13].

Boost Logic achieves energy recovery at high frequencies due to several design features. First, the boost converter stage in Boost Logic does not require diodes to perform energy recovery and can therefore operate efficiently at relatively higher frequencies. Being an n-n logic, Boost Logic eliminates the use of pMOS evaluation trees, greatly reducing capacitive loading of gate inputs in spite of being a dual-rail logic and enhancing speed. Also, Boost gates pre-charge to nearly $1/2V_{dd}$ which reduces the output swing of the gate and therefore the energy dissipated in the boost

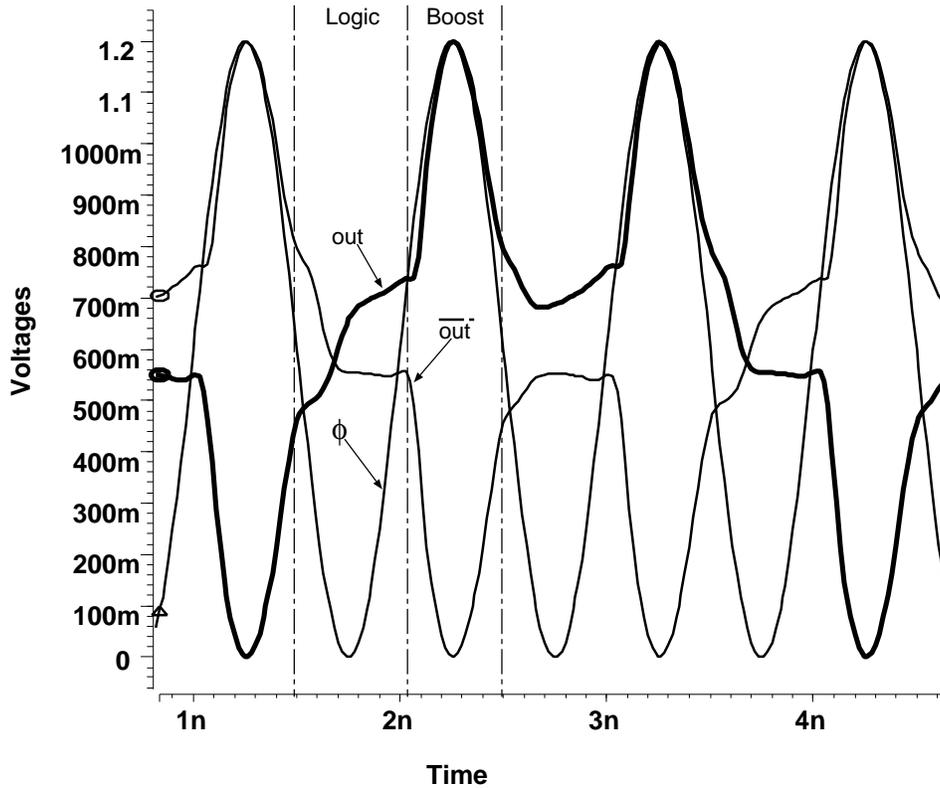


Figure 3: SPICE waveforms of a Boost Logic inverter

stage. By not having to follow the power-clock when it transitions at its fastest rate ($1/2V_{dd}$ for sinusoidal clocks), higher operating frequencies are possible for a given energy efficiency. This form of pre-charge also provides more time for the logic stage of the gate to evaluate as compared to energy recovery designs that pre-charge to nearly V_{dd} or Gnd .

Another feature of Boost Logic that enables its high frequency operation is the fact that the pseudo nMOS structure in the logic stage produces complementary output nodes with a voltage difference of nearly V_c . Thus, the gate outputs are not left unresolved at the onset of boost conversion precluding any “fight” between the output nodes of the energy recovering sense amplifier and resulting in efficient boost conversion. The absence of any conflict in the sense amplifier during the operation of the Boost stage also provides a data-independent capacitance to the clock generator, minimizing data-dependent jitter.

The intermediate voltage rails for the logic stage of the gate offer a body-biasing advantage to Boost Logic. Substrate contacts for all nMOS devices are made to V'_{ss} and the well contacts for the pMOS devices are made to V'_{dd} , providing a forward body bias to the boost converter transistors and improving energy recovery and fanout capability. At the same time, such body contacts ensure that the performance of the logic stage transistors is not degraded due to the body effect.

The transistor count of Boost gates is $2n + 8$ where n is the number of logical inputs. This transistor count presents a relatively low area overhead, since each Boost gate typically performs a complex logical operation (2 gates form a full adder, for example), amortizing the overhead of extra transistors. Furthermore, the evaluation tree is made up only of nMOS transistors, reducing the gate area considerably. Finally, Boost Logic is a dynamic logic family and does not require the use of pipeline registers to achieve high throughput.

