AN ABSTRACT OF THE THESIS OF

Thomas L. Ruggeri for the degree of Master of Science in

Electrical and Computer Engineering presented on April 19, 2012.

Title: TIMR: Time Interleaved Multi Rail

Abstract approved: _____

Patrick Y. Chiang

This work presents a new energy saving technique for modern digital designs. We propose Time Interleaved Multi-Rail (TIMR) - a method for providing two dynamic supply rails to a circuit. This technique uses the first supply rail to mask the transition delay while changing the voltage of the second rail. We examine the design of TIMR as well as the implementation and considerations. We propose a number of control schemes that range from traditional DVFS to "race to sleep". This thesis also shows simulations of the technique using a existing voltage regulator in order to find the time and energy overhead of implementing the design. We find a $100\mu s$ switching time delay and $118\mu J$ energy overhead associated with changing the voltage rail. This work concludes with comparisons to current energy saving techniques.

TIMR: Time Interleaved Multi Rail

by

Thomas L. Ruggeri

A THESIS

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Master of Science

Presented April 19, 2012
Commencement June 2012

Master of Science thesis of Thomas L. Ruggeri presented on April 19, 2012.

APPROVED:

_____

Major Professor, representing Electrical and Computer Engineering


_____

Director of the School of Electrical Engineering and Computer Science


_____

Dean of the Graduate School


I understand that my thesis will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my thesis to any reader upon request.


_____

Thomas L. Ruggeri, Author

# ACKNOWLEDGEMENTS

# CONTRIBUTION OF AUTHORS

Ben Goska and Ryan Albright work extensively on previous work that developed into this project. They contributed ideas and insight that made this thesis possible.

TABLE OF CONTENTS

# TABLE OF CONTENTS (Continued)

## LIST OF FIGURES

# LIST OF FIGURES (Continued)

## LIST OF TABLES

# DEDICATION

For my lifelong mentor, Erman Mays. May we all live life as he has.

## Chapter 1 – Introduction

Digital electronics are a large part of recent popular culture. We continually see new products that amaze and inspire us to do new things. With each new product released, be it a phone, tablet or computer, the limits of what was previously thought possible are pushed further. These advances in technology are made possible by modern day digital design in embedded systems.

## 1.1   Embedded Systems

Embedded systems are present in everyday life in ways we can see and ways we cannot. The author of [19] defines an embedded system:

> An embedded system is a microprocessor-based system that is built to control a function or range of functions and is not designed to be pro-grammed by the end user in the same way that a PC is.

We can find embedded systems in a number of places including popular devices such as smart phones, tablets and TV's as well as places we would not think of such as cars, washing machines, toys and kitchen appliances [19]. These systems perform operations over and over during their lifetime often without ever having user direct input.

Embedded systems are comprised of a processor, memory, power supplies, periph-eral devices and software. Each of these components perform a critical role in the op-

eration of the system. The microprocessor is the brain of the design that performs all of the computation necessary to carry out the function of the system. The increasing performance of these microprocessors in past years has led to an increasing demand for computational power and energy efficiency [37].

## 1.2 Energy Efficient Design

The need for energy efficiency in digital designs is increasing. This has been fueled by increased demand for wireless multimedia devices [37]. The increase in health monitoring devices has also required energy efficient design [14, 33].

The contributors of [33] have identified that in order to develop wearable health monitoring devices, low energy designs must be realized to reduce the battery size. With energy savings comes a direct correlation to improved functionality of future bio-medical devices. [33] finds that digital systems will need to achieve medium computational throughput while operating at ultra-low energy. Achieving this will unlock a number of bio-medical devices that can be worn on the body and provide new information on our health.

Work in [14, 33, 37] cite the need for adaptive supply scaling to meet the requirements of the system while saving energy. In this work we will propose a new scheme for adapting supply voltages of digital circuits to achieve energy savings. We will review current energy saving techniques, propose our design for energy savings and show simulated overhead for the design.

## Chapter 2 – Energy in Digital Design

This chapter will review energy consumption and energy saving techniques in digital circuits.

## 2.1  Energy Consumption

With the increased capabilities and demand of mobile devices, there has been an increased need for longer battery life in such devices [37]. A longer battery life can be achieved through a reduction in the energy consumed by the digital components on the device. In order to lower the total energy consumption of a device we must first understand where and how the energy is being consumed.

Energy consumption in a digital circuit can be broken into two parts: dynamic ($E_{dynamic}$) and leakage ($E_{leakage}$). The combination of the two comprises the total energy consumption:

$$E_{total} = E_{dynamic} + E_{leakage}$$

We now must examine each type of energy to find where possible energy savings exist.

## 2.1.1 Dynamic Energy

Dynamic energy is the energy that is consumed through charging of the load capacitance of a circuit. This energy is consumed from the supply rail when the output of the circuit transitions from a low value (0) to a high value ($V_{DD}$). This can best be seen with a simple CMOS inverter example shown in Figure 2.1. When the input of an inverter goes from high to low ($V_{DD} \to 0$) the output will toggle from low to high. This transition requires that current from the supply rail passes through the PMOS transistor and charges the load capacitance, $C_L$.
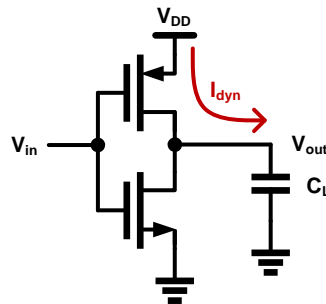
Figure 2.1: Output charging through PMOS transistor of a simple inverter.

When the input transitions from low to high ($0 \to V_{DD}$), there is no energy consumed from the supply because the energy being discharged through the NMOS is from the load capacitance, as seen in Figure 2.2.

The authors of [38] found that the total energy consumption from the supply during one transition in the following manner:
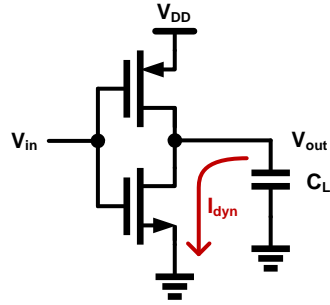
Figure 2.2: Current discharging through NMOS transistor of a simple inverter.

$$
\begin{aligned}
E_{V_{DD}} &= \int_0^\infty i_{VDD}(t)V_{DD}dt \\
&= V_{DD}\int_0^\infty C_L\frac{dv_{out}}{dt}dt \\
&= C_LV_{DD}\int_0^{V_{DD}} dv_{out} \\
E_{dynamic} &= C_LV_{DD}^2
\end{aligned}
$$

This equation shows that there are two ways of lowering the dynamic energy: decrease the load capacitance or decrease the supply voltage. Because energy has a quadratic relationship with supply voltage ($E_{V_{DD}} \propto V_{DD}^2$), the best way to decrease dynamic energy is to lower the supply voltage.

## 2.1.2   Leakage Energy

The other major component to total energy is leakage. Leakage energy is energy that is consumed by a circuit due to current leaking from the supply to ground. Because all

digital CMOS circuits are configurations of MOSFETs tied from the supply to ground, there is always a path of finite resistance from the supply to ground. This finite resistance is very large, but it is never infinite. As a result, there is a small leakage current through the circuit at all times $\left( I_{leak} = \frac{V_{DD}}{R_{trans}} \right)$. Leakage energy tends to increase as the feature size of the transistor is decreased making leakage energy an increasing problem. This leakage energy can be seen in the simple inverter in Figure 2.3.
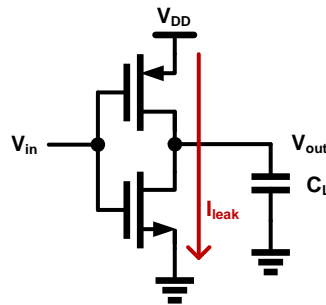
Figure 2.3: Leakage current through a simple inverter.

The authors of [38] found the leakage energy to be:

$$E_{leakage} = P_{leak} \times t = I_{leak} \times V_{DD} \times t$$

Again, the main way to reduce leakage energy is to lower the supply voltage ($V_{DD}$); however, there are additional techniques for lowering leakage that we will see later.

## 2.2   Energy Saving Techniques

There are multiple techniques for energy savings at both the circuit and architectural level. This section reviews some of the more common techniques.

## 2.2.1   Dynamic Voltage and Frequency Scaling

While running a circuit at its maximum possible frequency and nominal supply voltage provides the fastest possible operating conditions, energy savings can be achieved by operating the same circuit at lower values of frequency and supply voltage. This can be seen from previous analysis of the total energy consumed in a circuit per clock cycle:

$$
\begin{aligned}
E_{total} &= E_{dynamic} + E_{leakage} \\
&= \left( C_L V_{DD}^2 \right) + \left( I_{leak} V_{DD} t \right)
\end{aligned}
$$

As the supply voltage is lowered, the dynamic energy consumption is reduced quadratically and the leakage energy is reduced in a linear fashion. The same circuit operating at less than nominal supply voltage will require a longer time to complete computations (have a stable output at logical '0' or '1') resulting in a slower operating frequency.

