.

Towards An Efficient Low Frequency Energy Recovery Dynamic Logic

Submitted in partial fulfillment of the requirements
for the

Computer Science and Engineering

Preliminary Examination

by

Sujay S. Phadke

Advisor:
Prof. Marios Papaefthymiou

Advanced Computer Architecture Laboratory
Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor
September, 2005

Abstract

A large number of energy recovery circuits have been described to reduce the power consumption of CMOS logic circuits. Many of these circuits use diode connected transistors, accurately timed multiple clocks and self resonating power clocks. Boost logic is a 2 phase high speed energy recovery family that voltage scaling and high gate overdrive to achieve highly efficient energy recovery.

In this report, we investigate 3 new designs to improve the energy recovery efficiency of Boost logic at low frequencies. The first design involves the use of CMOS stacks in the evaluation stage of Boost logic. Simulations show a significant reduction in crowbar current and hence energy dissipation, at the cost of high area overhead.

The second design uses CMOS stacks and an inverter to create differential outputs in evaluation. Simulations indicate lesser energy dissipation but a lower operational range of frequencies because of sub-threshold operation of the inverter.

Next we investigate a new design based on an energy recovery domino logic topology which achieves significant reduction in energy dissipation with small area overhead. We discuss the experimental results of HSPICE simulations with a 32-bit adder using the proposed circuit topologies. The logic proposed achieves energy saving of 65% to 25% in frequency range 20MHz to 150MHz over Boost logic. Simulations indicate that it is more robust to power supply variations compared to conventional CMOS.

1 Introduction

As technology scales into the sub 100nm region, power consumption is increasingly becoming a dominant factor, sidestepping its traditional partner speed. The performance of future VLSI systems is going to be dictated primarily by power. Low power design techniques attempt to solve the problem of dynamic power and leakage by various techniques, viz. algorithmic, system level design, architectural and circuit design. Energy recovery circuits are a class of circuits which try to lower the dynamic power dissipation. The fundamental idea is to use resonant clocks which serve as the power rails. Part of the energy spent in charging capacitance can be recovered back and stored into an element like an inductor. This transfer of charge back and forth between the inductor and capacitance can lead to very low power consumption if performed efficiently.

One of the fundamental needs of an energy recovery system is gradually transitioning clock signals. Various circuits based on trapezoidal clocks and diode connected transistors have been proposed. Difficulty in generating the required clock signals and slow speed of operation are the drawbacks associated with them. The trapezoidal clock can be replaced by sinusoidally

varying clocks which can be generated naturally by using a LC oscillation circuit. The major design concerns are the ability to tune the circuit to the desired operational frequency and the efficiency of energy recovery. Data dependent capacitance offered by a circuit poses a challenge to resonating such systems and can offset any gains obtained through recovery.

A high speed energy recovery family, Boost logic, operating on a 2-phase clock was proposed by Sathe, [1] et. al. and was shown to operate upto 1.3 GHz. While the hybrid logic can achieve significant energy savings over a conventional voltage scaled CMOS design, it has certain drawbacks at lower frequencies. In this paper, we propose a modified energy recovery logic, which is much more efficient at lower frequencies. We provide results from simulations indicating the benefits of this logic in the desired operation regime.

This remainder of this report is organized as follows. Section 2 presents an overview of related work and the necessary background. Section 3 gives an introduction to Boost logic, its structure and operation. Section 4 mentions the contributions of this report. Section 5 discusses about a variant of Boost logic with its pros and cons. Section 6 introduces the idea of a energy recovery domino style topology and examines its performance. In Section 7 we compare different circuit styles with experimental results and give an insight into the merits of the proposed logic. In Section 8 we conclude and give information about future research.

2 Related work and background

Power dissipation today is one of the most important issues in VLSI system design. For high performance design, power dissipation can be the limiting factor to clock speed and circuit density because of the inability to get power to the circuits or to remove the heat that they generate. The challenges faced by designers of hand held computers and mobile phones iterate the importance of low power design. The problem is finding a good tradeoff between performance and power dissipation. The prevalent approach to trading performance for power has been to reduce the chip-wide supply voltage, which has the effect of decreasing the operational energy levels and switching speeds. The basic idea is to recycle charge transfer between a power supply and a load. In order to turn a transistor ON, we need to transfer energy. The energy dissipated during this transfer need not be related to the energy transferred, but in ordinary CMOS logic circuits both quantities are on the order of CV_{dd}^2, where C is the capacitance of a typical node, and V_{dd} is the operating voltage. This level of dissipation is unavoidable if all the needed electrons are extracted from the V_{dd} terminal of the power supply and returned to the ground terminal. [13] The essential idea of adiabatic computing is to construct circuits that allow each needed electron to be extracted at the lowest feasible voltage and returned at the

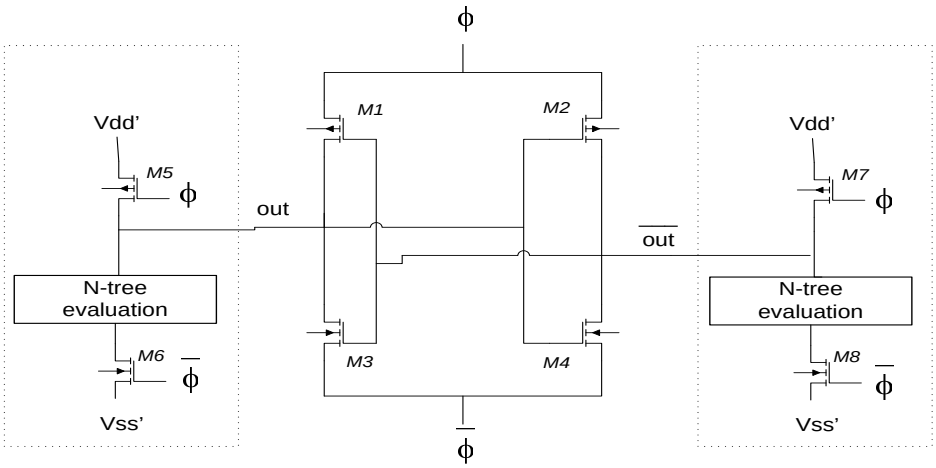

Figure 1: Structure of the Boost logic

highest feasible voltage.

Energy recovery circuits offer an alternative approach to the reduction of dynamic energy dissipation. Several energy recovery logic styles have been proposed [3–5, 7, 14]. These energy recovery techniques have been shown to achieve the significant energy gains when compared to voltage scaled CMOS. Some of the main factors that lead to lower speeds in energy recovery circuits are the use of diode-connected transistors [8, 9], the use of pMOS devices in evaluation trees [10]. The most relevant work here, is the logic family called Boost Logic [1]. 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, without the use of diodes or multiple clocks. However, as will be shown in later sections, the recovery efficiency is best in frequency ranges of 800MHz-1GHz, and increases in the lower frequency ranges (50MHz-200MHz). In order to solve this, we first look at the basic structure and operation of Boost logic. We then propose modifications to the gate structure which will enable us to have better performance in our domain of interest (ie: low frequency) while maintaining a good quality factor.

This lower range of frequencies is particularly important in DSP applications. Core functional units in JPEG or MPEG processors, like variable length decoders, FFT, IDCT or motion compensation units operate in the order of a few hundreds of MHz. Another potential application area is embedded and mobile computing, for which battery life is a primary concern.