Cascading Boost gates is straightforward. Since the boost conversion of a gate is required to occur concurrently with the logic evaluation stage in its fan-out gates, gates are cascaded by driving the boost stages of subsequent gates with alternating clock phases ϕ and $\bar{\phi}$, as shown in Figure 1. The connection for a chain of inverters is shown

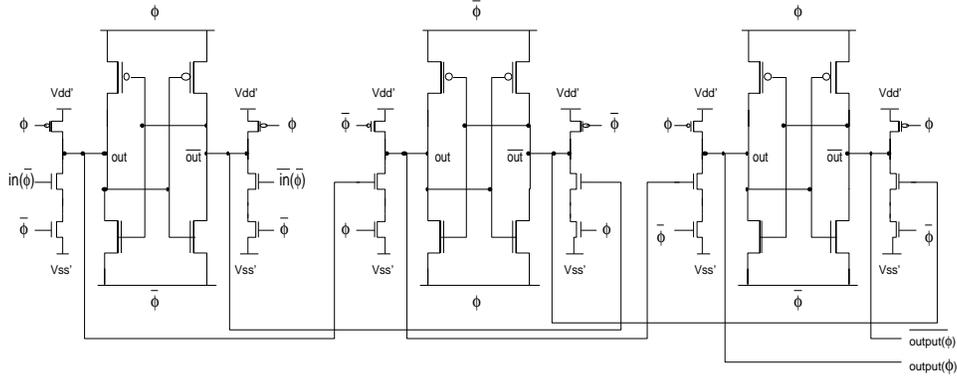


Figure 4: Cascade of Boost Logic inverters

in Figure 4. It is useful to observe that from a timing (and to a large extent, functional) perspective, a boost gate consists of a conventional gate driving a level-converting latch. As in latch-based design, Boost Logic is cascaded with alternating ϕ and $\bar{\phi}$ gates.

2.3 Energy and delay

In this section we consider the equations that govern the energy dissipation of Boost Logic and the delay through the logic stage of the gate. We also highlight the low delay variation of a Boost gate upon scaling V_c .

Given that the transistors in the evaluation tree operate in the linear mode, the delay δ in the logic stage of the gate can be approximated as:

$$\delta \propto \frac{C \cdot V_c}{\left[\left(\frac{V_{dd}}{2} + \frac{V_c}{2} - V_{th} \right) V_c - \frac{1}{2} V_c^2 \right]}, \quad (4)$$

where V_c is the voltage swing of the gate, and V_{dd} is the amplitude of the power-clock. This expression simplifies to:

$$\delta \propto \frac{C}{\frac{V_{dd}}{2} - V_{th}}. \quad (5)$$

Considering first-order transistor effects, this result implies that unlike CMOS, the delay of the logic stage of the gate does not depend on its supply voltage. This delay insensitivity to the conventional power supply is due to the fact that while an increase in V_c increases the current drive of the gate, the required voltage swing also increases. Since the transistors in the logic stage operate largely in the linear mode, the delay trade-off resulting from voltage swing and current drive is balanced. Thus, the supply voltage of the logic stage can be reduced so as to decrease the energy consumption in the gate to a certain extent. Indeed, the extent to which this beneficial energy-delay correlation can be exploited is limited by noise susceptibility considerations and boost conversion efficiency.

The effect of channel length variations and the associated variations in threshold voltages on Boost Logic performance is an important practical consideration. Although Boost Logic uses a near-threshold power supply to power its logic stage, it is important to note that the transistors in this logic stage do not perform logical evaluation in the sub-threshold regime. Instead, they operate in the linear mode, where the sensitivity of gate delay to V_{th} is comparable to its voltage-scaled CMOS counterpart.

The boost converter is implemented in energy recovery logic. Therefore, the energy dissipation of the boost stage can be shown to be approximately:

$$E_{boost} \approx \frac{\pi^2 \tau}{8T} C_{boost} V_{dd}^2, \quad (6)$$

where $\tau = RC_{boost}$ is the product of the resistance in the boost stage looking into a power clock terminal and the total capacitance of the gate. V_{dd} is the amplitude of the power clock and T is the clock period of the clock. Since $V_c = V_{th}$

by design, Equation (1) can be rewritten as:

$$E \approx \frac{3}{4} \cdot C_{logic} V_{th}^2 + \frac{\pi^2 \tau}{8T} C_{boost} V_{dd}^2. \quad (7)$$

Equation (7) is a good approximation of the actual energy dissipation in the Boost gate, because the boost stage output follows the power-clock closely and does not contain any additional energy dissipation terms due to diode drops in the gate. The scaling factor of 3/4 for the dissipation of the logic stage is higher than the expected value of 1/2 due to the crowbar current that flows in the pseudo nMOS logic when the output is evaluated low. If a complementary pull up tree was employed instead, the scaling fraction would have been 1/2. Nevertheless, the energy dissipation in the logic stage remains proportional to V_{th}^2 (unlike several low output swing logic families in which the energy dissipation varies linearly with the swing voltage), since the charge in the logic stage is actually provided by a supply with potential difference V_{th} . Although the term E_{boost} contains the factor V_{dd}^2 which is much greater than V_c^2 , the scaling factor $\pi^2 \cdot \tau / (8T)$ is significantly smaller than 1/2, even at operating frequencies of 1GHz. While Equation (7) assumes a clock amplitude of V_{dd} , this amplitude can be reduced for more efficient operation at lower frequencies, as will be seen in Section 3.4.

3 Simulation results

In this section, we present various performance and energy characteristics of Boost Logic. In Section 3.1 we investigate the robustness of Boost Logic to clock skew. In Section 3.2, we consider the delay variation in Boost Logic as a result of power supply variation. Monte Carlo simulation results performed on Boost Logic to investigate its sensitivity to process variation are presented in Section 3.3. In Section 3.4, we present the simulation results obtained from the 8-bit energy recovery multiplier along with Built-in Self Test designed entirely in Boost Logic. We also compare the energy consumption of the Boost Logic multiplier with pipelined, voltage-scaled CMOS implementations of the same multiplier with post-layout simulations.