The general technique that achieves these energy savings is Dynamic Voltage and Frequency Scaling (DVFS). Under this traditional approach, the voltage and frequency are dynamically reduced whenever the throughput of a the given circuit permits. The challenge with this technique is the control scheme for selecting when and to what level to switch. There are a number of proposed control schemes such as those based on workload [30], leakage current [23] and thermal sensors [29]. While the management of DVFS may differ from design to design, the principle remains the same: lower the voltage and frequency whenever possible to save energy. This theme will reappear in other more complex energy saving techniques.

### 2.2.2 Clock Gating

Clock gating is a technique that reduces energy of timed circuits by turning off the clock signal to circuits that are not active. This attempts to save the consumption of dynamic energy in a sequential circuit by removing the clock, which is the signal to the circuit to refresh [17]. If there is no change in output, there is no need to refresh and thus energy will be saved by removing the clock signal for that cycle. As seen in Figure 2.4, when the clock enable signal ($clk\_ena$) is set to 0, the clock signal is terminated before it enters a sequential circuit, thus eliminating the energy consumption from changing the output on the clock edge.
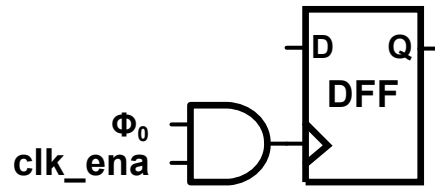


Figure 2.4: Clock gating of a D-Flip flop.

There are a number of potential issues with this method. One is the introduced overhead and complexity involved with determining when and where to gate the clock [17]. There have been attempts to improve upon the efficiency of clock gating controls such as [6]. Another shortcoming of traditional clock gating is its locality. Clock gating a small number of sequential circuit can save power in that circuit, but the entire clock tree is still oscillating. A larger energy savings can be achieved if more of the root clock structure is gated as was shown in [43]. This method attempts to gate the clock closer to the source and thus save more energy.

Clock gating is a fundamental technique for reducing energy consumption in sequen-

tial circuits. It can have a complex control overhead, but the results can be significant if there are large sections of sequential circuits such as in a design with a large pipeline depth.

### 2.2.3   Sub-Threshold Operation

Power savings can be realized by operating a circuit at a lower supply voltage as was seen in DVFS. These power savings need not be dynamic; the supply voltage of a circuit can be lowered statically to provide a constant energy savings. If the supply rail is close to the threshold voltage, $V_{th}$, of the CMOS transistors that compose the circuit it is known as operating in the "near-threshold" regime. If it operates below the value of $V_{th}$ then it is said to operate in "sub-threshold" regime.

Conventional first-level analysis of a CMOS transistor indicates that there are three regions of operation for the transistor to operate:

$$
\begin{aligned}
V_{GS} - V_T < \quad & 0 \quad (cutoff) \\
V_{GS} - V_T < \quad & V_{DS} \quad (triode) \\
V_{GS} - V_T > \quad & V_{DS} \quad (saturation)
\end{aligned}
$$

This would indicate that operating with a supply voltage less than the threshold voltage would result in non-functioning logic, however, this is not the case as current still flows due to weak inversion of the channel [18]. This allows for operation of the

circuit all the way to the theoretical minimum of $\approx 36mV$ (found in [31]). While the circuit will still function at this low supply voltage, there are a number of issues that arise with this regime of operation.

The first, and perhaps most important issue with sub-threshold operation is the delay induced by operating in this regime. Because the current charging the output of the circuit is only from a small weak inversion, there is very little current to charge the output capacitance and thus it takes a long time for the circuit to complete its operation. This increased delay means that the maximum clock frequency of a system using this circuit must now be lowered to account for the delay. A decrease in clock frequency cannot be tolerated in all applications which limits the practicality of implementing sub-threshold circuits.

While the power of sub-threshold circuit is significantly less than that of a super-threshold circuit, the increased time delay leads to a diminished return on lowering energy. This is mostly because of the leakage energy that dominates at super low supply voltages. The increased time delay allows for a larger window where leakage current is burned in the circuit. This becomes significant at supply voltages less than 300mV where the total energy is almost entirely leakage [18].

A final consideration of sub-threshold operation is the risk of process variations. Process variations such as slow or fast devices, random dopant fluctuation and body biasing can result in variation of the performance of circuits on the same die (intra-die) and across dies. These effects can also be systematic or dynamic in nature. Systematic variations are those that present a constant offset to the performance while dynamic variations depend on the operation that the circuit is performing. All of these variations

are present in the circuit regardless of what region the supply voltage is in, but the effects of variation are significantly increased when operating in sub-threshold. This has been found in empirical works [18] as well as in real world designs [34].

There are a number of circuit and architectural level techniques that have been proposed to combat some of the issues with sub-threshold operation. One of the most well-known techniques is Razor [16] which is an error detection circuit that can identify when errors have occurred and can assist a designer in recovering from such errors. Detecting and correcting errors can allow for a circuit to operate at sub (or near) threshold voltage where more errors will occur [13]. Other techniques for improved performance in sub-threshold include micro-rollback, tunable replica circuits (TRC) and modular redundancy [13].

### 2.2.4   Voltage Islands

We have seen that a circuit can benefit from different supply voltages, such as in DVFS and sub-threshold. Voltage islands partition a circuit layout based on supply voltage needs. With this technique, designers identify supply voltages for portions of the circuit, then route the floor plan of the design to partition circuits with the same supply voltage onto the same "island" [35]. An example floor plan with partitioned voltage islands can be seen in Figure 2.5. Voltage islands provide power savings by operating circuits at near ideal supply voltages for their desired performance per power.

There are also a number of challenges in designing voltage islands. The choice of floor plan, the physical layout of the circuit on a silicon die, can make or break the
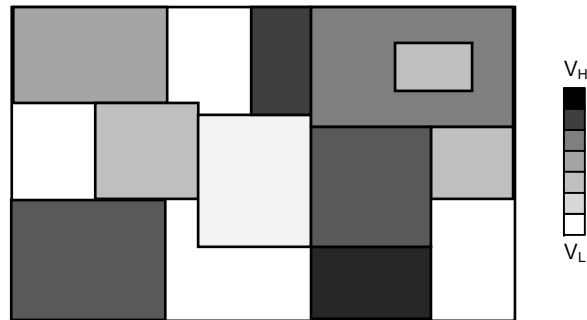
Figure 2.5: An example floor plan with different voltage islands (shown in shades of gray).

potential energy savings. The authors of [20] propose a unique routing algorithm that attempts to optimize energy and die area. Work done in [40] provides an algorithm for voltage assignment and floor planning based on a given application. There is also a large consideration for the level shifters that translate signals between islands. Work in [35] analyzes the need for an efficient converter between high and low voltage islands.

While the benefits of voltage islands can be great in terms of energy savings, there is a large amount of complexity in the implementation of the technique.

### 2.2.5 Power Gating

When a combinational circuit is not in use, the most prevalent energy that it is consuming is leakage energy. Power gating is a circuit level technique that attempts to eliminate this leakage energy by placing a PMOS transistor in between the supply rail and a virtual supply [32]. A simplified version of this circuit can be see in Figure 2.6. A PMOS power gate can be turned "on" or "off" to disconnect the supply from the circuit thereby eliminating leakage energy.
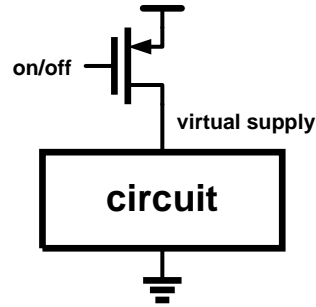
Figure 2.6: A simple power gate implemented with a PMOS transistor.

While power gating can significantly reduce energy consumed during idle periods of operation, there are a number of considerations for this technique [24], the first of which is the sizing of the device. If a minimum sized PMOS transistor is used then the current through the device will cause the circuit to operate slower than the nominal case. The current through a PMOS device in saturation is modeled by [38] as:

$$I_D = \frac{k_p'}{2} \frac{W}{L} \left(V_{GS} - V_T\right)^2 \left(1 + \lambda \left(V_{DS} - V_{DSat}\right)\right)$$

This equation shows that as the width of the PMOS power gate is reduced so is the driving current that can be delivered to the virtual supply and thus the circuit. Because the transistor exhibits an "on" resistance, a minimum channel length transistor will have a high resistance and thus a large voltage drop between the supply and virtual supply [2]. For both of these reasons, the size of the PMOS power gate should be large.

By making the PMOS transistor large there are negative trade offs [2, 39]. One of these is the increased delay in turning the MOSFET "on" or "off." The increased size requires a large amount of energy to charge (or discharge) the gate and thus it takes longer to transition states than with a minimum sized gate. Another negative is the

increased area required for a larger gate. Making a very large gate takes up space that could otherwise be used for logic.

The result of all of these considerations is a trade-off between power and performance [24, 2]. Large power gates will take up die space and operate slowly but provide maximum power savings while small power gates allow for faster operation at the cost of energy savings. These trade-offs have no golden ratio; it is up the designer of the circuit to determine what is optimal for each circuit block.

### 2.2.6 Multi-Rail Supplies

In a similar fashion to voltage islands, multi-rail supplies attempt to supply parts of a circuit with different supply voltages, however, the implementation of this technique is more dynamic through the use of power gates. As seen in the example in Figure 2.7, multi-rail supplies are connected to a single virtual supply through multiple power gates. This technique achieves energy savings with a moderate control overhead.