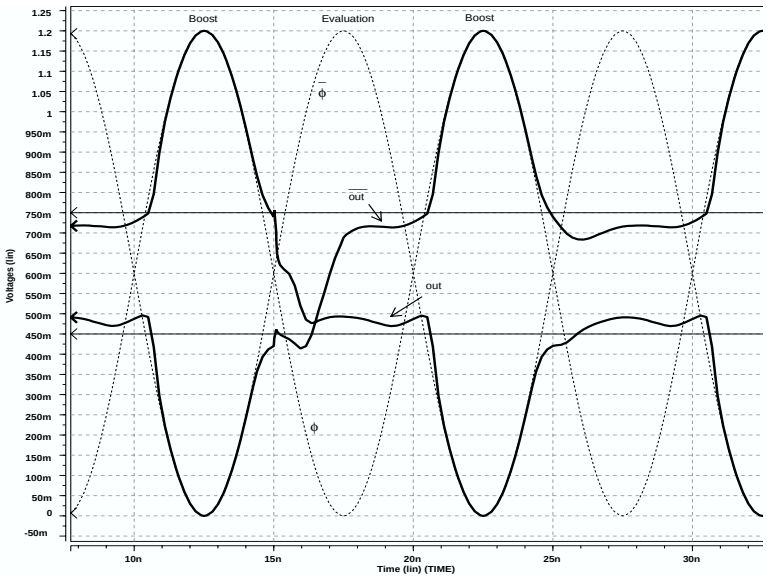

Figure 2: Operation of the Boost gate (Type 1)

3 Fundamentals of Boost logic

The structure of Boost logic is shown in Fig. 1. 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 pseudo nMOS-like gate is chosen to reduce the loading on the gate thereby improving performance. The potential difference between the voltage supply rails in the logic stage is $V_c = V_{th}$. The boost stage, which is essentially an energy recovering sense amplifier, resembles back-to-back CMOS inverters. The only difference is that the rails are replaced by ϕ and $\overline{\phi}$. 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.

3.1 Operation of the Boost gate

The operation of the Boost gate is shown in Fig. 2. When ϕ is low and $\overline{\phi}$ is high, the gate is in evaluation mode. The pseudo-NMOS tree evaluates and drives the *out* and $\overline{out}$ lines to the rails V'_{dd} and V'_{ss}. This small difference is then amplified by the "Boost" stage when ϕ goes high and $\overline{\phi}$ goes low. The output lines *out* and $\overline{out}$ now follow the power clock and return to the rails at the end of Boosting.

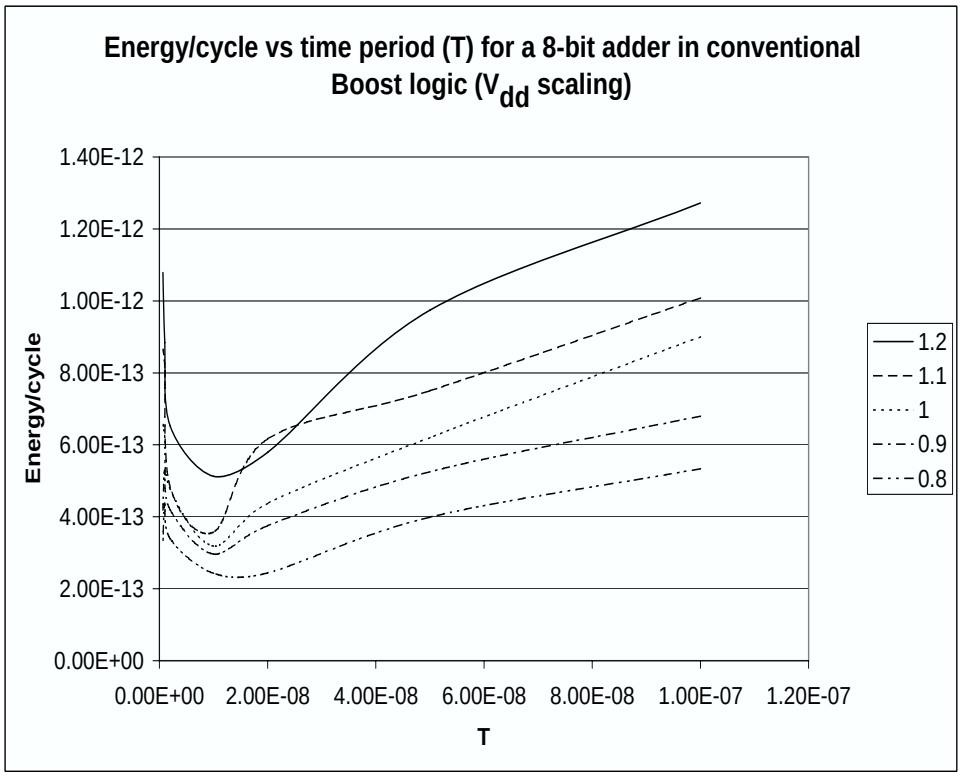

Figure 3: Increasing energy dissipation at low frequencies (for Type 1)

3.2 Performance of Boost Logic at Low Frequencies

Boost logic has excellent performance over vanilla CMOS in the near 1GHz frequency range. It is reported to have 65% energy saving over a similar voltage scaled pipelined design. [1]. The design is also shown to be robust to clock skew and power supply variations, much more than CMOS. The delay penalty due to the header and footer can be reduced by sizing. The major drawback of Boost logic is the high crowbar current in the evaluation circuit, which becomes significant at low frequencies. When the gate evaluates, both the header and footer turn ON. Due to this, there is always a straight path from V'_{dd} to V'_{ss} and a fight between M5, M7 and their corresponding pull down networks. The increased energy dissipation at low frequencies is seen in Fig. 3. In an energy recovery system, the energy dissipated in charging or discharging a capacitance is shown [17] to be equal to

$$E_{diss} = \frac{RC}{T} CV^2,$$

where C, V, R, T are the total effective circuit capacitance, power supply voltage, effective circuit resistance and the time period respectively. As the time period T keeps increasing, the energy dissipated keeps on decreasing. This can be seen in Fig. 3, where the energy drops initially

at high frequencies (lower T). In this range of frequencies, the recovery is efficient and the crowbar current and leakage in the evaluation stage are much less. However, as the time period is increased, the crowbar current contribution to the energy grows linearly as $I_{crow} \cdot V \cdot T$. Thus, we see the rising parts of the curves for large T.

4 Contribution: 3 Types of Circuit Configurations Analyzed

The focus of our research was to design efficient energy recovery circuits as the lower frequencies, using a sinusoidal power clock. One of the main issues to be addressed was the increased energy dissipation at low frequencies. We have tried out various approaches to this problem. Some of these are replacing the NMOS stack with a complementary MOS stack, use of multiple threshold devices, dynamic logic evaluation and pulsed evaluation.

(Type 1) This is the classic pseudo-NMOS based BOOST logic. [1]

We refer to this reference as Type 1 throughout this report. The circuits that we designed and analyzed are the following:

(Type 2) Replaced the evaluation tree with complete CMOS trees on both sides, as shown in Fig. 4(a).

(Type 3) CMOS tree on one side and a clocked inverter on the other side, as shown in Fig. 4(b). We designed 2 variations in this style:

1. Use of nominal V_{th} devices
2. Use of zero V_{th} devices in the inverter stack

(Type 4) Modified domino logic evaluation stage and a sense amplifier stage, as shown in Fig. 4(c).

We also performed experiments by keeping the circuit unchanged, but pulsing the header and footer of the evaluation stacks with square wave pulses, instead of ϕ and $\overline{\phi}$. Interestingly, this did not give any power gains compared to the Type 1 circuit. The problem lies in using a square wave pulse to switch on the header and footer, by which we increase the voltage level term in