3.1 Robustness to clock skew

Boost gates depend on the power-clock for driving the boost converter of the gate as well as providing timing information for the correct operation of the gate. Robustness to clock skew is therefore a strict requirement for fine-grained energy recovery logic. It should be noted that the balanced, dual-rail design of Boost Logic ensures that the clock tree always drives nearly the same load regardless of its state, thus reducing the time-varying skew that can exist in the power clock. Nevertheless, robustness to clock skew is necessary because of the absence of buffers in the clock tree to re-power the clock and control skew.

In a cascade of gates, the phase difference between the power-clock driving a gate and the power-clock driving its fan-out gate can affect the energy efficiency and functionality of the energy recovery gate. We refer to this kind of clock skew as external clock skew. Since Boost Logic requires two clock phases, 180° out of phase, to perform any computation, another kind of skew is possible wherein there exists a phase difference between ϕ and $\bar{\phi}$ for a given gate. We refer to such a phase difference between ϕ and $\bar{\phi}$ as internal skew.

To determine the robustness of Boost gates to both kinds of skew, we evaluated a parallel arrangement of basic Boost gates. Providing random inputs to the gates, we verified functional correctness in each gate while varying the amounts of both types of clock skew. The clock signals used in the experiments were forced signals. Random input vectors were generated by a Linear Feedback Shift Register (LFSR) which we designed in Boost Logic. These vectors served as inputs to the parallel arrangement of gates that we designed in Boost Logic. The test gates were AND, OR, XOR, INV and AOI. A FO4 load was applied to the output of each test gate. A functional error in any of the gates would be detected in the signature of the signature analyzer. Simulations were carried out over the range of different internal skew and external skew values from -45% to $+45\%$ of the clock period.

Figure 5 shows the schmoop plot obtained. The points marked '+' indicate that all gates operated correctly at the corresponding values of internal and external skew. The skew values are given as a percentage of the cycle time. It can be inferred from Figure 5 that Boost Logic operates correctly over a large range of possible conditions of internal and external skew. In particular, all Boost Logic gates simulated correctly under simultaneous internal and external skew, each amounting to 15% of the clock cycle.

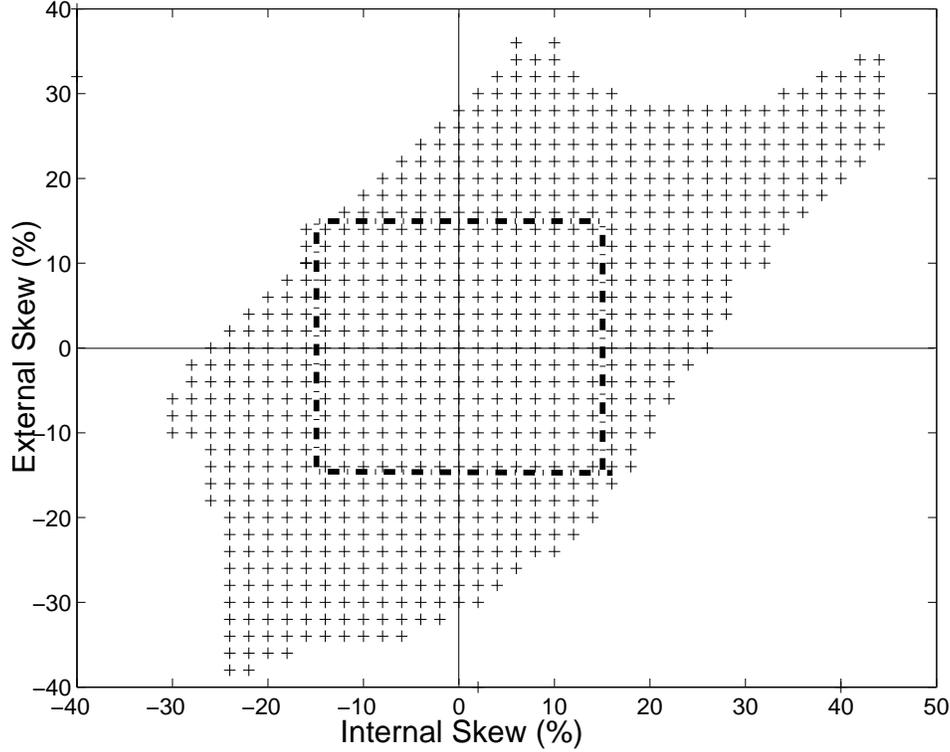


Figure 5: Schmoop plot for functional correctness over a range of internal and external skew values

3.2 Power Supply Variation

The sensitivity of Boost Logic to power supply variation is an important property from an operational standpoint. Boost Logic is powered by two supplies: The power-clock and the ultra low DC supply voltage. Voltage fluctuation in either power supply affects the performance of Boost Logic. From Equation (5), it can be inferred that the delay in the logic stage of Boost Logic is independent of V_c and inversely proportional to the power-clock amplitude. The sensitivity of Boost Logic to the power-clock amplitude is comparable to that in a voltage-scaled CMOS circuit which varies as $V_{dd}/(V_{dd} - V_{th})^\alpha$ ($1 < \alpha < 2$). However, the somewhat counter-intuitive delay independence of Boost Logic to V_c as predicted from first-order transistor behavior needs further verification. As such, we evaluated the sensitivity of a Boost NAND gate to variations in V_c . The load driven by the NAND gate was another identical gate.

