Boost Logic : A High Speed Energy Recovery Circuit Family

Visvesh S. Sathe, Marios C. Papaefthymiou
Department of EECS, University of Michigan
Ann Arbor, USA
{vssathe,marios}@eecs.umich.edu

Conrad H. Ziesler
MultiGig Inc.
Scotts Valley, USA

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 simulation results from an 8-bit carry-save multiplier in a $0.13\mu m$ *CMOS process with* V_{th}*=340mV. At 1.4GHz and a 1.1V supply voltage, the Boost multiplier dissipates 3.44pJ per computation, achieving 57% energy savings with respect to its static CMOS counterpart. Using low* V_{th} *devices, Boost Logic has been verified to operate at 2GHz with a 1.2V voltage supply and 3.76pJ 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 higher operating frequencies 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, 5, 6, 9, 10]. 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 [2, 3], the use of pMOS devices in evaluation trees [4, 8], and the excessive time required to resolve the complementary outputs of the dual-rail gates during evaluation [5, 6].

In this paper, we propose a novel dynamic n-n logic family called Boost Logic. This family is a fine-grained, two-phase hybrid logic that consists of conventionally 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

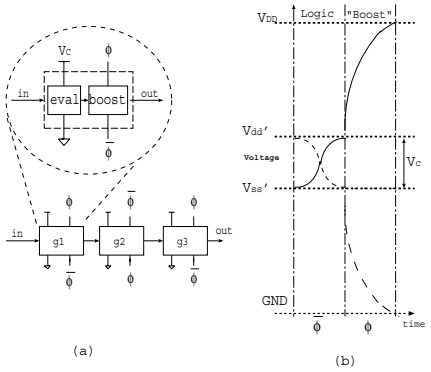

Figure 1. Boost Logic (a) Cascade and (b) Operation

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-recycling "Boost" stage. The logic stage operates at an ultra-low DC voltage supply and provides Boost Logic with greater voltage scalability as compared to fully energy recovering logic. An efficient amplifying stage ("Boost") is used at the output of the logic stage to boost the voltage level of the output nodes from $V_{dd}{}'$ to the nominal voltage V_{dd} and from $V_{ss}{}'$ to GND, as shown in Figure 1(b). V_c is approximately equal to V_{th}. The logic and boost stages of a Boost Logic gate operate in complementary clock phases.

In Boost, both dynamic and leakage power in the evalu-

Proceedings of the IEEE Computer Society Annual Symposium on VLSI
New Frontiers in VLSI Design
0-7695-2365-X/05 $20.00 © 2005 **IEEE**

ation 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 the n-type trees in the evaluate stage by the boost stage.

The idea of providing greater gate overdrive has been previously proposed [1, 7] in which bootstrapping was used. Such techniques lack the robustness offered by the boost stage however, and are limited in the amount of gate overdrive that can be achieved.

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

$$E = \frac{1}{2} \cdot C{V_c}^2 + E_{boost}, \tag{1}$$

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 carry-save multiplier with BIST was designed in an industrial 0.13μm process. At 1.4GHz, the Boost Logic multiplier dissipated a total of 3.44pJ in the logic and clock generator.

To compare the performance of Boost Logic with other design styles, we also implemented a pipelined, voltage-scaled CMOS multiplier. An industrial synthesis tool was used to generate a pipelined CMOS carry-save multiplier optimized for minimum energy dissipation at 1.4GHz. Energy comparisons between the two multipliers were made at the frequency of 1.4GHz. From the schematic simulations of the multipliers, Boost Logic achieved energy savings of 57% over its pipelined static counterpart. Using low V_{th} devices, Boost Logic has been verified to operate at 2GHz with a 1.2V voltage supply and 3.76pJ energy dissipation per cycle.

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 alloted for logical evaluation in each cycle.

The remainder of the paper is organized as follows: In Section 2, we present Boost Logic and 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 the robustness of Boost gates to clock skew 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 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 Structure

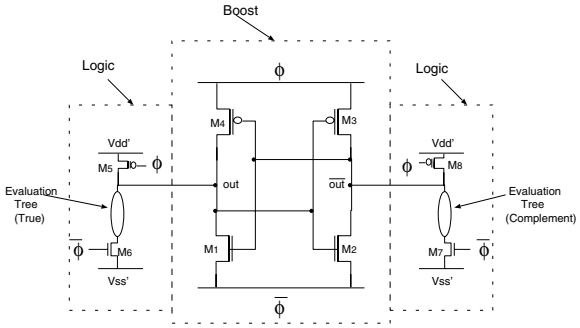

Figure 2. Boost Logic

Figure 2 shows a typical Boost Logic logic gate. Boost Logic is a two-phase, dual-rail, partially energy recovering n-n 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. For the purpose of robustness, the weak pMOS pullup can be made strong and a complementary pullup pMOS evaluation tree be added in series. The power supply rails are at voltages:

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

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

The choice of voltage values is motivated by the operation of the boost stage and will be discussed in greater detail in Section 2.2. The potential difference between the voltage supply rails in the logic stage is therefore $V_c = V_{th}$. The boost stage,

Proceedings of the IEEE Computer Society Annual Symposium on VLSI
New Frontiers in VLSI Design
0-7695-2365-X/05 $20.00 © 2005 **IEEE**

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 $\overline{\phi}$.

Boost Logic is a dual-rail logic that provides a balanced and data-independent capacitance to the power-clock by the gate, thus 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 $M5$, $M6$, $M7$, and $M8$. Since 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 $M5$ and $M8$.

2.2 Operation

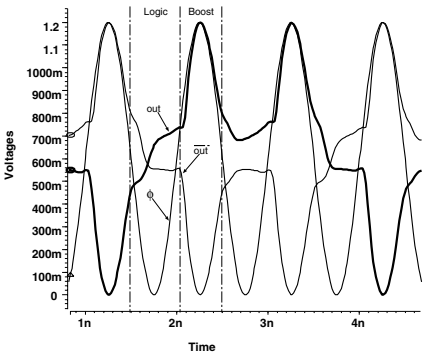

Figure 3. SPICE waveforms of a Boost Logic inverter

Figure 3 illustrates the operation of a Boost inverter. The complementary clock waveform $\overline{\phi}$ 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 $\overline{\phi}$ 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 $\overline{\phi} = V_{dd}$. 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 450mV in Figure 3), the logic stage is deactivated, disconnecting out and $\overline{out}$ 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 $\overline{\phi}$ (ϕ) goes past ${V_{ss}}'$ (${V_{dd}}'$), causing $\overline{out}$ (out) to subsequently follow $\overline{\phi}$ (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 $\overline{\phi}$, respectively. These outputs will drive the next gate during its logical evaluation stage.

As ϕ and $\overline{\phi}$ 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 all four transistors of the boost stage stops and the logic stage once again begins to evaluate.

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 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 logic evaluation of the gate 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 logic stage provides the complementary output nodes with a voltage difference of nearly V_c. Thus, the gate outputs are not 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-induced jitter.

The intermediate voltage rails in 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 fan-out capability. At the same time, the body contacts avoid performance degradation of the logic stage transistors 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 gate area considerably. Finally,

being a dynamic logic family, Boost Logic does not require pipeline registers to achieve high throughput.

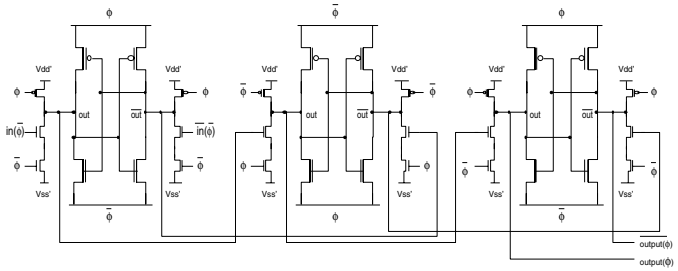

Figure 4. Cascade of Boost Logic inverters

Cascading Boost gates is straightforward. Since the boost conversion of a gate occurs 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 $\overline{\phi}$, as shown in Figure 1. A Boost Logic inverter chain is shown in Figure 4. 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 $\overline{\phi}$ 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 unconventional 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 by:

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

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

$$\delta \propto \frac{C}{\frac{V_{dd}}{2} - V_{th}}. \tag{5}$$

Considering first-order transistor effects, this result implies that unlike CMOS, the delay of the logic stage of the gate does not depend on the supply voltage of the conventional logic. This delay insensitivity to the conventional power supply can be explained by the fact that the transistors in the logic stage conduct in the linear mode and therefore behave like resistors. Since the delay incurred in charging and discharging the load through a resistor is independent of the power supply, the delay in the logic evaluation stage is insensitive to fluctuations in supply voltage considering first order transistor effects. 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 V_{th} variation on Boost Logic performance is an important practical consideration. Although Boost Logic uses a near threshold power supply to power its logic stage, the transistors in its logic stage do not operate in the sub-threshold regime. Instead, the transistors operate in the linear mode, where the sensitivity of gate delay to variations in 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:

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

where $\tau = RC_{boost}$ is the product of the resistance in the boost stage looking into either ϕ or $\overline{\phi}$ 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:

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

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 where the energy dissipation is proportional to V_{th}) 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 higher than ${V_c}^2$, the scaling factor $\pi^2\tau/(8T)$ is significantly smaller than $1/2$, even at operating frequencies of over 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 3.2.