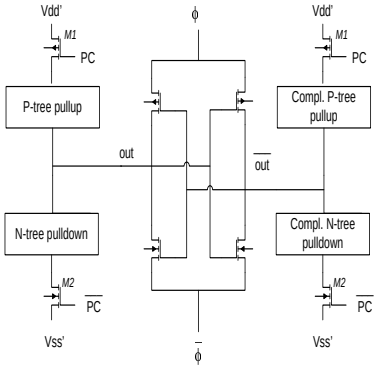

(a) Type 2

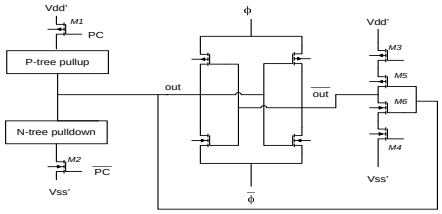

(b) Type 3

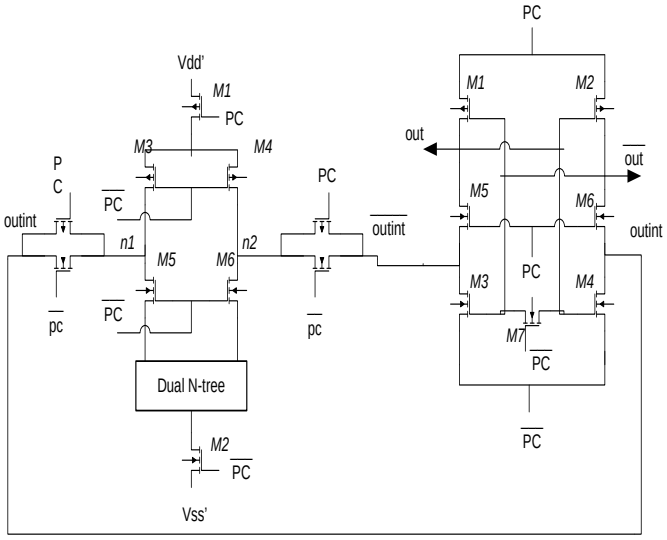

(c) Type 4

Figure 4: Different circuit configurations

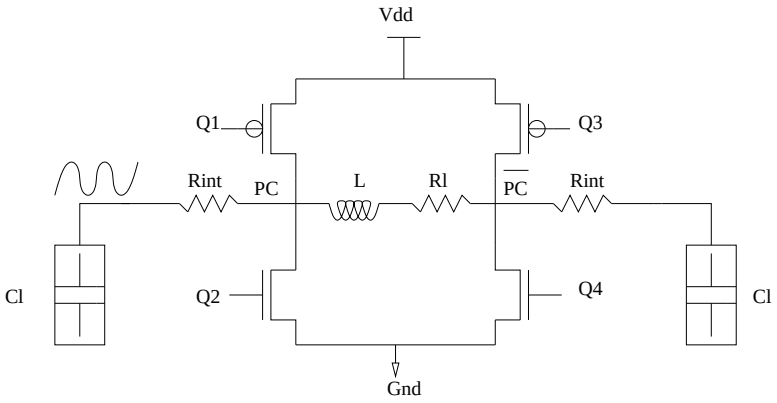

Figure 5: Structure of the clock generator used for simulations

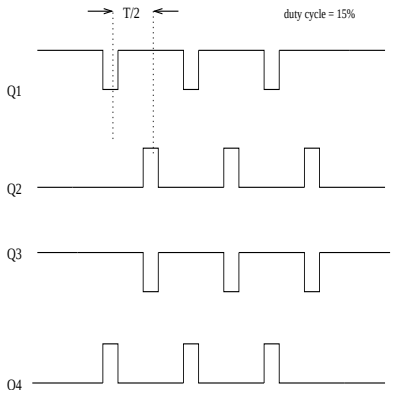

Figure 6: Pulsing waveforms applied to the replenishing switches

the short circuit current power expression $I_{crow} \cdot V \cdot T$. We tried to offset this by decreasing the duty cycle of the pulse, but the logic cannot be clocked below a certain duty cycle, if it has to switch within the given cycle time. Moreover, supplying this separate square wave pulse would require generating and distributing another clock line which is impractical.

The following sections will discuss the 3 circuits in detail, with HSPICE simulations results and comparisons.

4.1 Structure of the Clock Generator used

For the purposes of generating the sinusoidal power clock, we used a clock generator as shown in Fig. 5. For resonating the circuits in the frequency range of 20MHz-200MHz, it would be appropriate to introduce an off-chip inductor with inductance in tens of nanohenries and a small series resistance(from a quality standpoint). Such components available commercially have a

resistance of less than 1Ω for an $L < 50nH$. The inductor used here is L=40nH. Approximate values for the inductor series resistance is $R_L = 2\Omega$ and interconnect resistance is $R_{int} = 2\Omega$. The switches serve as replenishing switches and pull the voltage at the outputs of the inductor to full rail in a near adiabatic fashion. Each switch has been designed as a large transistor with multiple fingers. The replenishing switches are clocked by square wave pulses shown in Fig. 6.

5 Type 2 and Type 3 Circuits: Use of CMOS Stacks

The pseudo-NMOS evaluation tree performs well at higher frequencies, with regards to energy dissipation, when compared to vanilla CMOS, by reducing the switched capacitance and area. The energy profile for Type 1 shows a trend of achieving an energy minimum around 1GHz. For low frequencies however, there is increased crwobar in the evaluation circuit. As the time period keeps on increasing, the crowbar current manifests itself as a larger contribution to the overall energy dissipation, leading to large amounts of losses. For lower frequencies, the crowbar current in the evaluation stack needs to be reduced. In the Type 2 and Type 3 circuits, we make use of CMOS stacks for evaluation, which would help in reducing the crowbar current.

5.1 Type 2 Circuit

5.1.1 Structure of the Gate

The structure of this gate is shown in Fig. 4(a). The evaluation trees receive complementary inputs. The header and footers transistors are clocked by ϕ and $\overline{\phi}$ and disconnect the evaluation stacks from the sense amplifier. The gate uses a dual rail structure to provide a data independent capacitance to the clock network. In the evaluation circuit, the power rails are as follows:

$$V'_{dd} = \frac{V_{dd}}{2} + \frac{V_c}{2}$$

$$V'_{ss} = \frac{V_{dd}}{2} - \frac{V_c}{2}$$

5.1.2 Operation

The circuit operates in 2 phases - charge/discharge and sense phase. During the first phase, ϕ is low and the headers and footers in the evaluation stage are turned ON. The CMOS stacks either evaluate high or low in a complementary fashion, thus setting the voltage on the *out* and $\overline{out}$ lines. During the next phase, ϕ goes high, $\overline{\phi}$ goes low, and the sensing circuit amplifies

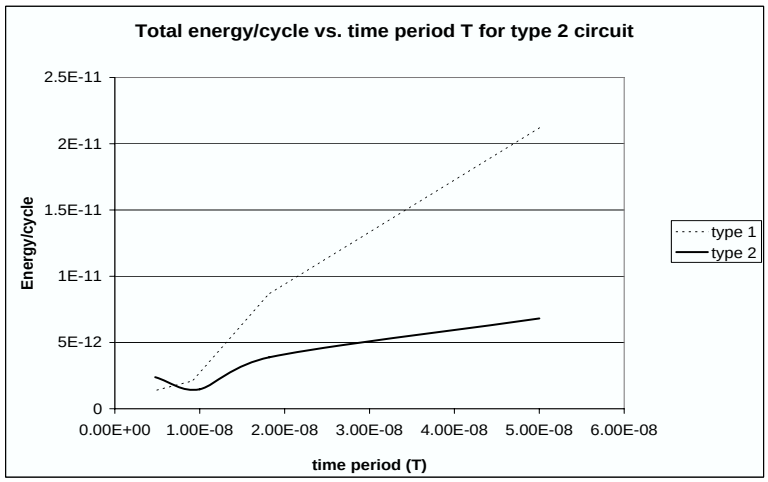

Figure 7: Energy dissipation of a 32-bit ripple carry adder in Type 2

the difference in an energy recovery manner, thus driving the output lines to their full rails following the clock.

5.1.3 Simulation Results

We designed a 32-bit ripple carry adder using this circuit type. The simulation results using the clock generator are shown in Fig. 7. As seen, in the frequency range of 20MHz-200MHz (T from 5×10^{-8} to 0.5×10^{-8}), the circuit dissipates much lower energy than Type 1, primarily due to very low crowbar current. The percentage contributions to the total energy for Type 1 and Type 2 is shown in Fig. 8. The contribution of crowbar current is relatively less in the Type 2 circuit. This reduces the total energy/cycle significantly at lower frequencies. However, the main drawback of this circuit is the increased area overhead. We pay a price in using double the area. Sizing of the PMOS and NMOS devices also needs to be done to ensure equal rise and fall times. The increase in area (double the number of transistors as compared to Type 1) leads to more amount of capacitance to resonate, affecting the quality of the circuit. This problem can be addressed if we can create a differential voltage during evaluation by avoiding the use of a complete CMOS stack on the other side. This technique is discussed in the following section.