Figure 6 illustrates the effect of power supply variation on the delay of Boost Logic and CMOS at supply voltages of 1.2V and 0.8V. Figure 7 illustrates the effect of power supply variation on the energy dissipation of Boost Logic and CMOS at supply voltages of 1.2V and 0.8V. In this experiment, the conventional supply of the CMOS and Boost NAND gates was varied over a range of $\pm 30\%$ and the percentage change in delay was measured. The results indicate that Boost Logic delay and energy dissipation vary in the range $[-13\%, +12\%]$ and $[-10\%, +30\%]$, respectively, for the reported variation in power supply. This variation in delay and energy dissipation is significantly lower than that observed in CMOS even though the Boost Logic power supply operates at V_{th} .

3.3 Process Variation

To investigate the robustness of Boost logic to process variation, we performed Monte Carlo simulations on a sample circuit containing NAND boost gates. A similar experiment was conducted for CMOS NAND gates. The channel length of the FETs in the NAND gate was assumed to be a normally distributed random variable, with a standard deviation of 5% of the mean channel length. It was assumed that channel length variations within a gate are negligible in a $0.13\mu\text{m}$ process.

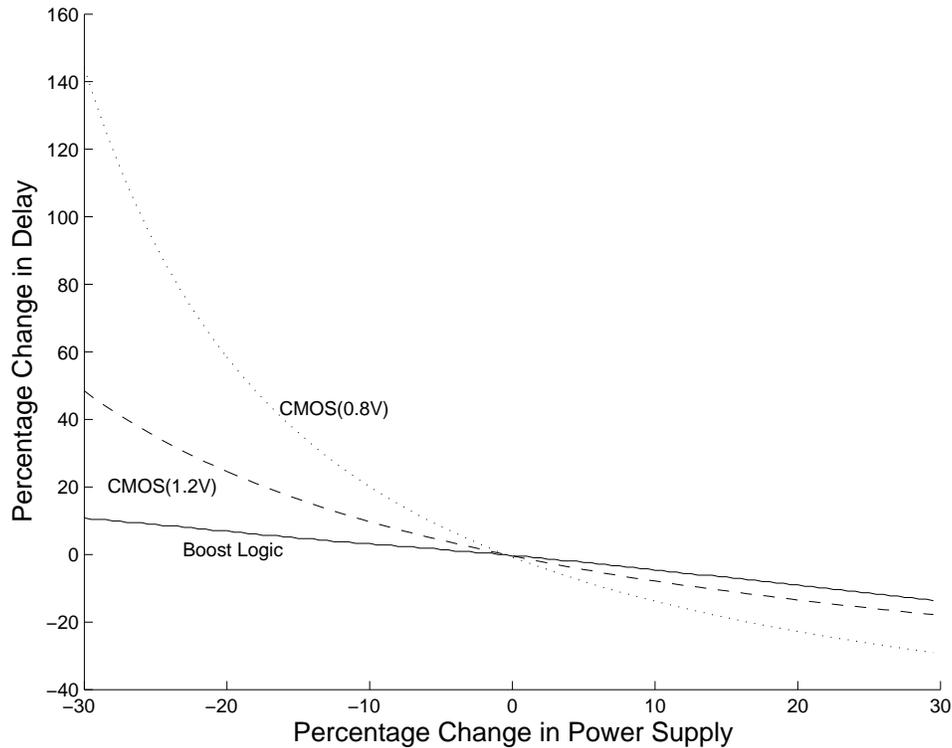


Figure 6: Effect of power supply variation from nominal values on delay in Boost and CMOS NAND gates

From Monte Carlo simulations, the 3σ values of the Boost and CMOS logic delays were found to be 3.15% and 13.7% of their respective mean values. While the sensitivity of delay to channel length variation seem to be lower for Boost in comparison to CMOS, it must be noted that the impact of channel length on the delay of CMOS logic depends not on one gate alone but on the variation along an entire path. The variation in the delay of a collection of gates is expected to be lower than that of a single gate. Consequently, the effect of channel length variation on the cycle time of a conventional CMOS logic circuit strongly depends on the number of gates in the stage in question and can potentially be lower than implied by the simulation results.

The 3σ values of the resulting distribution of the energy dissipated by the Boost and CMOS gates was found to be 13.75% and 2.33% of their mean values respectively. Predictably, the sensitivity of energy dissipation to channel length variation is greater for Boost compared to CMOS for two main reasons. First, the energy dissipated due to leakage (a major cause of variation in energy dissipation in CMOS) accounted for a small fraction of energy consumed in the simulation. Second, for the given boost design, with its pseudo nMOS evaluation stage, the amount of crowbar current (and therefore total energy dissipation) depends greatly on the channel length of the transistors in the boost stage. The sensitivity of Boost Logic energy dissipation to channel length variation can be greatly reduced by introducing complementary pMOS transistors in the logic evaluation stage.