In a multi-rail design there are two or more power rails routed to each subcircuit. Each subcircuit also has a virtual supply rail that is connected to all of the power rails through an individual PMOS power gate. The control system must decide which of the power rails the subcircuit should be connected to while disabling the others. The result is a circuit that can operate at a supply voltage that best suits the needs of the system: high, medium or low in this example.

While the power savings of this approach can be great, there are are also a number of considerations. The authors of [15] find that there are two main factors to the supply
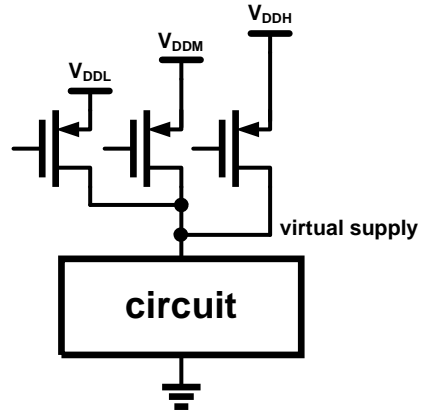
Figure 2.7: Three power rails (low, mid and high) connected to a circuit through three PMOS power gates.

choice: switching delay and switching energy. The switching delay is the amount of time that it takes to transition from one rail to another. In [15], the authors determine that the delay can be modeled with an RC delay that results in:

$$t_{switch} = -ln\left(1 - \frac{V_H + V_L}{2V_L}\right)\tau$$

where

$$\tau = \frac{2C_oW + C_{ox}WL + 2\lfloor C_jL_SW + 2C_{jws}L_S \rfloor + C_L}{|\mu_p|C_{ox}\left(\frac{W}{L}\right)\left(|V_{GS}| - |V_{tp}| - \frac{|V_{DS}|}{2}\right)}$$

Notice that the delay is a function of the difference between the supply rail voltages and PMOS power gate widths. As was seen with a single power gate, making the power gates larger will dramatically decrease switching delay at the cost of area.

The second factor in supply choice is the switching energy. With each transition between supply rails there is some energy consumption in the control overhead. This energy comes from charging and/or discharging the gates of the power gate PMOS'.

Finally, the authors of [15] find that there is a break even point where more energy was saved than expended in lowering the supply voltage. They found this point to be:

$$N_{BE} = \frac{(E_{high} - E_{low})}{E_{switch}}$$

The authors then find that this break-even number is less than $1$ for their test adder and multiplier from $500mV - 900mV$. This means that there is no energy penalty for switching from a higher to lower rail and it is always advisable. The concept of switching delay ($t_{switch}$) and switching energy ($E_{switch}$) are two key components in any multi-rail design.

### 2.2.6.1 Multi-rail adaptations

Panoptic Dynamic Voltage Scaling (PDVS), introduced in [36], offers the benefits of multi-rail with the advantages of architectural level control. By scheduling the voltage level based on knowledge of upcoming instructions, Panoptic DVS is able to achieve energy savings at a medium ($40\% - 60\%$ of maximum) workload. This is done by operating at the energy optimal voltage for units with time insensitive tasks. Take an example seen in Figure 2.8 where an adder unit can be delayed to save energy.

In order to reduce the switching energy, the authors of [25] implement a unique control scheme they call Stepped Supply Voltage Switching. This technique uses the intermediate (or medium) rail when switching from low to high or high to low. With this small change in overhead, the authors claim that they achieved between $15\%$ and $50\%$ switching energy savings.
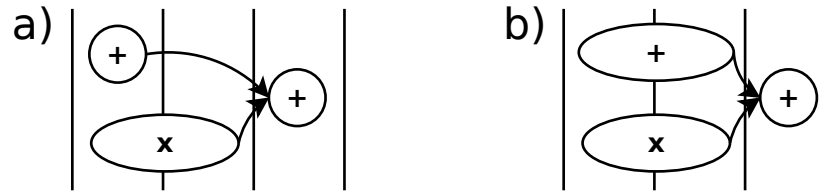
Figure 2.8: An example of Panoptic DVS. By lowering the first adders supply voltage and using slack time while the multiply completes, energy is saved in case (b) compared to (a) without an impact on execution time.

Work in [41] implements a single voltage supply, multi-supply and PDVS on a single 90nm test chip. They also implement a sub-threshold voltage level for comparison based on work in [5]. The authors compare the results in terms of energy savings and area overhead.

### 2.2.6.2 Voltage dithering

Work in [5] showed a multi-rail design with two rails, $V_{DD_L}$ and $V_{DD_H}$. The interesting usage of the multi-rails that this work introduces is local voltage dithering (LVD). With this approach, the control signals can be used to not just select one voltage, but it can also dither between the supply voltages to reach the energy optimal $V_{DD}$. Work in [1] also implement voltage dithering.

## Chapter 3 – Proposed Design

In this chapter we introduce the Time Interleaved Multi-Rail (TIMR) design for energy savings. We also show implementations of the design with on and off chip power supplies and propose a number of control schemes. The chapter concludes with considerations for the design.

## 3.1   TIMR: Time Interleaved Multi Rail

We have seen the need for energy reduction in Chapter 1. We have also seen a number of energy saving techniques from Chapter 2. Here we introduce Time Interleave Multi Rail which reduces energy by managing multiple power rails with an adaptive power supply control. This method attempts to offer the energy savings seen in both Dynamic Voltage and Frequency Scaling (DVFS) and Multi-Rail Supplies (Section 2.2.1 and Section 2.2.6 respectively) without the need to stall the execution flow as in DVFS or operate at a static supply voltage as in Multi-Rail Supplies.

The methodology of TIMR is a combination of DVFS and Multi-Rail Supplies to get the benefits of both without the overhead of either. DVFS offers the ability to operate a circuit at an energy optimal supply voltage provided that the control overhead exists and the circuit can tolerate stalling while the voltage regulator that feeds the supply line changes from one voltage to another. Multi-Rail Supplies offer two or more static

supply voltages to operate with a very small switching time, but it does not offer the ability to dynamically change the voltage level of these supplies. As such, a designer must currently choose between fast switching times or achieving a fine granularity of supply voltages. TIMR will allow for both fast switching times and fine granularity between supply voltages.

TIMR uses the physical design of Multi-Rail Supplies to provide each circuit with two power rails. As seen in Figure 3.1, two power rails are routed to each circuit and connected to a virtual supply via PMOS power gates. These power gates serve the same function as those in traditional power gates, seen in Section 2.2.5, to connect or disconnect the power supply from the virtual supply that sources the circuit. The TIMR design does not use the power gates to implement voltage dithering as seen in Section 2.2.6.2, so the control lines to the gates of the PMOS headers are boolean $0$ or $V_{DD}$ signals.
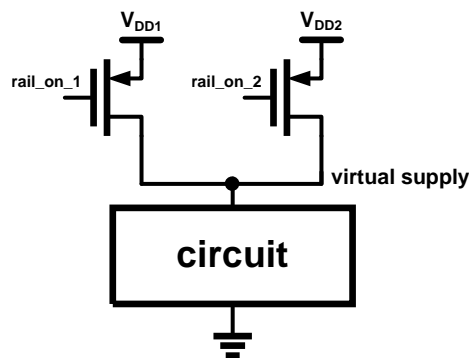
Figure 3.1: A simplified schematic of Time Interleaved Multi Rail.

The supply voltage on each rail is not static, as with Multi-Rail Supplies, rather it varies with the circuit's needs and the energy optimal point as is seen in DVFS. The

typical operating condition is as follows: the first rail ($V_{DD1}$) is connected to the virtual supply by setting the control signal $rail\_on\_1$ to a zero and thus turning on the PMOS power gate as seen in Figure 3.2. While the first rail is actively sourcing power to the circuit, the second rail ($V_{DD2}$) is disconnected from the virtual node as signal $rail\_on\_2$ is tied to $V_{DD}$ and the PMOS power gate is off. Because the rail is not active, it may undergo a voltage change without an induced time penalty on the circuit.
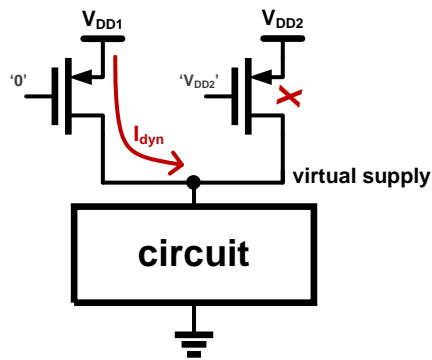


Figure 3.2: TIMR circuit operating with no change in supply voltage.

If the power management hardware determines that the circuit benefits from operating at a different supply voltage (higher or lower than $V_{DD1}$) it signals the voltage regulator to change its output voltage on rail two ($V_{DD2} \rightarrow V'_{DD2}$) as seen in Figure 3.3. While the regulator performs the change in voltage, the circuit continues to operate normally (no stalls are introduced) because it can continue to operate on the first supply. When the second supply has settled to its new voltage level, the power management hardware toggles the values of $rail\_on\_\{1, 2\}$ and the circuit transitions to use the second rail at the energy optimal supply, seen in Figure 3.4.