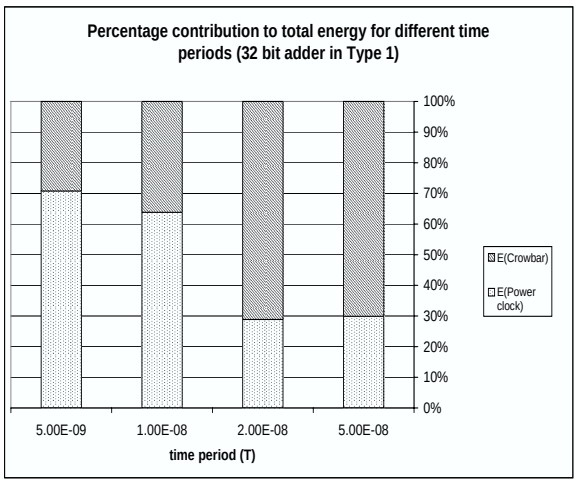

(a) Type 1

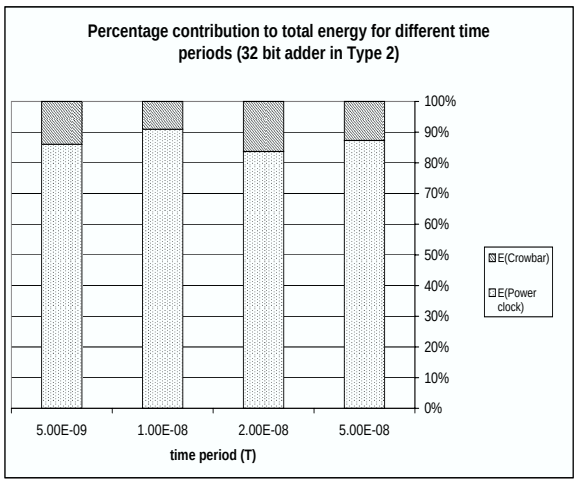

(b) Type 2

Figure 8: Percentage contribution to total energy/cycle for different time periods (T)

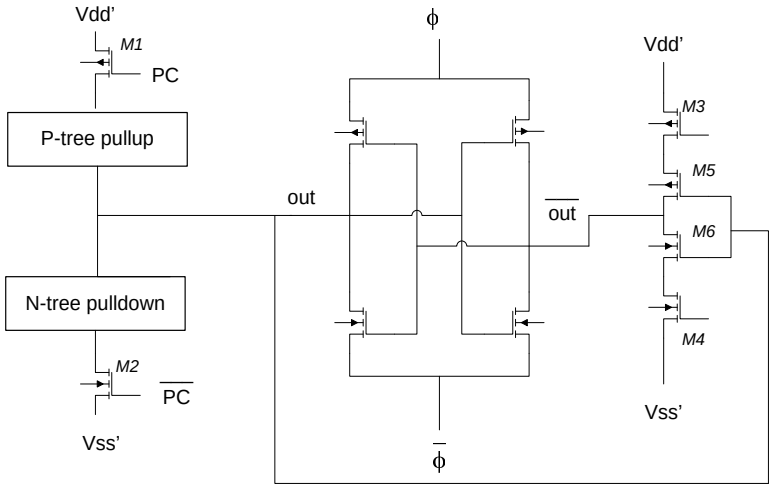

Figure 9: Structure of the Type 3 gate

5.2 Type 3 Circuit

5.2.1 Structure of the Gate

The structure of the Type 3 gate is shown in Fig. 9. In order to reduce the crowbar current, an appropriately sized CMOS stack can be used for evaluation. One of the major drawbacks with the classical Boost gate is that there is always a fight between the weak PMOS pull-up and the evaluation logic pull down. This configuration eliminates the straight path from V'_{dd} to V'_{ss} and reduces crowbar current. In order to create a voltage differential across the *out* and $\overline{out}$ lines, it is not required to have a complete complementary stack on the other line, which would require double the number of transistors. The *out* can be fed to an inverter to create the necessary difference as shown. The gate uses a dual rail structure to provide a data independent capacitance to the clock network. In the evaluation circuit, the power rails are as follows:

$$V'_{dd} = \frac{V_{dd}}{2} + \frac{V_c}{2}$$

$$V'_{ss} = \frac{V_{dd}}{2} - \frac{V_c}{2}$$

The reason for choosing these voltage levels is that we want the evaluation logic to discharge the nodes only through a small voltage.The sense circuit is powered by the clock ϕ and $\overline{\phi}$. After evaluation, the *out* and $\overline{out}$ lines follow the power clock as they slowly get pulled to the rails.

5.3 Operation

The circuit operates in 2 phases - charge/discharge and sense. During the first phase, ϕ is low and the headers and footers in the evaluation stage are turned ON. The CMOS stack either evaluates high or low, thus setting the voltage on the *out* line. This is then fed to the inverter which drives the $\overline{out}$ line and this produces the voltage differential. During the next phase, ϕ goes high and the sense circuit amplifies the difference in an energy recovery manner, thus driving the output lines to their full rails following the clock signals ϕ and $\overline{\phi}$.

5.3.1 Simulation Results

We designed a 32-bit ripple carry adder using this circuit type. Fig. 10 shows the trend energy dissipation with frequency, compared to the Type 1 circuit. As seen, in the frequency range of 20MHz-200MHz (T from 5×10^{-8} to 0.5×10^{-8}), the circuit dissipates much lower energy than Type 1. It is higher than the Type 2 circuit, because of crowbar in the inverter stage

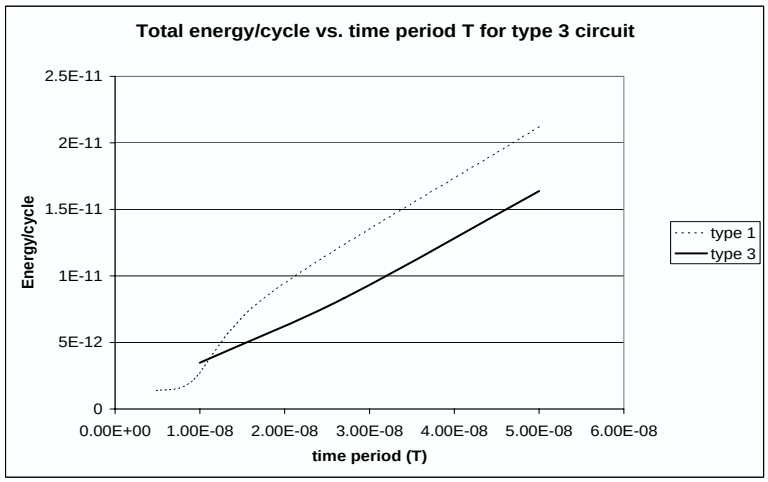

Figure 10: Energy dissipation of a 32-bit ripple carry adder in Type 3

(comparison in Section 7). The flip side of it is that we have a lower area overhead and hence less capacitance to resonate. The percentage contributions to the total energy for Type 2 and Type 3 is shown in Fig. 11. The HSPICE simulation waveforms (with the clock generator) for the Type 3 circuit are shown in Fig. 12. The main drawback of this circuit topology is the operation of the inverter. The inverter, being supplied V'_{dd} and V'_{ss} through the header and footer is driven by the *out* at its input. The gate overdrive, $V_{GS} = V_G - V'_{ss} < V_{th}$. Thus it operates in the sub-threshold region and is unable to drive the $\overline{out}$ to its full rail of V'_{dd}. The voltage difference as a function of frequency is shown in Fig. 12. As seen, the output being driven by the inverter is slow in reaching its full rail value. The maximum differential created is less than $100mV$ for $f > 100MHz$, assuming a loading of one inverter on the output lines. Thus this circuit is sensitive to coupling noise which would invariably be introduced by aggressor nets.