3.4 8-bit Boost Logic array multiplier

We have designed an 8-bit carry-save array multiplier suited for use in FIR filters which are not latency critical. The accompanying BIST logic was also entirely designed in Boost Logic. As shown in Figure 8, an LFSR was used to provide pseudo-random input vectors to the multiplier. Outputs to the multiplier were processed by a signature analyzer. The power-clock signals were derived using an H-bridge clock generator. Pulses a and b were used to control switches in order to replenish the energy in the clock generator. Being a periodically driven oscillator, no special start-up circuitry was required, and stable oscillations in the multiplier were observed within 2 cycles. In the

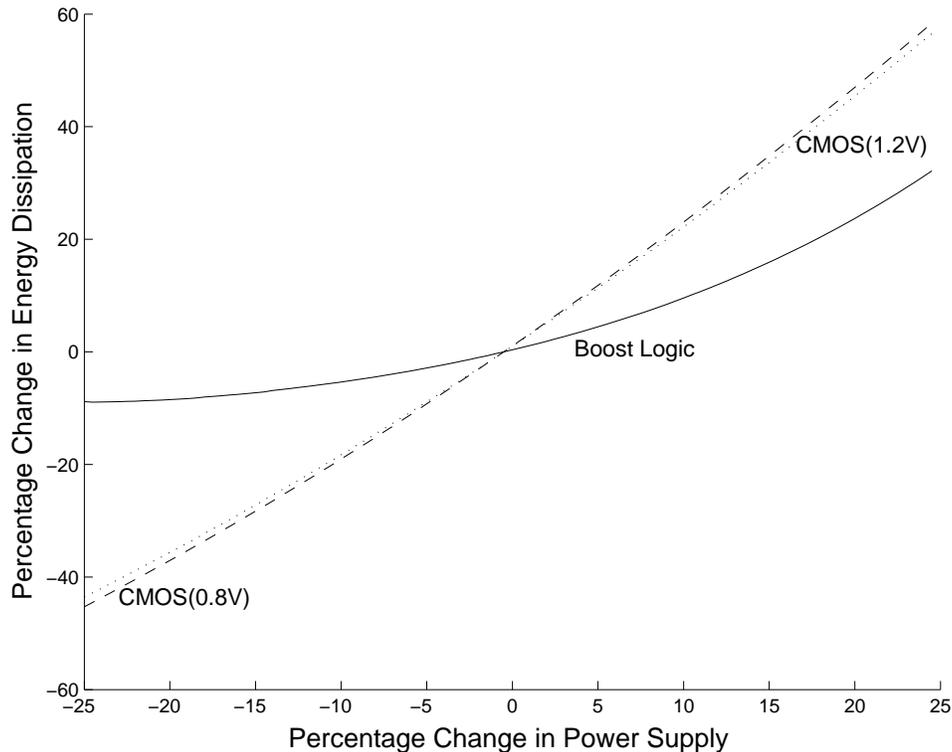


Figure 7: Effect of power supply variation from nominal values on energy dissipation in Boost and CMOS NAND gates

experimental setup, the capacitance driven by the clock generator (including the parasitic capacitance of the inductor) was approximately 20pF per phase. The value of the inductor used in the circuit depended on the operating frequency. Functional performance was verified by recording the signature output of the analyzer at a predetermined time and comparing it to the expected signature. To verify the significance of lower V_{th} devices, we also designed an identical multiplier using low V_{th} devices.

In Section 3.4.1, we consider the effects of power clock voltage-scaling on the energy-delay relationship in a Boost multiplier. In Section 3.4.2, we compare the effect of using low V_{th} devices on the Boost multiplier over a range of frequencies up to 2GHz. Finally, in Section 3.4.3, we compare the energy dissipation between the Boost multiplier and the voltage-scaled pipelined CMOS multiplier at 1.6GHz. From post-layout simulations, the energy dissipation of the multiplier with BIST and the clock generator was found to be 8.11pJ per computation at 1.6GHz. Simulations of the CMOS multiplier and Boost multiplier resulted in 68% energy savings of the Boost multiplier over its CMOS counterpart.

3.4.1 Boost Logic: Voltage Scaling

In this subsection, we consider the conflicting trends between the Boost Converter efficiency and the Logic Stage energy with respect to operating frequency at a given clock voltage amplitude. From an energy perspective, this trade-off results in an optimal operating frequency for a given clock amplitude. Furthermore, this optimal frequency varies with the power-clock amplitude. We compare the energy dissipation at this optimal frequency to the energy dissipation in a multiplier operating at a lower clock amplitude at the same frequency.

Figure 9 illustrates the Energy-Delay relationship for a given Boost multiplier at different operating voltages using normal V_{th} devices. The lowest possible energy dissipation for each frequency forms the energy delay curve for the multiplier. Delay is defined as the time period of the power-clock.