This approach allows the circuit to transition to its energy optimal supply voltage
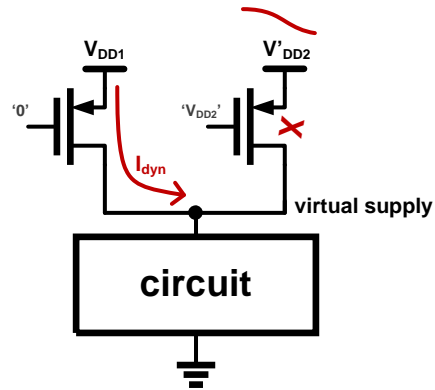
Figure 3.3: TIMR circuit changing voltage of second rail.

(within the constraints of the power supply, as we will see later) without incurring the stall time delay seen in DVFS by continuing to operate on the first rail while the transition occurs on the unused second rail. Because the switching time for the power rails is much less than the stall time [15], much less operating time is wasted in changing supply.
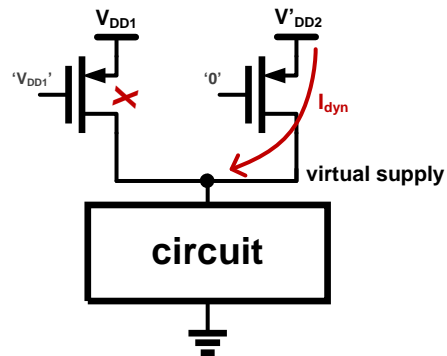


Figure 3.4: TIMR circuit switching to second supply rail.

## 3.2    Implementation

Now that TIMR has been introduced, we will examine some implementations of the design. These implementations include the use of off and on chip voltage regulators and some of the possible control schemes.

### 3.2.1    Off-Chip Regulators

One major component to TIMR is the power supply. Most modern day embedded systems consist of a central processing chip and one or more power supply chips that exist on the same board. This configuration, shown in Figure 3.5, is known as off-chip regulators. These regulator's size and topology depends on the needs of the embedded system, but are often "buck" (or DC-DC) supplies for the main processing core [42]. The choice of supply can have a major impact on the implementation of TIMR as a number of factors affect the design. These factors include voltage level granularity, control interface, pad requirements, loading capacitance, settling time and feedback stability.

A system using TIMR will need to consider the following factors as part of the design process. The power supply control interface is a critical part of the methodology of TIMR. The power management hardware must be able to control the voltage level of the regulator through either the regulator itself or using a VID controller [7, 8]. The control lines, along with the two power supply lines for each block, must be considered when laying out the chip and the accompanying printed circuit board. This may lead to an increased number of pads on the chip.

The loading capacitance is another factor that plays into the design. As seen in Fig-
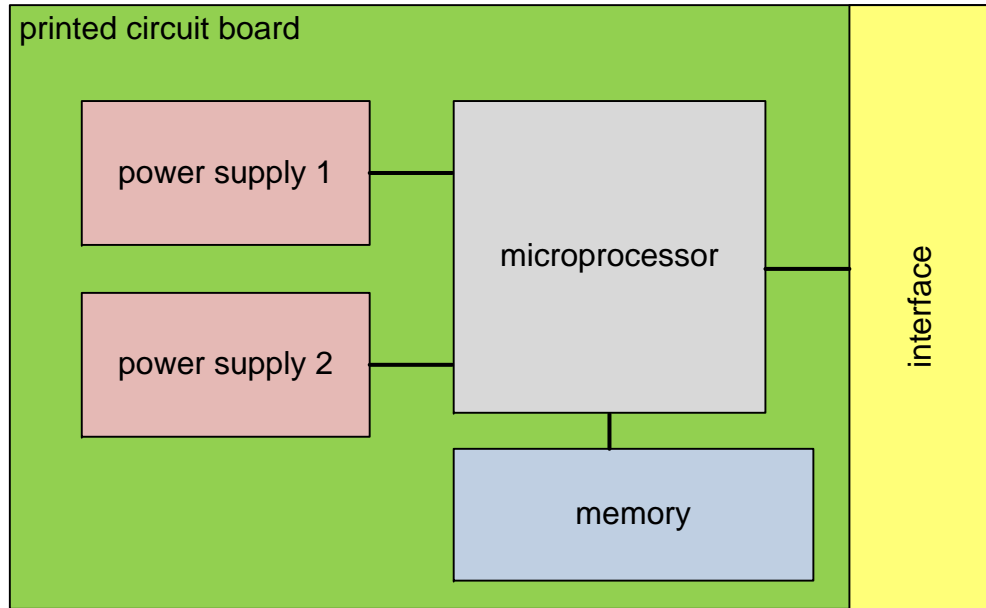
Figure 3.5: Typical off-chip power supply configuration.

ure 3.6, the design of a typical power supply circuit includes a somewhat large capacitor, $C_L$, which acts as a decoupling capacitor and filter for the power supply rail. The sizing of this capacitor is dependent upon the power supply choice as well as the load, desired ripple and output transient response [21, 22]. Sizing this capacitance, and the proper choice of capacitor type, will impact the performance as the settling time and stability of the power supply will be affected.

While there are a number of design choices that must be balanced, the use of an off-chip regulator for TIMR can be very beneficial. Because separate regulators are common in an embedded design, very little would need to be done at an architectural, layout or board level to add support for TIMR. It is not uncommon for regulators to have multiple rails available on a single chip [21], so only one part would be necessary for
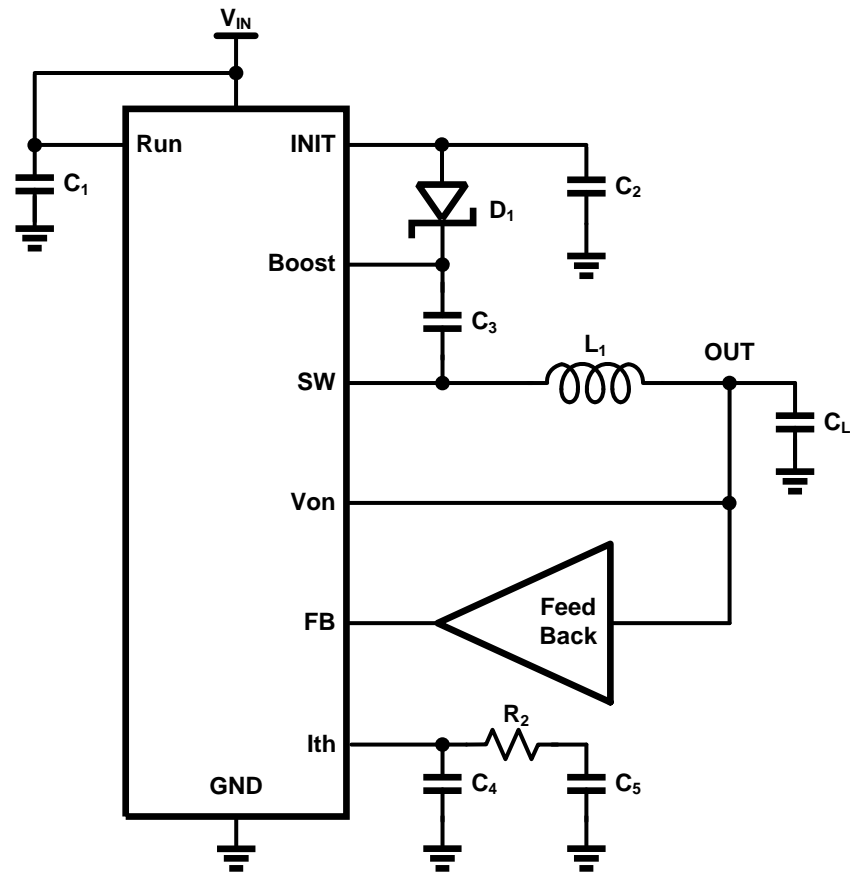
Figure 3.6: Typical off-chip power supply schematic [9, 10].

the TIMR rails.

## 3.2.2 On-Chip Regulators

Another option for power supplies is an integrated on-die regulator. While taking up silicon space on die is expensive (using space that could be used by other circuitry), the integration of the voltage regulator on-die removes the need for an expensive item off-chip as the voltage regulator is often one of the most expensive items on the bill of

materials (BOM) in an embedded system. The use of an on-die regulator also removes a number of pads off the chip (both power supply and control lines). The integration offers faster speeds in both communication and changes in voltage as the on-die regulator has shorter lines that take less energy to drive and can thus be done at a faster rate [38].

This approach can only be done if a practical voltage regulator exists that can meet the requirements of the system. This is becoming more viable option as shown by a number of research works. In [26] the authors demonstrate a DC-DC power supply that focuses on fast switching for Dynamic Voltage Scaling. This work builds a test chip in 130nm CMOS that can provide up to $1W$ of output power and operate at voltages as low as $400mV$. This is ideal for the TIMR design as it could operate in the sub to near-threshold regime. This proposed design attempts to combine both a buck regulator and a switched capacitor design.

Work done in [28] presents another option for an on-die regulator done in 28nm technology, a very modern technology that is currently being used for digital designs. The work presents a design that is capable of operating at sub $1V$ level and produce more than $500mW$ of output power. This regulator could be used for TIMR on a large circuit within a modern digital design.

The authors of [27] present a design that is capable of up to $90\%$ efficiency while operating in either buck or boost mode. This design also has dual outputs which would be ideal for a TIMR design as the same regulator can provide both power supply rails to the target circuit. The proposed design is interesting because it does offer a boost capability so it could be useful for an embedded system that has the need for low input voltage operation (such as a voltage harvesting system). In such a case, the input voltage

could be boosted for operation.