5.3.2 Use of Low Threshold Devices in Type 3 Circuit

The small voltage at the output of the inverter is due to the limited drive of the *out* input (to the inverter) and subsequent sub-threshold operation of the inverter side transistors. The output response of the inverted side is very sluggish. We performed experiments using zero V_{th} devices in the inverter stack. This modification would slightly increase the energy dissipation but help in increasing the output level of the weakly driven inverter, and perform the evaluation faster. The IBM $0.13\mu m$ library provides zero V_{th} NMOS devices. The V_{th} of the PMOS devices has to be lowered by body biasing. Nominal V_{th} for this technology is around 340mV. Body biasing was used to bring down the V_{th} of the PMOS devices to around 260mV. Fig. 13 shows

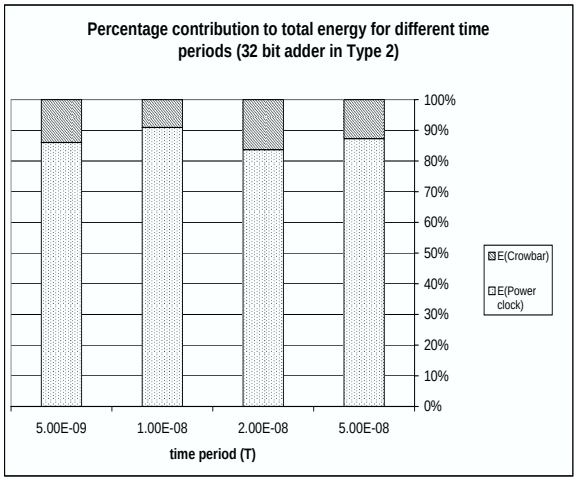

(a) Type 2

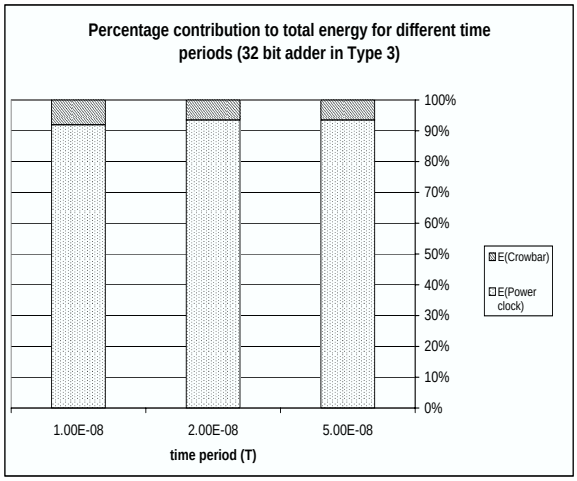

(b) Type 3

Figure 11: Percentage contribution to total energy/cycle for different time periods (T)

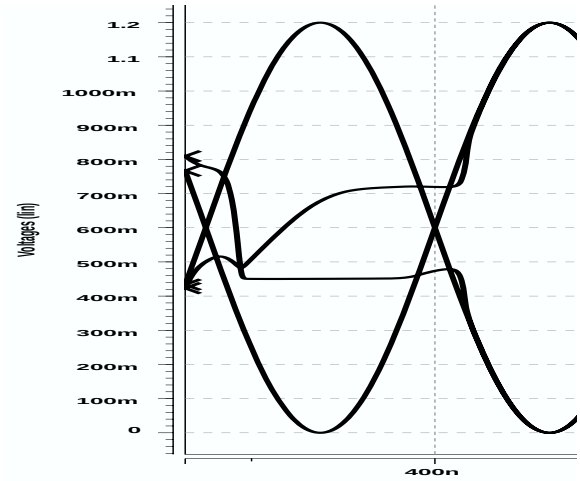

(a) at 10MHz

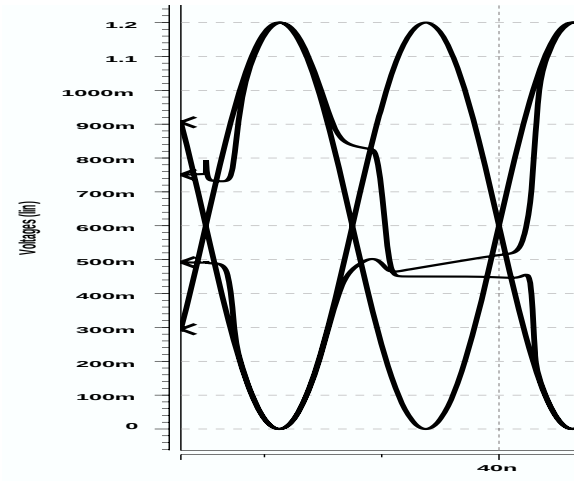

(b) at 100MHz

Figure 12: HSPICE simulations of the Type 3 gate at 10MHz and 100MHz

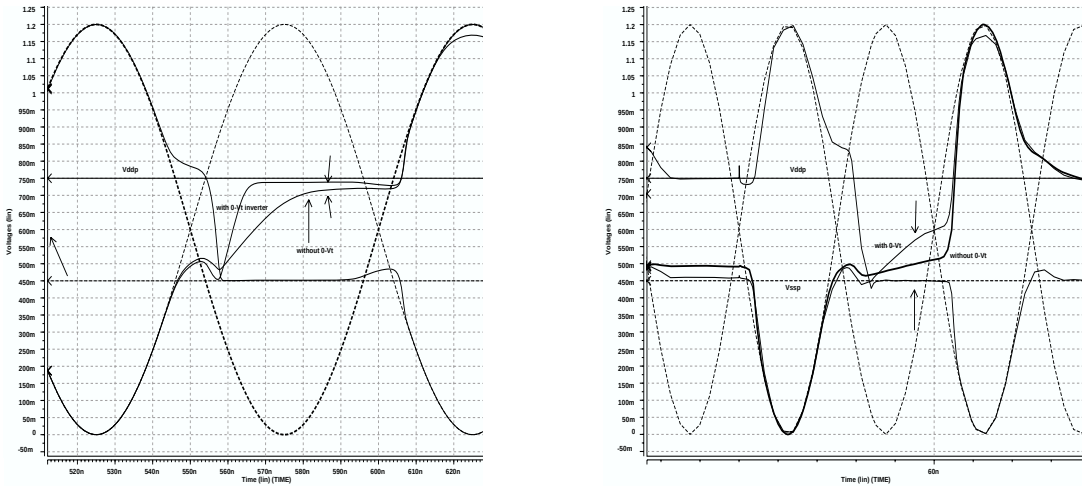

(a) at 10MHz (b) at 100MHz

Figure 13: HSPICE simulation for the Type 3 gate with and without zero V_{th} devices at 10MHz and 100MHz

the output waveform at the inverter side with and without the zero V_{th} devices, at 10MHz and 100MHz. Referring to 13(a), at 10MHz, the output differential voltage at the end of evaluation is very close to V'_{dd} for both cases. However, at 100MHz, as seen in 13(b), this differential value drops to only $V_c/2$ for the stack with zero V_{th} devices, and an abysmally low 50mV for the stack without zero V_{th} devices. Practically, it would be difficult to sustain this difference, taking into account circuit parasitics and coupling noise. The gains due to reduction in area obtained by using an inverter (instead of a complete complementary CMOS stack on the other side) and the subsequent lower capacitance, are offset by the reduced drive. This would further worsen with increased fanout. The circuit is operable practically only for $f < 50Mhz$, where the differential output voltage is at least $100mV$.

This circuit configuration helps in reducing total energy consumption significantly as compared to the pseudo-NMOS, especially at lower frequencies (between 10MHz-200MHz), without much area overhead. However, it suffers from high noise susceptibility and a lesser range of operable frequencies. In the next section, we explore a new design which is found to create a more robust voltage differential, independent of fanout capacitance and keeping the area overhead minimal.

6 Type 4: Domino CMOS with Reduced Internal Node Swing

The main drawback of the pseudo-NMOS evaluation logic at low frequencies is the increased energy/cycle dissipation as the frequency is decreased. This is because at low frequencies, the crowbar current becomes a dominant factor of overall power dissipation, offsetting the gains due to recovery. Using a CMOS stack and inverter had certain advantages, as outlined in Section 5. However, the low swing at the output of the inverter due to sub-threshold operation results in only a small differential voltage which is highly susceptible to coupling noise. In this section, we describe the structure of modified circuit using domino logic for the evaluation.