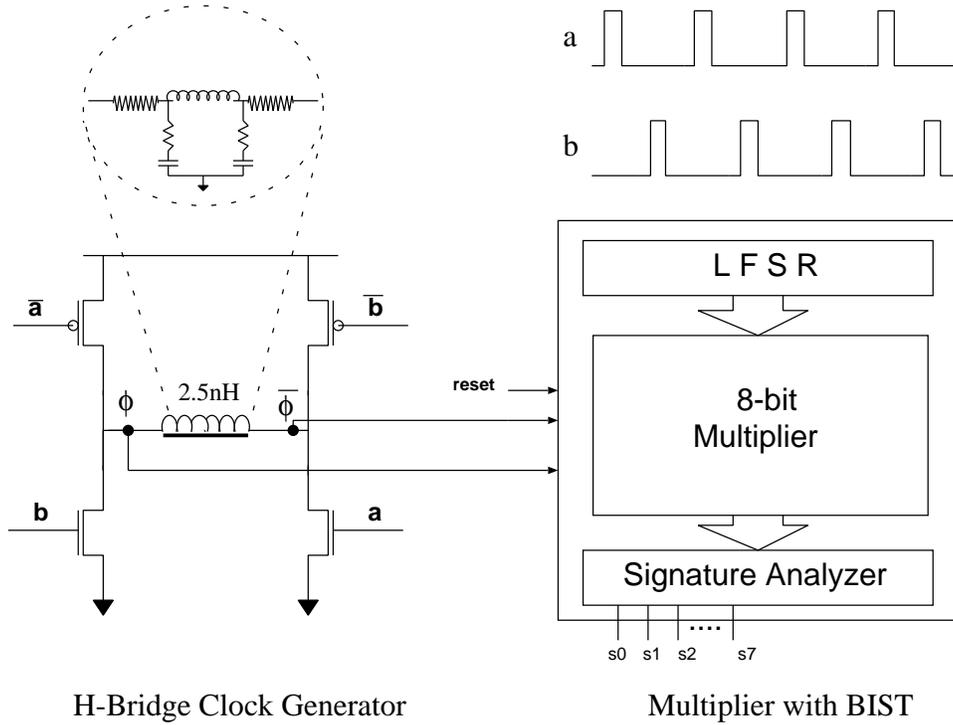


Figure 8: Overall simulation setup

For large clock periods, the energy consumption of Boost gates per computation increases with increasing time period due to the "crowbar" current that flows in the logic stage of the Boost gate when the gate evaluates low. For systems designed to operate at such lower frequencies, the logic stage of the boost gate should employ complementary evaluation trees as opposed to a pseudo nMOS logic. As the clock period decreases, the energy wasted by the crowbar current after logical evaluation is reduced, thus reducing energy dissipation with decreasing time periods of operation. A further reduction in the time period beyond a certain value degrades the energy efficiency of the Boost stage, which relies on a gradually-slewing power-clock. Consequently, the circuit consumes more power per computation. At lower clock supply voltages, the multiplier was observed to fail before the energy penalty due to an inefficient Boost stage became dominant. Nevertheless, it is apparent that the energy benefits derived from operating at a lower clock amplitude are greater than the energy penalty arising from operating the circuit at a "sub-optimal" frequency.

Therefore, the operating frequency resulting in minimum energy dissipation at a given supply voltage is not necessarily the lowest energy dissipation achievable at that frequency. It is observed that minimum energy is achieved for a given design by operating at the lowest possible clock amplitude.

3.4.2 Low V_{th} design

By using low V_{th} devices in the design of Boost gates, it is possible to greatly improve their performance and energy dissipation. Not only do low V_{th} transistors enable faster evaluation in the logic stage of the Boost gate, but they also increase the window of time for which header and footer devices remain on, allowing more time for logical evaluation and providing an opportunity for higher throughput or lower latency of computation. To illustrate the impact of using low V_{th} devices, we designed an 8-bit Boost multiplier using $V_{thl} = 200mV$ offered by the process.

Post-layout simulation was performed over a range of operating frequencies. Once again, voltage scaling was used to reduce the energy dissipation of the multipliers for lower frequencies. In the energy delay curves shown in Figure 10, it is observed that at 1.6GHz, the low V_{th} design obtained energy savings of 29% over its nominal V_{th} counterpart. Furthermore, the use of low V_{th} transistors allows the Boost multiplier to operate at frequencies beyond 2GHz. The operation of Boost logic is possible even with zero V_{th} devices in the logic stage, since the transistors in

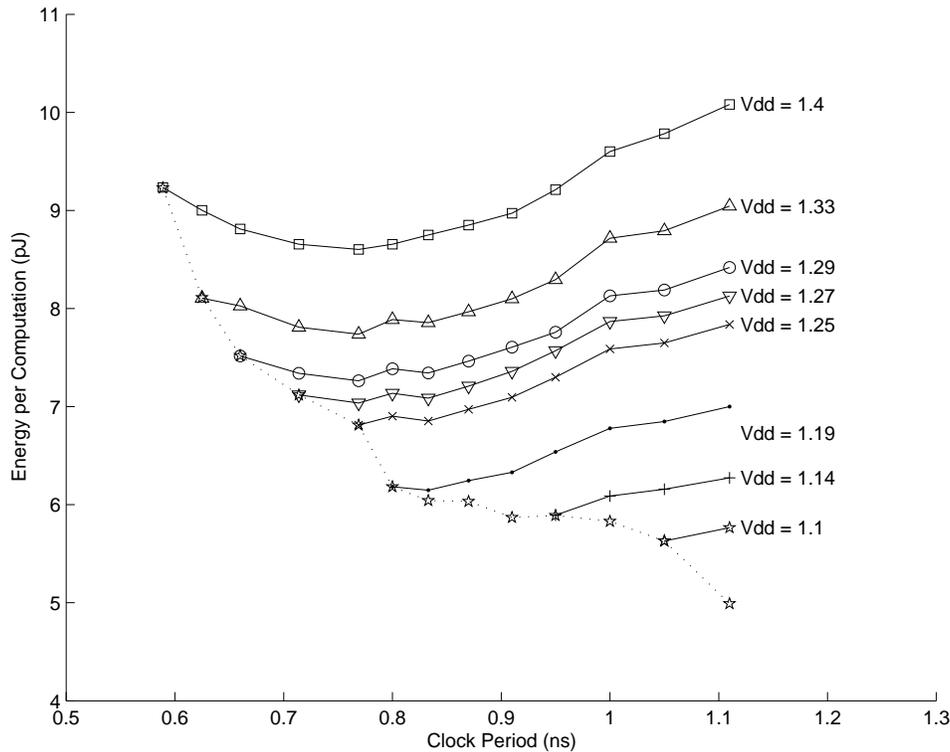


Figure 9: Energy-Delay variation in 8-bit array multiplier

the logic stage are always strongly in cutoff ($V_{gs} < 0$) when not conducting.