While on-die regulators are not currently commonly used in modern digital systems, their use may become more prevalent in the future. This can be seen as the technology is developing and constantly being improved upon. There are a number of proposed designs that could be used for a system designed with TIMR [26, 27, 28]. These designs offer a number of benefits over off-chip designs that include faster switching times, fewer off-chip lines and components and lower part costs for board designers.

### 3.2.3   Control Scheme

The final part of implementation of TIMR is the control scheme. Once the hardware exists and the power supply has been designed for the system, the last piece is controlling the two power gate transistors along with coordinating the changes of voltage on the inactive rail. Here we propose a number of possible control schemes including one based on traditional DVFS, one that aggressively saves energy and one that is based on "race to sleep".

### 3.2.3.1   Traditional DVFS Based

This control scheme attempts to mimic traditional DVFS. When the processing load of the circuit is found to be less than full, the circuit will attempt to operate at a lower voltage and clock frequency that still finishes the computation on time but does so with less energy. The method by which the computational throughput is measured need not

be specific to this scheme; rather, it will be an architectural or circuit level choice. Load monitoring can be done by thermal sensor, performance counter or energy/power monitor. Once it is determined that the circuit is operating with some slack (the throughput is less than maximum and the circuit is completing the required work with time to spare each cycle) then the voltage can be lowered to save energy.

With this control scheme, TIMR would perform much the way traditional DVFS does because it follows the same algorithm to apply voltage changes. The difference will be in the speed that the voltage switches can occur. As seen in Figure 3.7, traditional DVFS will stall the execution while the voltage on the supply rail changes, whereas TIMR does not need to stall the execution because the change occurs on the second rail. Notice that there is a small time delay to switch the power gate, but this is very small (on order of nanoseconds) when compared with the power supply switching (on order of hundreds of microseconds). Because little execution time is lost in making a supply transition, only the timing and energy constraints of TIMR limit the number of transitions that can occur as can be seen later.
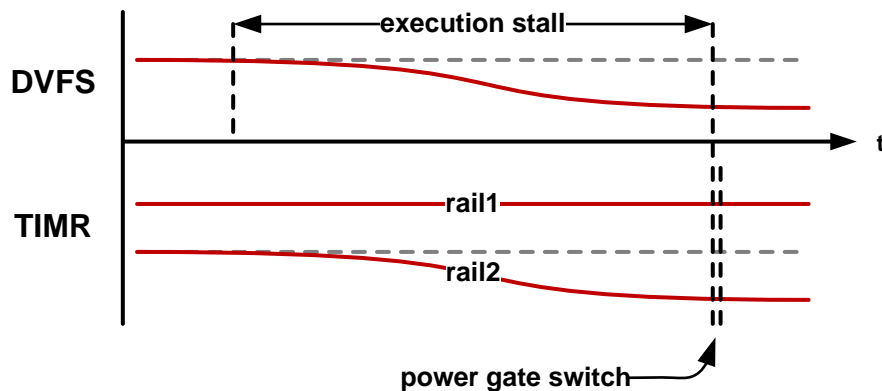


Figure 3.7: Supply rails of both traditional DVFS and TIMR.

### 3.2.3.2    Energy Aggressive

While the Traditional DVFS scheme does offer moderate energy savings, it is possible to operate TIMR with a much more aggressive approach towards energy savings. While DVFS moves to a lower supply voltage when sensors or control indicate that slack exists, Energy Aggressive attempts to operate at a lower voltage without any indication of slack and then report if the resulting operating conditions are too low. This can be done in a number of different ways. We will propose three techniques in this section: Bottom up, Top down and Aggressive rail.

**Bottom up**

This approach starts out at the lowest supply voltage possible and then checks the throughput sensors (thermal, performance counters, energy monitors, etc.) to ensure that the new voltage meets the throughput requirements of the circuit. Take the example in Figure 3.8: the voltage starts at its lowest level on both rails and rail two immediately goes up one level of granularity to accommodate the next step up. Because the throughput sensors report that the throughput level is acceptable (simplified in this diagram), the supply stays at the lowest level. As soon as the sensors report an unacceptable throughput the supply transitions to rail two. With this scheme, an extreme importance is put on energy savings as the control attempts to go as low as possible and then move the supply up as necessary.
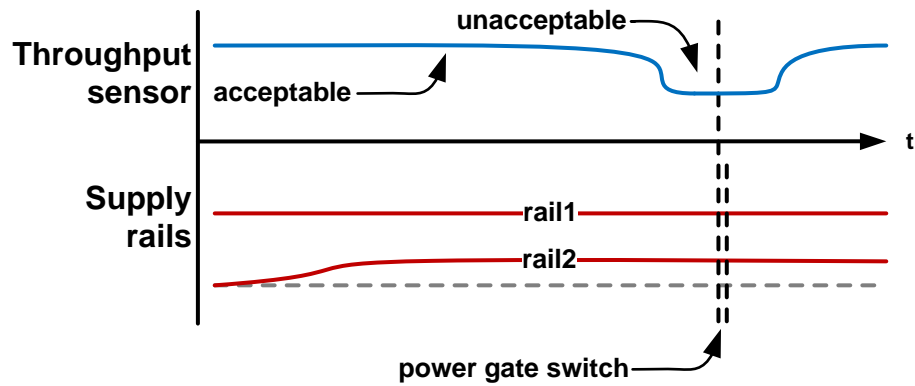
Figure 3.8: Supply rails and throughput sensor with TIMR under bottom up control scheme.

### Top down

This approach is very similar to Bottom up with the supply voltage starting at the nominal value and then moving downward. This scheme will immediately move down by one voltage step and then rely on sensors to check the throughput as seen in Figure 3.9. If the throughput level is still acceptable, then the voltage scales down by another step. This approach puts importance on energy savings, but does not sacrifice throughput to achieve it.

### Aggressive rail

One thing that has not been seen to this point is a control scheme that uses the time interleaving of the rails to pursue energy savings. Aggressive rail is a method that moves one rail to a lower voltage followed directly by the next rail. As seen in the example in Figure 3.10, the voltage on rail two is lowered. The system switches to rail two as soon as the voltage stabilizes on it. Once the circuit begins to operate on rail two, rail one immediately lowers below rail two. This method moves from the highest voltage down

Figure 3.9: Supply rails and throughput sensor with TIMR under top down control scheme.

as in Top down without checking that the throughput is acceptable after each transition. This means that Aggressive rail can suffer from over shooting the energy optimal supply voltage.

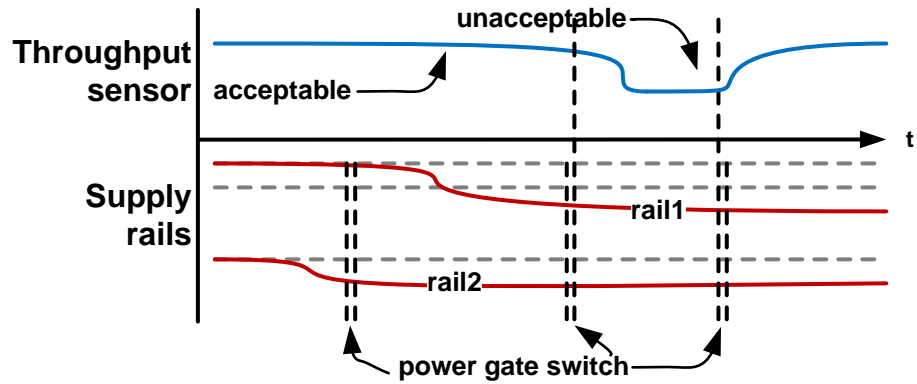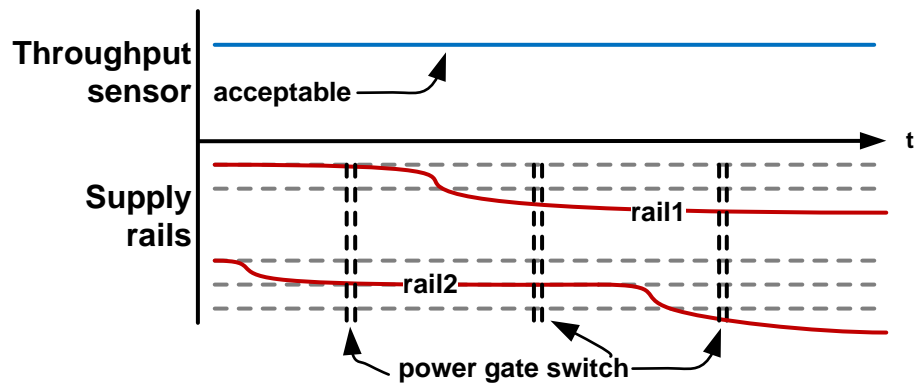

Figure 3.10: Supply rails and throughput sensor with TIMR under aggressive rail control scheme.

### 3.2.3.3  Race to Sleep Based

The final control scheme that we propose is one that implements race to sleep. Race to sleep is a method that operates a circuit at its fastest possible settings and then goes to sleep when the computation has completed [4]. Race to sleep attempts to save energy by avoiding numerous power state changes and spending a maximum amount of time in sleep where minimum energy is consumed. It has been a well known power saving technique in embedded systems. Because race to sleep focuses on minimizing power state transitions, TIMR does not seem like a good fit to implement race to sleep.