6.1 Structure of the logic

The structure of the energy recovery domino gate is shown in Fig. 14. In the evaluation circuit, transistors M1 and M2 clock the logic. Transistors M3 and M4 are used for precharging the internal nodes high. M5 and M6 prevent the internal nodes of the dual N-tree from being charged to full rail, thus significantly reducing power consumption in this stage. The transmission gates are used to transfer charge from the data lines and isolate the circuit from the sense circuit during the sense phase. The sensing circuitry is modified compared to the classical Boost logic. The power supplies are sinusoidal clocks ϕ and $\overline{\phi}$ which are out of phase with each other. There are 2 proxy *outint* and $\overline{outint}$ lines, which are low capacitance lines are driven by the evaluation logic. Transistors M5 and M6 isolate these lines from the higher capacitance *out* and $\overline{out}$ lines during evaluation. M7 is used for equalization of the 2 outputs, during the evaluation phase. The gate uses a dual rail structure to provide a data independent capacitance to the clock network. In the evaluation circuit, the power rails are as follows:

$$V'_{dd} = \frac{V_{dd}}{2} + \frac{V_c}{2}$$

$$V'_{ss} = \frac{V_{dd}}{2} - \frac{V_c}{2}$$

The reason for choosing these voltage levels is that we want the evaluation logic to discharge the nodes only through a small voltage. The voltage differential between the lines will then be amplified in a recovery manner by the sensing circuit. V_c has to be chosen so that we have an optimum trade-off between energy dissipation and speed of operation, and there is no reverse flow of charge into the supplies V'_{dd} and V'_{ss}. Here, we choose $V_c = V_{th}$. The bulk connection for all NMOS devices is made to V'_{ss} and that of PMOS devices is made to V'_{dd}. This has the

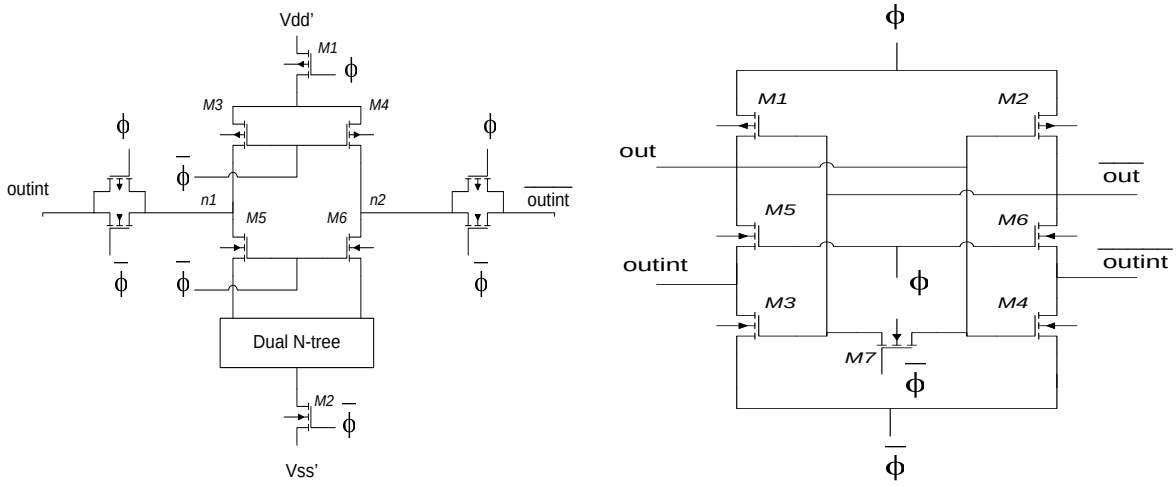

(a) Dynamic logic evaluation stage (b) Modified Sense amplifier circuit

Figure 14: Structure of the Type 4 gate

advantage of providing a forward body biasing advantage and enhancing operation.

6.2 Operation

The dynamic Boost gate operates in 2 phases. During the 1st phase, ϕ is low and $\overline{\phi}$ is high. This is the evaluate/hold phase. In the sense circuit, transistors *M5* and *M6* are switched off. The outputs *out* and $\overline{out}$ in the sense circuit are held at approximately $VDD/2$ with the help of the equalizing transistor *M7*. The header and footer in the evaluation circuit are switched on. The internal nodes *n1* and *n2* were precharged high during the precharge phase (described later). In this phase, the N-tree pulls down one of these nodes and sets up a voltage differential between *n1* and *n2*. The transmission gates are also switched on and the charge is transferred onto the low capacitance lines *outint* and $\overline{outint}$. The low capacitance of these lines helps in setting up a sufficient differential voltage (independent of output loading and fan-out). In the next phase precharge/Boost, ϕ goes high and $\overline{\phi}$ goes low. At this time, the evaluation circuit is cut-off from the Boost circuitry by the transmission gates. The intermediate transistors *M5* and *M6* are now switched on, and the circuit behaves as a cross-coupled inverter. The voltage difference between *outint* and $\overline{outint}$ is sensed and amplified. As ϕ goes high and $\overline{\phi}$ goes low, *out* and $\overline{outb}$ follow the output sinusoidal waveforms. At this time, in the evaluation circuit, transistors *M5* and *M6* are OFF. *M3* and *M4* precharge the internal nodes to V'_{dd}. The transistors *M5* and *M6* prevent the internal nodes of the pull down N-tree from charging up high, and hence reduce the energy consumption during evaluation.

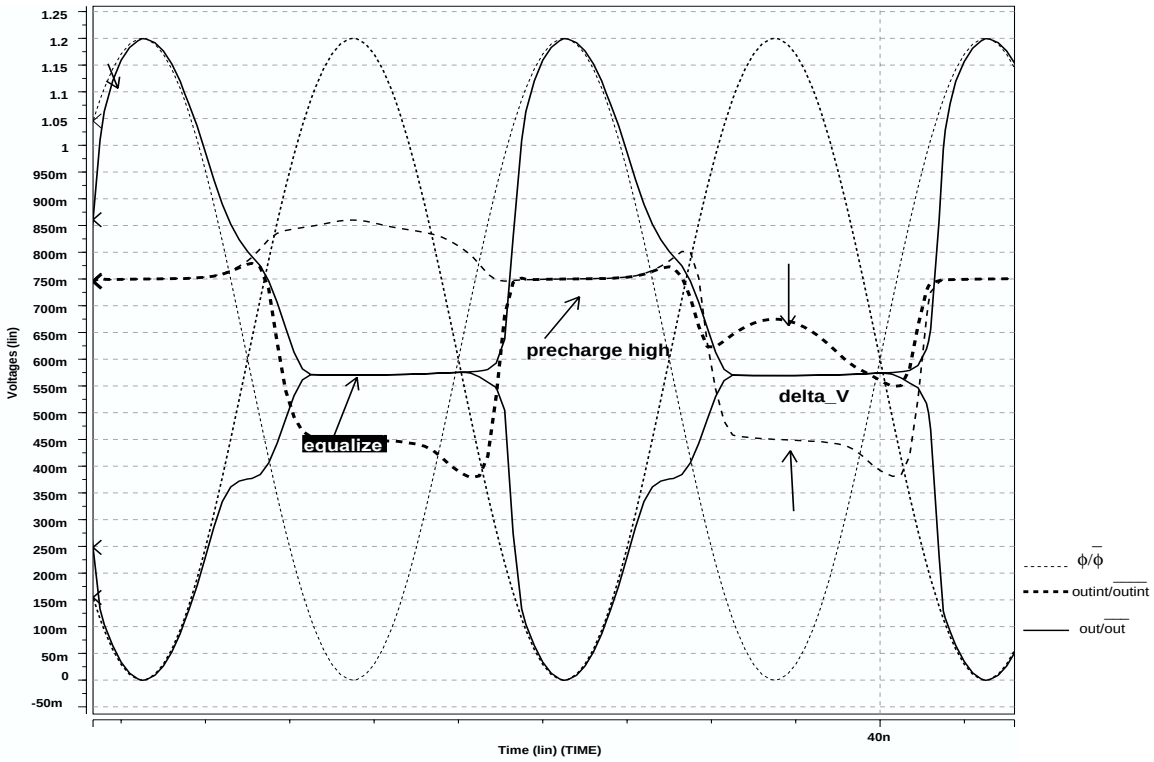

Figure 15: HSPICE simulation waveforms showing the operation of the Type 4 gate

6.3 Simulation Results

We simulated a number of circuits with this circuit topology; 2-NAND, 8-bit ripple carry adder and a 32-bit ripple carry adder. The frequency range of interest was 12MHz-200MHz. The circuit performs well in this range of operating frequencies. Voltage scaling was used to lower energy consumption at lower frequencies. The amplitude of the power clock, Vdd was varied from 1.2V down to 0.8V. The evaluation logic rails V'_{dd} and V'_{ss} were scaled according to the equations describe before, keeping V_c constant. Fig. 16 shows the results for a 32-bit ripple carry adder in this topology. The circuit is loaded with an inverter on each of the output lines.