3.4.3 Energy Comparisons

In order to compare the energy efficiency of Boost Logic with respect to CMOS multipliers, an industrial tool was used to synthesize a pipelined multiplier. Unlike the boost multiplier, which was designed to be an array multiplier, the synthesis tool was allowed to perform logical optimization of the conventional multiplier netlist. The depth of the pipelined multiplier was determined on the basis of meeting a throughput of 1.6GHz with minimum energy dissipation. Synthesizing multipliers of various pipeline depths resulted in the selection of a nine-stage pipeline as the optimal pipe-depth for operation at 1.6GHz. Using a lower number of pipeline stages resulted in excessive dissipation due to the high operating voltage required to meet the throughput constraint. Using more pipeline stages resulted in increased overall energy dissipation due to the energy overhead of the latch dominating over the potential savings possible from voltage scaling. The conventional multiplier design obtained did not include clock buffers and therefore, the reported energy of the pipelined multiplier does not account for the energy dissipation due to the clock tree buffers.

The Boost multiplier simulation includes the energy dissipation in the multiplier as well as energy dissipated in clock generation and distribution. An on-chip inductor was designed for the clock generation circuit, and the extracted 13-element lumped RLC model for the inductor was used in the clock generator for simulations. In addition, the clock tree capacitance was estimated from the layout for the netlist simulation of the Boost multiplier. The energy results reported in the Boost multiplier simulation therefore include the dissipation in the clock generator and in clock distribution. The multipliers were not redesigned for different throughputs. Instead voltage-scaling was performed on the power-clock supply voltage of the multipliers to achieve lower energy dissipation at lower operating frequencies.

Figure 11 shows our simulation results. The curves depicted in the figure are energy delay curves for the synthesized CMOS multiplier and both versions of the Boost multiplier, nominal V_{th} and low V_{th} . As expected, the low DC supply voltage of the Boost Logic gate allows for significant power savings over pipelined, voltage-scaled CMOS de-

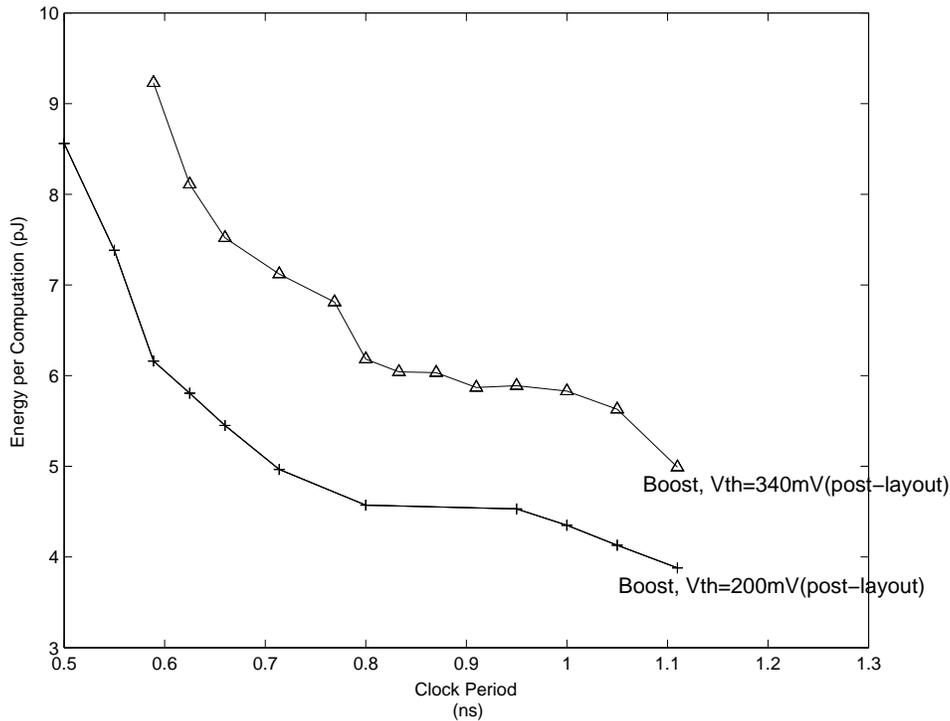


Figure 10: Energy delay curves for nominal V_{th} and low V_{th} 8-bit Boost multipliers

signs. When comparing simulation results at 1.6GHz, the Boost Multiplier offers 68% savings over the voltage-scaled CMOS multiplier. Using a low V_{th} design, these energy savings increase to 78% over the CMOS multiplier. Also, note that the use of low V_{th} enabled operation frequencies of up to 2GHz.

Being a fine-grained logic, Boost Logic has a latency of 12 cycles while the static CMOS design has a latency of 9 cycles. Therefore, Boost Logic may be more suitable for applications where latency is not critical.

4 Conclusion and future work

In this paper, we have proposed Boost Logic, a high-speed low-power energy recovery logic. In our analysis and simulations, we have addressed practical considerations involved in the design of Boost Logic through the characterization of clock skew (both internal and external), supply and process variation. Boost Logic was verified for correct operation with simultaneous internal and external clock skew amounting to 15% of the clock period.

We designed two 8-bit carry-save multipliers in Boost logic, using nominal and low V_{th} devices respectively. Our simulations indicate that Boost Logic achieves energy savings of 68% compared to voltage scaled CMOS at 1.6GHz. Using a lower V_{th} devices result in energy savings of 78% over CMOS. Thus, Boost Logic represents a significant step toward a structured, systematic approach to high-speed energy recovery design.