One difficulty in practically implementing race to sleep is the granularity of the operations to be completed before going to sleep. If fine grained operations can be selected for this method, then a greater energy savings can be achieved. As seen in Figure 3.11, a traditional example performs an entire program (functions 1-5) and then goes to sleep. This means that the each function of the program operates at a single supply voltage. With a TIMR approach, each function of a program could operate on a much finer grained race to sleep. The functions could operate with their optimal supply voltage (found through simulations or trial and error) and then sleep in between function calls. This approach could allow for energy savings especially if there were long waits for cache or memory access that could take advantage of TIMR based race to sleep to save energy.
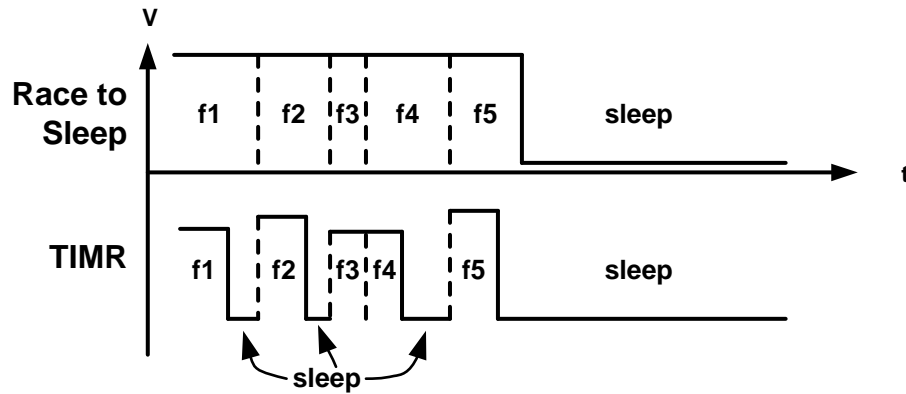
Figure 3.11: Function execution (1-5) under traditional race to sleep and TIMR based race to sleep. Note that function 3 (f3) is too short to go to sleep based on the switching time of TIMR and each function can now operate at a unique supply voltage.

## 3.3 Considerations

Along with the implementation choices that are made with TIMR, there are also a number of design considerations that must be addressed. These considerations include the finite power supply switching time, the loading of the power supply and the control overhead introduced by the TIMR circuitry.

### 3.3.1 Power Supply Switching

When there is a desired level change on one of the power supply rails, the control hardware of TIMR must communicate to the voltage regulator (either on or off chip) the new output voltage level. Once the power supply receives this signal it will begin changing the output voltage. The time from when the output voltage begins to transition from its original value until it settles to its new desired level is the switching time, $t_{switch}$. The

timing of this sequence can be seen in Figure 3.12. This time is affected by the power supply, the change in voltage output and the power supply loading which is discussed later.
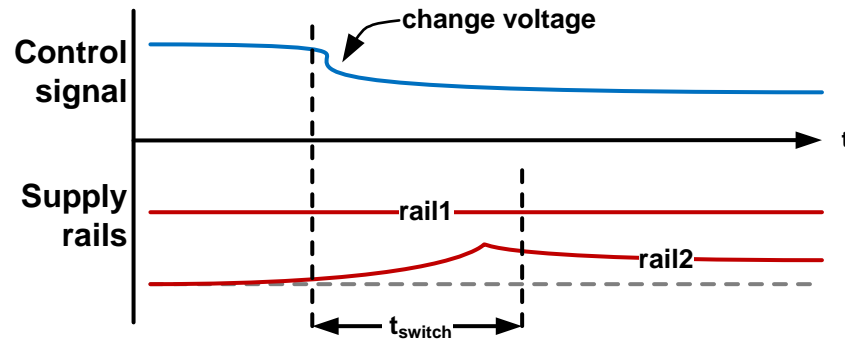


Figure 3.12: The switching time, $t_{switch}$, illustrated in a TIMR design.

The choice of power supply topology affects the switching time. The best way to adapt to this consideration is to simulate with the exact topology that will be used in the design. Also, the size of the output voltage change will have an effect on the settling time. Simulating with the largest possible voltage change (from $V_{min} \rightarrow V_{max}$ and $V_{max} \rightarrow V_{min}$) will find the worst case switching time from this effect. We will show later that simulations with an accurate power supply model with the largest voltage change can identify the settling time.

The implication on the design of the value of $t_{switch}$ is the speed (or frequency) at which the power supply rails can be changed. Because the circuit must wait for the rail to settle before changing the active rail, it cannot switch the supplies faster than $t_{switch}$ plus a buffer time. This means that the upper limit on time interleaving is set by $t_{switch}$.

## 3.3.2   Power Supply Loading

The loading of the power supply is another consideration of the design of a TIMR system. The circuit acts as a load on the power supply that will draw current from the active rail up to a certain limit, $I_{max}$. The power supply wires from the voltage regulator to the circuit will add to the effective loading of the supplies they have a parasitic capacitance to other components on the silicon die [38]. The power supply must be able to supply the circuit with enough power for operation (as with any design), but it also must not have an increased switching time due to the loading.

Simulation of the circuit with accurate loading models is important to ensure that the switching time will be sufficient. Proper guard bands added to the switching frequency of the TIMR control can also help reduce the effect of the loading. Solid routing practices on the chip can help alleviate the parasitic loading due to the power line routing.

It is important to note that while the switching time will be impacted by the loading of the power supply, a voltage rail that is being changed is not sourcing any dynamic current to the circuit. In the TIMR design when a voltage regulator changes its output it is the regulator that is not active. This means that the switching will occur on a rail that has ideally zero $I_{load}$. There is non-ideal leakage through the power gate PMOS that allows for some finite $I_{load}$, but this value is much less than the current through the circuit during normal operation. As such, the transients introduced on the power rail by switching the output voltage should have no effect on the operation of the circuit and should not be made worse due to the loading on the rail.

### 3.3.3 Control Overhead

The final consideration that we will examine is the control overhead introduced in the design by TIMR. This overhead comes from two places: the control hardware and the second power supply line that must be routed to the circuit.

The control hardware is a digital circuit that can be synthesized from register transfer level (RTL) code written in Verilog or VHDL. This circuitry will continually monitor the throughput sensors (thermal, performance counter, energy monitor, etc) and make decisions on switching the voltage regulator. If a change is needed, the circuitry must communicate to the power supply the new desired output level. Once the supply has settled, the control circuitry must toggle the power gate headers. The control hardware circuit can be made through a simple state machine, the implementation of which is beyond the scope of this work. The resulting circuitry takes up area on the die that could otherwise be used for other circuitry and consumes energy to operate which are both overhead to the TIMR design.

Because TIMR requires a second power supply rail for the circuit, there is increased area used in routing the design and increased energy in the leakage of the second rail. A traditional multi-rail design, as seen in [15, 25, 36], has static rails so it does not require that the operating circuit have its own unique power supply lines; multiple circuits can share the same power supply lines on the same silicon die. TIMR requires that each circuit have its own dedicated pair of power supply lines on the same silicon die, which is unique to the TIMR approach. There is also an energy overhead in changing the output voltage of the power supply which we will discuss later.

The area overhead must be an acceptable cost of the implementation of TIMR. The energy overhead must be less than the energy saved by the design or there is no point in implementing TIMR.

## Chapter 4 – Simulation and Results

We will present our simulation setup and results in this chapter. We will conclude with comparison to other techniques.

## 4.1 Off-Chip Regulator Simulations

As we have seen in Section 2.2.6 and Section 3.3, there are two main factors that make the implementation of TIMR possible. These two factors are the switching time, $t_{switch}$, and the switching energy, $E_{switch}$. Both are overhead to the implementation and must be less than the time and energy savings or the implementation of TIMR is impractical. To find values for $t_{switch}$ and $E_{switch}$, we propose and execute a number of SPICE simulations. These simulations, which use existing power supply models, find values that show that the timing and energy overhead of TIMR is less than the potential savings and demonstrates that the implementation is practical.

### 4.1.1 Simulation Setup

The first part of the simulation is the choice of voltage regulator. While there are a number of choices of voltage regulators that could be used, we chose a regulator from Linear Technologies, the LTC3605 [9] with specifications found in Table 4.1.

The LTC3605 regulator was chosen for a number of reasons which include input

Table 4.1: Specifications of LTC3605 voltage regulator.

| Type | Specification |
|---|---|
| Input voltage | $4 \leftrightarrow 15V$ |
| Output voltage | $600mV \leftrightarrow V_{in}$ |
| Output current | $5A$ |
| Topology | Synchronous buck regulator |
| Efficiency | $96\%$ |
| Output tracking | Yes |
| Reference | $0.6 \pm 1\%$ |
| Integrated FETs | $70m\Omega$ top and $35m\Omega$ bottom |

voltage range, output voltage range, internal MOSFETs, output power and topology. These design choices are detailed as follows:

**Input voltage range**  The input voltage range ($4 \leftrightarrow 15V$) makes this part acceptable for power sources of both $5V$ and $12V$, both common source voltages from batteries and larger power supplies [42].