6.4 Energy Dissipation

There are a number of factors leading to very low power dissipation in this circuit style. The charging of internal nodes of the evaluation tree to very small voltages causes small dissipation during evaluation. Charge transfer from the evaluation stage onto the sense circuit is done in 2 steps through low capacitance lines. This helps in reducing the drive required by the pull down stack to create the required voltage swing. During the Boost sense stage, the evaluation

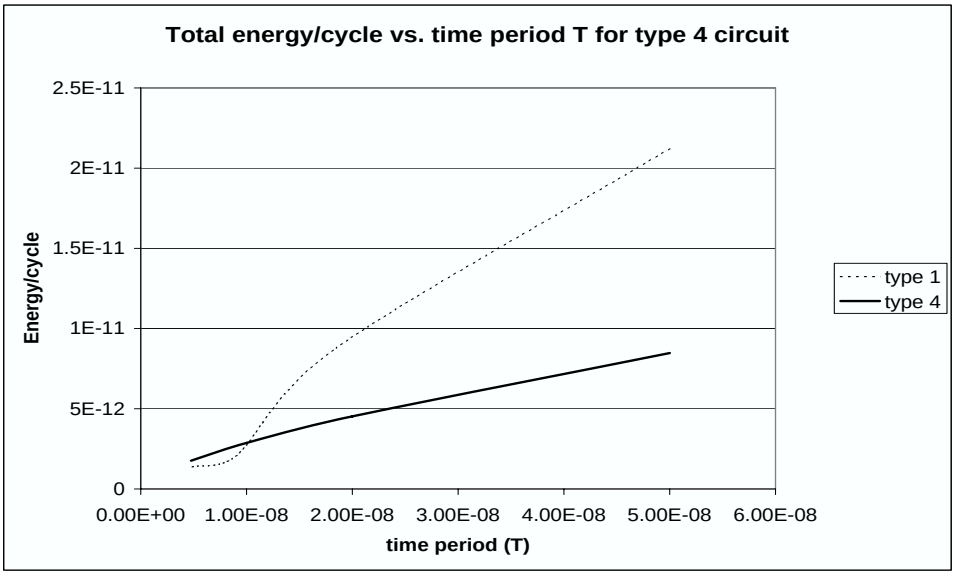

Figure 16: Energy dissipation of a 32-bit ripple carry adder in Type 4

stage circuit is virtually cutoff from the *out* and $\overline{out}$ lines, which in turn leads to lesser power dissipation during the sense phase.

As the frequency is decreased, the energy dissipated in the sense circuit (recovery) reduces according to $(RC/T)CV^2$. The energy dissipation in the evaluation circuit is initially a small contribution to the total energy. However, as T keeps on increasing, the crowbar in the evaluation stage does not keep increasing as $I_{crow} \cdot V \cdot T$, like in the previous cases. This is because there is no straight path from V'_{dd} to V'_{ss}. The total energy dissipation with this circuit is very less compared to the other circuit configurations, though slightly more than Type 2 (comparisons appear in Section 7). The circuit requires 7 additional transistors as compared to Type 1; but for a large gate with a high fan-in, this is lesser compared to the area overhead incurred in using cicuit Type 2 with 2 CMOS stacks. The energy dissipated can be calculated as shown in Fig. 17. In a given cycle, a gate will be in the energy recovery phase when ϕ is high. We consider 2 gates receiving the power clock in opposite phases (like when cascaded). Thus each gate is in the sensing phase. The ac current in the system is given by

$$i = jwC_{phase}v$$

where C_{phase} is the capacitance seen by the power clock per phase, w is the resonant frequency and v is the ac voltage. The energy dissipated in a cycle is given by

$$E = |i|^2 R_{eq} T$$

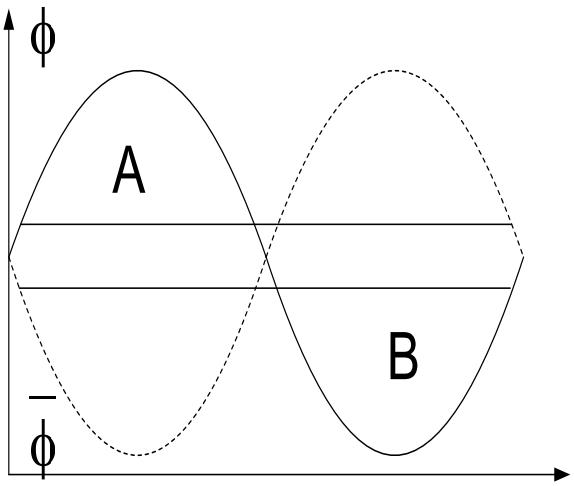

Figure 17: Waveforms for calculation of energy dissipated in Type 4 circuit per cycle

where T is the time period. R_{eq} is the resistance from the power clock to the capacitance. In this case, it is the R_{on} for the pull-up PMOS (either *M1* or *M2* in Fig. 14(b)) and R_{down} of the combination pull-down (*M3/M4* and *M5/M6* in Fig. 14(b)). The energy in a single gate is one-half of this. The ac voltage is calculated by considering the 2 crests - one of ϕ and the other of $\overline{\phi}$. It is approximated as a sinusoid and with an amplitude of $V_{dd}/2$ and centered at $3V_{dd}/2$. The ac voltage is given by $v = V_{dd}/4$. Substituting, we get

$$E_{recovery} = \frac{1}{2}(2\pi f)^2 C_{phase}^2 \left(\frac{V_{dd}}{4}\right)^2 R_{eq} T$$
$$E_{recovery} = \frac{\pi^2}{8}\frac{\tau}{T} C_{phase} V_{dd}^2$$

7 Comparison of Various Topologies

7.1 Energy Dissipation

In order to demonstrate the benefits of the energy recovery dynamic CMOS logic, we designed various circuits using the different circuit styles. In this section, we will compare the energy dissipation of the different circuit types. Fig. 18 shows a comparison of energy/cycle dissipation between the energy recovery domino CMOS and the other versions. The curves depict the total energy dissipation in the circuit including the clock generator. The circuit Type 3 has lower crowbar current in the evaluation stage compared to Type 1. Secondly, it uses only an inverter on the complementary side which reduces the overall capacitance and hence contributes to lower energy dissipation. The downside is the slow operation of the sub-threshold inverter and small voltage differential, susceptible to noise. The energy recovery domino has significantly

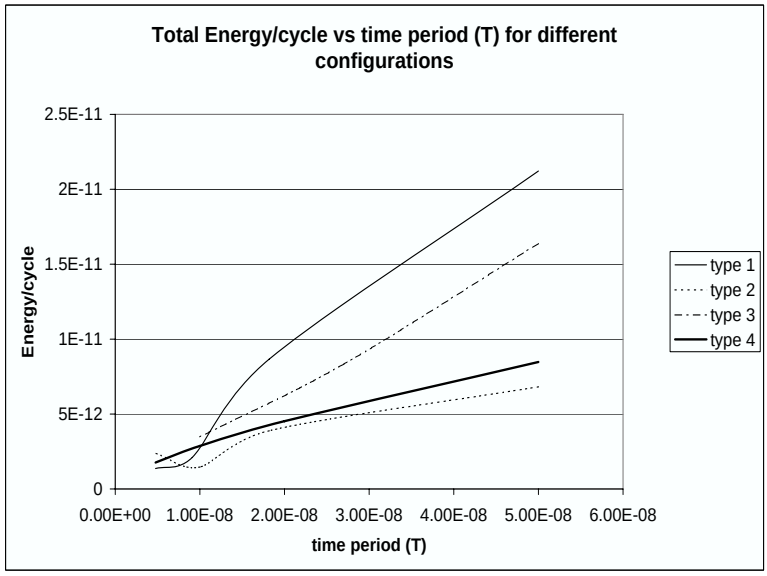

Figure 18: Comparison of energy dissipation between different circuit topologies

lower dissipation than the other two topologies in this frequency range. The reasons for this are the very low level charging up of internal nodes of the pull down stack in the evaluation tree, transfer of charge from the evaluation tree onto lower capacitance output lines and preventing any direct path from V'_{dd} to V'_{ss} during evaluation. The circuit has an area overhead of 3 transistors in the sense stage and 4 in the evaluation stage, which is much lesser than Type 2 or Type 3 for a gate with a large fan-in. The energy saving are around 65% to 25% in the range 20MHz to 150MHz over Type 1 Boost logic.