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 present simulation results obtained from the 8-bit energy recovery multiplier along with Built-in Self Test. We also compare the energy consumption of the Boost Logic multiplier with pipelined, voltage-scaled CMOS implementations of the same multiplier.

Proceedings of the IEEE Computer Society Annual Symposium on VLSI
New Frontiers in VLSI Design
0-7695-2365-X/05 $20.00 © 2005 **IEEE**

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.

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^o out of phase to perform any computation, another kind of skew is possible wherein there exists a phase difference between ϕ and $\overline{\phi}$ for a given gate. We refer to such a phase difference between ϕ and $\overline{\phi}$ as internal skew.

To determine the robustness of Boost gates to both kinds of skew, we evaluated a parallel arrangement of basic Boost gates such as INV, AND, OR and XOR. 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. Simulations were carried out over the range of different internal skew and external skew values from -45% to $+45\%$ of the clock period.

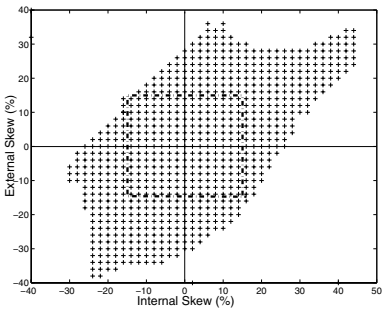

Figure 5. Schmoo plot for functional correctness over a range of internal and external skew values

Figure 5 shows the schmoo 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, the Boost Logic gates simulated all operate correctly with simultaneous internal and external skew amounting to 15% of the clock cycle.

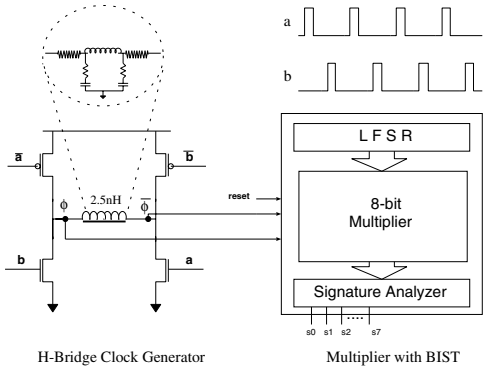

Figure 6. Overall simulation setup

3.2 8-bit Boost Logic carry-save multiplier

We have designed an 8-bit carry-save 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 6, an LFSR provides pseudo-random input vectors which were used by the multiplier as inputs. 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. In the experimental setup, the total capacitance driven by the clock generator (including the parasitic capacitance of the inductor and wiring capacitance of the clock tree) was approximately 20pF per phase. The value of inductance used depended on the frequency of operation. We also designed an identical multiplier using low V_{th} devices to evaluate the use of low V_{th} devices in Boost Logic gates . In this section, we compare the energy dissipation between the Boost and voltage-scaled pipelined CMOS multipliers. We also compare the energy-delay performance of a low V_{th} Boost multiplier with its nominal V_{th} counterpart.

To compare the energy efficiency of Boost Logic and CMOS multipliers, an industrial tool was used to synthesize a pipelined carry-save multiplier. The tool was constrained not to logically alter the multiplier netlist so as to maintain a fair comparison between the two multipliers. The CMOS multiplier was sized and pipelined on the basis of meeting a throughput of 1.4GHz with minimum energy dissipation. Synthesizing multipliers of various pipeline depths resulted in the selection of an 8-stage pipeline as the optimal pipe-depth for operation at 1.4GHz. Using a different number of pipeline stages resulted in higher energy dissipation. The reported energy of the CMOS multiplier does not account for the energy dissipation in the clock generation and distribution.