**Output voltage range**  The output voltage range of this part needs to be low enough to operate at (or at least near) near-threshold levels to demonstrate the potential energy savings. The selected part can have an output voltage ranging from $600mV$ to the input voltage, $V_{in}$. This part will be able to approach the threshold voltage to allow for low energy operation.

**Internal MOSFETs**  Internal MOSFETs (power FETs that are internal to the voltage regulator package) are desirable because they take the uncertainty of the MOSFET selection out of the process. With internal FETs, we can describe results that are dependent on one part, rather than three (voltage regulator and two FETs).

**Output power** The output power of the regulator must be able to power the circuit. For our tests, we assume a large load that will require up to $5W$. This means the system must be able to deliver $5A$ at $1V$. The selected part can deliver up to $5A$ output current.

**Topology** The topology choice is important because it needs to represent a typical device. The selected part is a high efficiency (up to $96\%$) synchronous buck regulator which is representative of the application space we are targeting (embedded digital systems) [42].

With the regulator selected, we now setup the simulation circuit. We use the circuit that was shown previously in Section 3.2.1. The author of [12] provided the foundation of the circuit along with the design SPICE file found in [10]. This base schematic requires two changes to meet the needs of this simulation. First, the feedback circuit must be modified to allow dynamic digital control. Second, the load capacitance, $C_L$ must be updated to match the simulated circuit.

The foundation circuit uses a resistor divider to provide feedback to the regulator. This technique is simple, but does not offer any method for digital control. To add this functionality, we add a VID controller. The LTC1706-61 [11] from Linear Technologies was chosen because of its feedback range and reference voltage ($600mV$) which matches that of the LTC3605 regulator. This part provides the reference voltage to the FB pin on the regulator based on a sample of the output voltage and a set-point which is determined by five digital control lines. This part is recommended for use with the AMD Opteron processor line [11] so it is a practical addition to our circuit, not a theoretical

linchpin.

The other change to the foundation circuit is the loading capacitance, $C_L$. This capacitance, as discussed in Section 3.2.1, is used to affect the ripple and transient response of the regulator [21, 22]. With the additional load from the circuit and the design practice found in [21, 22], we selected a large value for $C_L$ of $500\mu F$. This large value of capacitance will effectively reduce load transients and model the loading capacitance of the circuit.

With the adaptations to the foundation circuit, we arrive at the simulation circuit seen in Figure 4.1 with the values in Table 4.2. This circuit will be used for our simulations.

Table 4.2: Values of components in simulation circuit.

| Component | Value |
|:---:|:---:|
| $C_1$ | $22\mu F$ |
| $C_2$ | $2.2\mu F$ |
| $C_3$ | $0.1\mu F$ |
| $C_4$ | $10pF$ |
| $C_5$ | $220pF$ |
| $C_6$ | $1nF$ |
| $C_L$ | $500\mu F$ |
| $L_1$ | $1\mu H$ |
| $R_1$ | $162k\Omega$ |
| $R_2$ | $16k\Omega$ |
| $D_1$ | $CMDSH2-3$ |
| $V_{IN}$ | $5V$ |
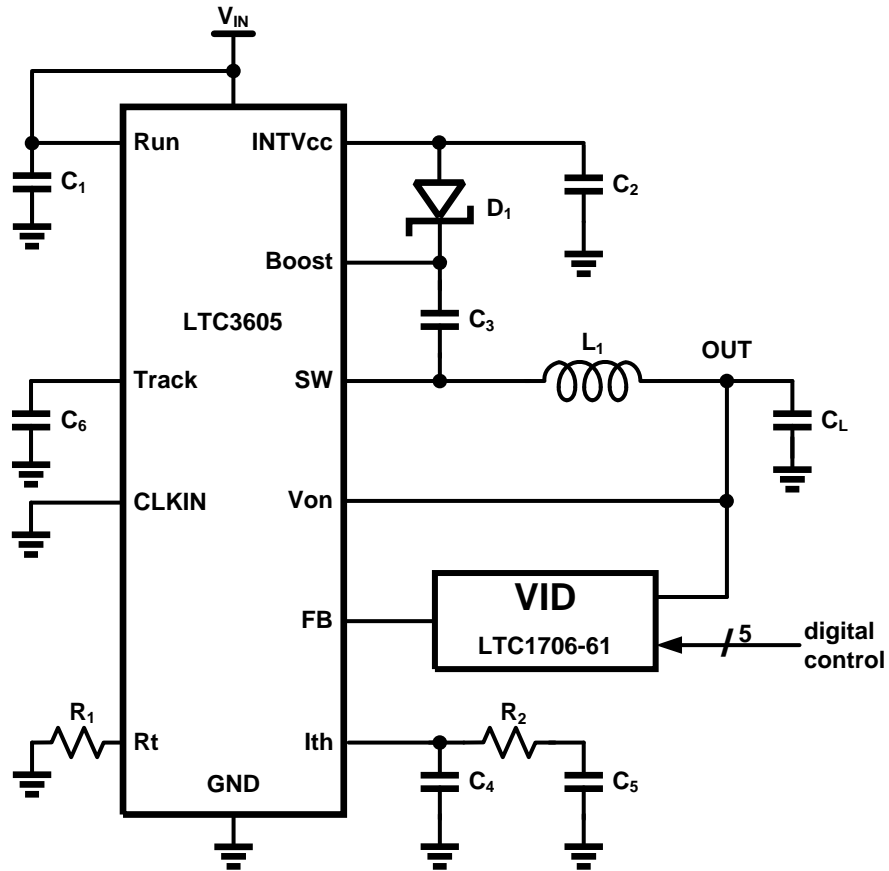
Figure 4.1: Circuit used for SPICE simulations.

## 4.1.2 Simulation for Time

We ran simulations to find the value to $t_{switch}$ with our simulation circuit. This was done by allowing the circuit to stabilize, then changing the value of the digital input on the VID controller to change the output voltage of the regulator. The time that it took to settle to the new output voltage is the switching time. This was done from $V_{min} \to V_{max}$ and $V_{max} \to V_{min}$ under $10\%$ and $100\%$ maximum load conditions. The resulting waveforms can be seen in Figure 4.2. The exact values of $t_{switch}$ can be found

in Table 4.3.

Table 4.3: Values of $t_{switch}$.

| Transition | 10% load | 100% load |
|---|---|---|
| $V_{min} \rightarrow V_{max}$ | $74.7\mu s$ | $64.8\mu s$ |
| $V_{max} \rightarrow V_{min}$ | $63.6\mu s$ | $61.6\mu s$ |

Note that in these simulations we assume that the load is connected during the voltage transition, thus the maximum possible value of $t_{switch}$ is observed. In a system described in the Implementation section, the voltage would transition and be exposed to a small transient followed by a large transient when the load was reconnected to the rail.
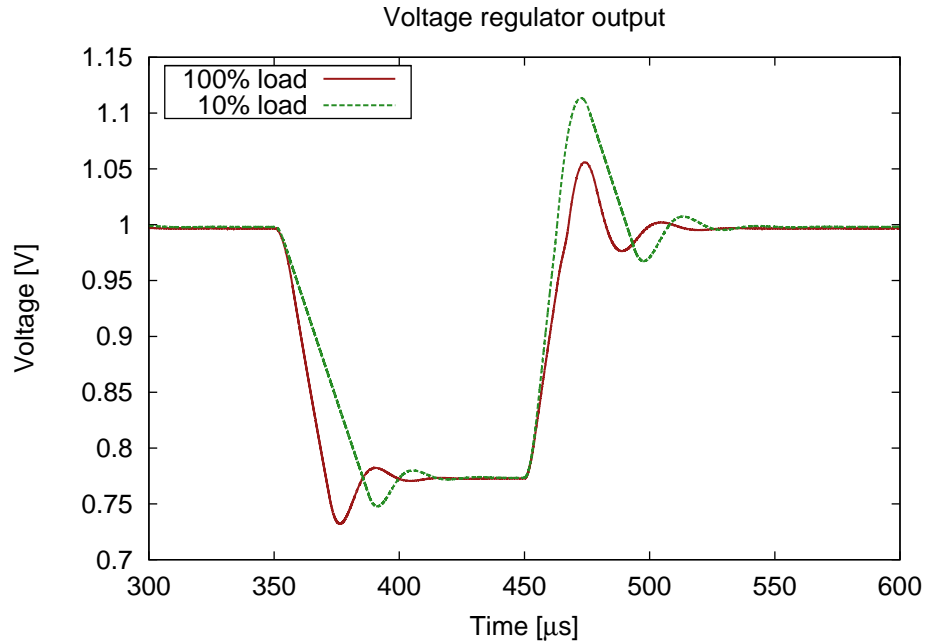


Figure 4.2: Simulation results for $t_{switch}$.

### 4.1.3   Simulation for Energy

We conclude the simulations with finding the switching energy, $E_{switch}$. This is the value of the energy consumed by switching the voltage regulator output. The value of this energy can be found with the following equation:

$$E_{switch} = P_{transition} \times t_{switch} - E_{active}$$

Where $P_{transition}$ is the power consumed by the voltage regulator during the transition period and $E_{active}$ is the energy consumed by the regulator during normal operation. Both of these values can be found from simulation.

As Figure 4.3 demonstrates, there is a large spike in power consumption by the regulator during both transitions of the output voltage. The simulation results for average power over this time and total energy consumed are shown in Table 4.4.