7.2 Variation in Delay due to Variation on Power Supply

We performed simulations to observe the impact of power supply variation on the delay of the domino structure. It can be seen from Fig. 19 that the percentage change in delay for the Type 4 circuit is much lower than regular CMOS. The Type 4 circuit is relatively more insensitive to variations in power supply.

7.3 Conformance to the Power Clock

A problem with the energy recovery domino topology is less degree of conformance of the output lines to the power clock waveform. This is shown in Fig. 20. When ϕ is low, the sense stage is in the equalization phase, and the outputs are held together by the equalization transistor. We are using $\overline{\phi}$ itself as an input to the transistor, instead of a separate out of phase clock pulse.

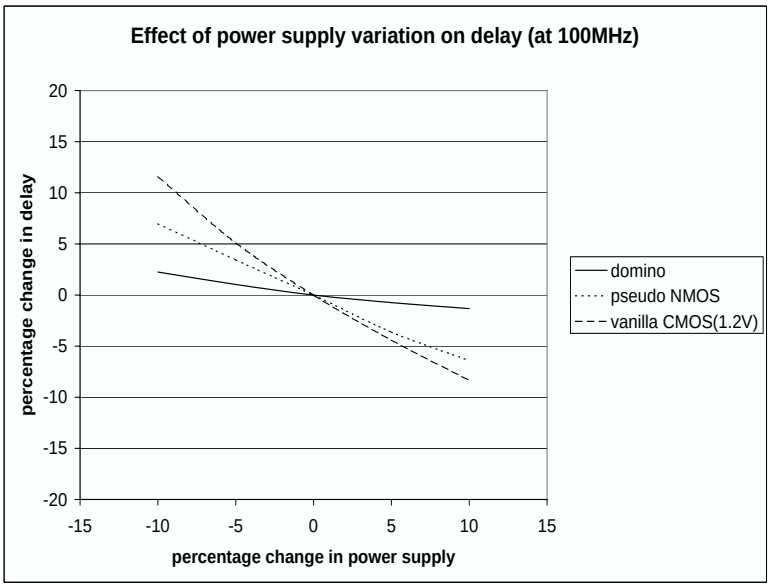

Figure 19: Change in delay with change in power supply

As can be seen from the figure, there is a delay when the equalization starts, during which the output deviate from the clock waveform. When ϕ goes high, the sense kicks in. However, the transistors that switch ON and transfer charge from the proxy output lines to the true lines have a delay associated with them. This again results in a delay before which the output lines start following the clock. This resulting deviation is the cause of higher energy dissipation, that would have been significantly decreased if the signal conformed perfectly with the clock. The delay increases with increasing frequency.

8 Conclusions and Future Work

In this report, we described the design of a energy recovery domino CMOS circuit built up on the idea of Boost logic. We performed experiments on several circuits and demonstrated the definite energy savings obtained. We performed extensive simulations on various circuits to ascertain the energy savings and robustness to power supply variation.

Our future work in this direction would involve improving the resonance performance of the circuit to achieve better conformance to the waveform. We need to compare results with a layout of the circuit with all parasitics and additional input/output buffers as required. The clock generator could also lend itself to a better design without the overhead of replenishing transistors, and this is to be looked at. We discussed various issues that arose with the design of the circuit and are exploring various alternatives in the design space.

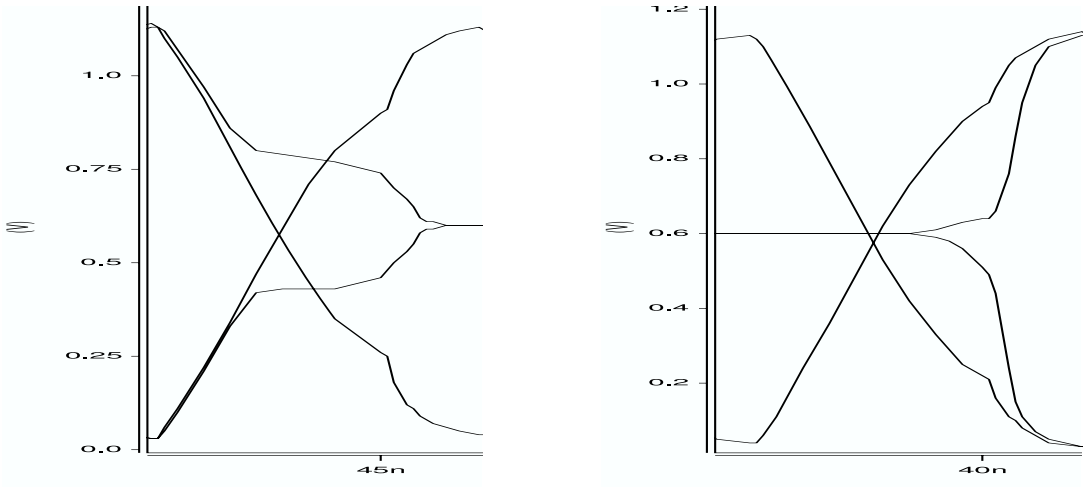

(a) equalization (b) switching into sensing

Figure 20: Deviation of outputs from the clock waveform

References

[1] V. S. Sathe, M. C. Papaefthymiou, C. H. Ziesler, *Boost Logic: A High Speed Energy Recovery Circuit Family*in IEEE Computer Society Annual Symposium on VLSI (ISVLSI) 2005.

[2] C. Ziesler, S. Kim, M. C. Papaefthymiou, *A Resonant Clock Generator for Single-Phase Adiabatic Systems*, ISLPED '01: Proceedings of the 2001 international symposium on Low power electronics and design, 2001. pp.159–164.

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

[4] Y. Moon and D. Jeong, A*n Efficient Charge Recovery Logic Circuit*, JSSC, April 1996.

[5] W. Athas, N. Tzartzanis, L. Svensson, and L. Peterson, *A low-power microprocessor based on resonant energy*, JSSC, Nov 1997.

[6] C. Ziesler, S. Kim, and M. C. Papaefthymiou, *Resonant clock generator for single-phase adiabatic systems*, in ISLPED, Aug 2001.

[7] 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.

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

[9] Alex G. Dickinson, John S. Denker, *Adiabatic Dynamic Logic*, IEEE Journal of Solid-State Circuits, Vol. 30, No. 3, March 1995. pp. 311–315.

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

[11] C. Ziesler, J.Kim, V.Sathe, and M.Papaefthymiou, *A 225 MHz Resonant Clocked ASIC Chip*, in ISLPED, Aug 2003.

[12] S. Kim, C. Ziesler, and M. Papaefthymiou, *A true single-phase 8-bit adiabatic multiplier*, in DAC, June 2001.

[13] A. Kramer, J. S. Denker, S. C. Avery, A. G. Dickinson, and T. R. Wik, *Adiabatic Computing with the 2N-2N2D Logic Family*, in 1994 Symposium on VLSl Circuits Digest of Technical Papers.

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

[15] D. Deleganes, M. Barany, G. Geannopoulos, K. Kreitzer, M. Morrise, D. Milliron, A. Singh, and S.Wijeratne, *Low-Voltage-Swing Logic Circuits for a 7 GHz x86 integer core*, in ISSCC, pp. 154-155, February 2004.

[16] J. Wood, T. Edwards, and S. Lipa, *Rotary traveling-wave oscillator arrays: a new clock technology*, JSSC, Nov 2001.

[17] W.C. Athas, N. Tzartzanis, *Energy Recovery for Low-power CMOS*, in the Proc. of the Conference on Advanced Research in VLSI, IEEE Press, Chapel Hill, NC, 1995.