A design advantage offered by the structure of Boost Logic is the considerable power benefit achievable from the use of low V_{th} devices in the evaluation tree of the gates. The use of zero V_{th} is also possible, since the evaluation tree devices are either strongly on or strongly in cutoff, with negative V_{gs} .

Although Boost Logic uses an ultra-low DC power supply for its logical stage, it does not operate in the sub-threshold regime. It is therefore less susceptible than sub-threshold circuits to threshold voltage variation.

We have designed test circuits for the evaluation of Boost Logic in an industrial $0.13\mu\text{m}$ CMOS process and submitted them for fabrication.

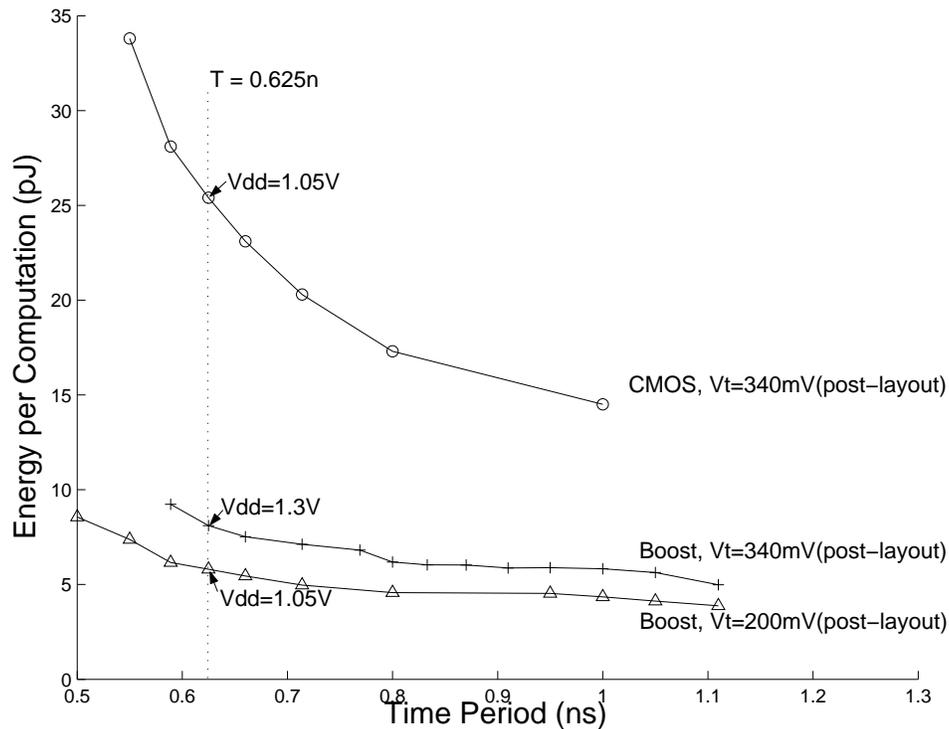


Figure 11: Energy consumption vs frequency for 8-bit multipliers

5 Acknowledgments

The authors would like to thank Sanjay Pant for his valuable input. This research was funded by the US Army office under Grant No. DAADA19-03-1-0122.

References

- [1] S. G. Younis and T. Knight, "Practical Implementation of Charge Recovering Asymptotically Zero Power CMOS," in *Symposium on Integrated Systems*, 1993.
- [2] Y. Moon and D. Jeong, "An Efficient Charge Recovery Logic Circuit," *JSSC*, April 1996.
- [3] W. Athas, N. Tzartzanis, L. Svensson, and L. Peterson, "A low-power microprocessor based on resonant energy," *JSSC*, Nov 1997.
- [4] D. Maksimovic, V. Oklobdzija, B. Nikolic, and K. Current, "Clocked CMOS adiabatic logic with integrated single-phase power-clock supply: experimental results," in *ISLPED*, Aug 1997.
- [5] C. Ziesler, J. Kim, V. Sathe, and M. Papaefthymiou, "A 225 Mhz Resonant Clocked ASIC Chip," in *ISLPED*, Aug 2003.
- [6] V. De and J. D. Meindl, "Complementary adiabatic and fully adiabatic mos logic families for gigascale integration," in *ISSCC*, Feb 1996.
- [7] A. Dickinson and J. Denker, "Adiabatic dynamic logic," *JSSC*, March 1995.
- [8] S. Kim, C. Ziesler, and M. Papaefthymiou, "A true single-phase 8-bit adiabatic multiplier," in *DAC*, June 2001.
- [9] Y. Yibin and K. Roy, "QSERL: Quasi-Static Energy Recovery Logic," *JSSC*, February 2001.
- [10] C. Seitz, "Hot-Clock nmos," in *Chapel Hill Conference on VLSI*, 1995.
- [11] D. Delegates, M. Barany, G. Geannopoulos, K. Kreitzer, M. Morrise, D. Milliron, A. Singh, and S. Wijeratne, "Low-Voltage-Swing Logic Circuits for a 7Ghz x86 integer core," in *ISSCC*, pp. 154–155, February 2004.

- [12] C. Ziesler, S. Kim, and M. C. Papaefthymiou, "Resonant clock generator for single-phase adiabatic systems," in *ISLPED*, Aug 2001.
- [13] J. Wood, T. Edwards, and S. Lipa, "Rotary traveling-wave oscillator arrays: a new clock technology," *JSSC*, Nov 2001.