The Boost multiplier simulation includes the energy dissipation in the multiplier as well as energy dissipated in clock generation and distribution. A post-layout extracted 13-element lumped RLC model for the inductor was used in the clock generator for simulations. The wiring capaci-

Proceedings of the IEEE Computer Society Annual Symposium on VLSI
New Frontiers in VLSI Design
0-7695-2365-X/05 $20.00 © 2005 **IEEE**

tance of a resonant clock distribution network is significant and cannot be neglected. Consequently, the clock tree capacitance was estimated from placement and included in all Boost multiplier simulations. The energy results reported in the Boost multiplier simulation therefore include energy dissipation in the clock generator and the clock distribution network. The multipliers were not redesigned for different throughputs. Instead,voltage-scaling was performed on the CMOS supply voltage and the power clock voltage of the CMOS and Boost multipliers respectively, to achieve lower energy dissipation at lower operating frequencies.

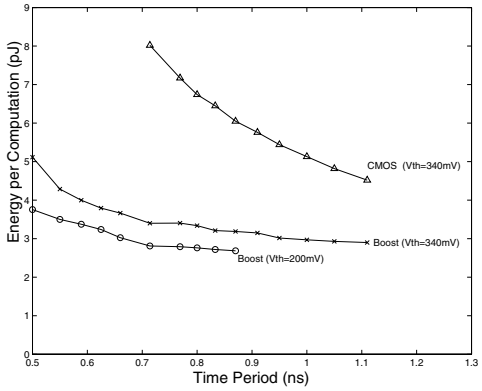

Figure 7. Energy consumption vs frequency for 8-bit multipliers

Figure 7 shows the results obtained from pre-layout simulation. The curves depicted in the figure are energy-delay curves for the synthesized CMOS multiplier and both versions of the Boost multiplier, normal V_{th} and a low threshold voltage version with $V_{th} = 200$mV. As expected, the low DC supply voltage of the Boost Logic gate allows for significant power savings over pipelined, voltage-scaled CMOS designs. When comparing pre-layout simulation results at 1.4GHz, the Boost multiplier offers 57% savings over the voltage-scaled CMOS multiplier.

Low V_{th} transistors in Boost multiplier gates enable faster evaluation in the logic stage of the Boost gate. 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. Using a low V_{th} design, pre-layout simulations at 1.4GHz indicate a decrease in power dissipation of 66% over the CMOS multiplier and 18% over its normal V_{th} counterpart. Furthermore, the use of low V_{th} transistors allows the Boost multiplier to operate at frequencies of over 2GHz (not shown in Figure 7).

Being a fine-grained logic, Boost Logic has a latency of 12 cycles while static CMOS has a latency of 8 cycles. Therefore, Boost Logic is 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-energy energy recovery logic. We have addressed practical considerations involved in the design of Boost Logic in our analysis and simulations through the characterization of Boost Logic operation with clock skew (both internal and external). Boost Logic was designed to provide a data-independent capacitive load to the resonant clock generator, minimizing data-dependent jitter. Simulations of an 8-bit carry-save multiplier indicate that Boost Logic achieved energy savings of 57% compared to voltage-scaled CMOS at frequencies over 1GHz.

A design advantage offered by the structure of a Boost Logic gate is the considerable power benefit achievable from the use of low V_{th} devices in the evaluation tree of the gates. The use of low V_{th} in the Boost multiplier achieved 66% energy savings over static CMOS. The use of zero V_{th} is also possible since the evaluation tree devices are either strongly on, or in cutoff with negative V_{gs}.

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

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] W. Athas et al. A low-power microprocessor based on resonant energy. *JSSC*, Nov 1997.

[2] V. De and J. D. Meindl. Complementary adiabatic and fully adiabatic mos logic families for gigascale integration. In *ISSCC*, Feb 1996.

[3] A. Dickinson and J. Denker. Adiabatic dynamic logic. *JSSC*, March 1995.

[4] S. Kim et al. A true single-phase 8-bit adiabatic multiplier. In *DAC*, June 2001.

[5] D. Maksimovic et al. Clocked CMOS adiabatic logic with integrated single-phase power-clock supply: experimental results. In *ISLPED*, Aug 1997.

[6] Y. Moon and D. Jeong. An Effi cient Charge Recovery Logic Circuit. *JSSC*, April 1996.

[7] C. Seitz. Hot-Clock nmos. In *Chapel Hill Conference on VLSI*, 1995.

[8] Y. Yibin and K. Roy. QSERL: Quasi-Static Energy Recovery Logic. *JSSC*, February 2001.

[9] S. G. Younis and T. Knight. Practical Implementation of Charge Recovering Asymptotically Zero Power CMOS. In *Symposium on Integrated Systems*, 1993.

[10] C. Ziesler et al. A 225 Mhz Resonant Clocked ASIC Chip. In *ISLPED*, Aug 2003.