Table 4.4: Result from $e_{switch}$ simulations.

| Transition | Average Power | Energy |
|---|---|---|
| None | $212mW$ | $6.57\mu J$ |
| $V_{min} \rightarrow V_{max}$ | $5.67W$ | $125\mu J$ |
| $V_{max} \rightarrow V_{min}$ | $2.61W$ | $104\mu J$ |

The power consumption of the regulator during a high to low transition is less than that of a low to high transition because the regulator is able to use some of the stored charge while lowering output voltage. This can be seen in the plot of power consumption of the load capacitance $C_L$, seen in Figure 4.4. During lowering, the load capacitor, $C_L$, supplies charge to the circuit and thus has a negative power consumption, while it has a very large positive power consumption during the charging of a low to high transition.
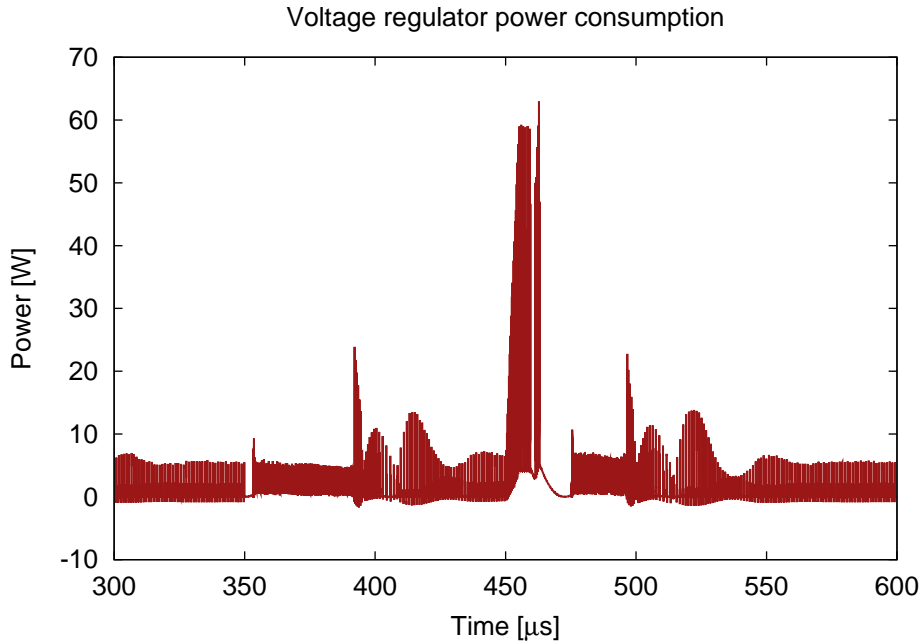
Figure 4.3: Simulation results for $E_{switch}$ showing the power consumption of the voltage regulator. Note the spike around $400\mu s$ as the output voltage is lowered and the large spike at $460\mu s$ as the output voltage is raised.

## 4.2   Results

In this section, we will discuss how the results of simulation effect the implementation of TIMR. Using the experimental values of $t_{switch}$ and $E_{switch}$ we can determine the practical limitations of the TIMR approach.

We will begin by looking at the time overhead. As discussed in Section 3.3.1, the switching time limits how often a transition can be made on one of the supply rails. We found that the worst case switching time is $74.4\mu s$ ($V_{min} \rightarrow V_{max}$ with 10% load). We
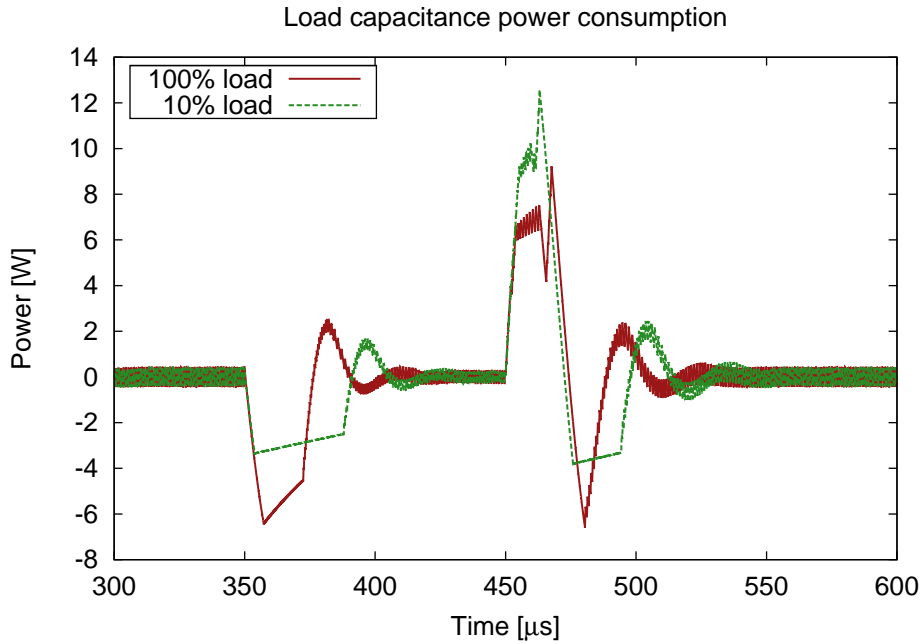
Figure 4.4: Simulation results for $E_{switch}$ showing the power consumption of the load capacitance.

cannot transition the supply voltage at a rate faster than the switching time. In practice, it makes sense to add an additional guard band to this number. Applying a $25\%$ padding to this time gives $100\mu s$ or $10kHz$ operation. This will be the limit of how often we can change the supply voltage on a single rail.

Now we will examine the energy overhead of TIMR. We found that the worst case energy was $118.43\mu J$ when switching from a low to high voltage. This energy overhead is much more significant than that of multi-rail designs [15], but the difference can still be made up by the energy savings. We find that the power savings is given by:

$$P_{savings} = (V_H - V_L) \times I_{load}$$

In our example case of switching from $1V$ to $750mV$ we find that the power savings is $1W$ or $1\mu J/\mu s$. This result gives the break even point for the transition, which is given in [15] as:

$$N_{BE} = \frac{E_{savings}}{E_{switch}} = \frac{P_{savings} \times t}{E_{switch}} = \frac{1\mu J/\mu s \times t}{118.43\mu J}$$

This leaves us with a time of $118.43\mu s \approx 120\mu s$ that we must stay at the new voltage level to break even. This is slightly longer than the time requirement, so it must be taken into consideration from the control logic.

The final results show that as long as the switching of the voltage rails is limited to greater than $120\mu s$ the time and energy requirements of TIMR are met. As long as the control logic of TIMR does not attempt to operate faster than this, we will see and energy savings and proper operation.

## 4.3   Comparison

With the overhead of TIMR discussed, we can now compare the results to other energy saving techniques.

### 4.3.1 Traditional DVFS

TIMR offers a significant advantage over traditional DVFS in both time and energy. As was described in Section 3.2.3.1, the timing of TIMR allows for nearly no time overhead as long as the control is not switched faster then $t_{switch}$. This behavior was seen in Figure 3.7. There is only a very small energy overhead. The voltage regulator changes its output in both techniques, so the only energy difference is the power gate. This is a very small overhead ($\approx 2pJ$ [15]). If the designer can afford to route the second power line and construct the power gates and control circuitry, TIMR can offer a significant time advantage to traditional DVFS.

### 4.3.2 Multi-Rail Supplies

While TIMR and multi-rail supplies are similar, there is a significant difference in their overheads. Multi-rail supplies offer static supply voltages with a very minimal time and energy overhead. TIMR offers dynamic supply voltages at a modest time and energy overheads; the area and control overheads are similar. The comparison between the two techniques depends on the implementation and application. Because multi-rail offers energy savings with a very low overhead, it may be the best option for conservative designs. When large energy savings are needed, TIMR offers dynamic supply rails that can offer greater energy savings if the design can afford more coarse grained rail switching. The choice between the two techniques depends heavily on the application as both offer energy savings.

### 4.3.3  Race to Sleep

While race to sleep attempts to eliminate leakage current, TIMR may offer a better energy savings and finer grained control over this technique. As illustrated in Figure 4.5, if the execution time of the sub-functions that make up the program (f1-5) is less than the switching time of TIMR then each function can operate at an individual supply voltage. This can offer energy savings on a finer grain when compared to traditional race to sleep. We can also see that there are a number of opportunities for "sub-sleep," sleeping after the end of function. This additional energy savings can be used in the case when there is a long memory access at the end of a function in processors where out-of-order execution is not possible (common to low power ARM architectures [3]).
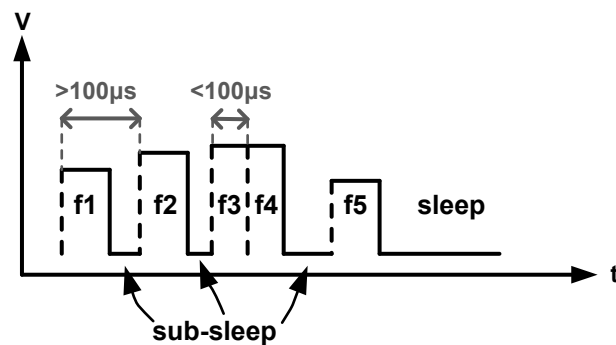


Figure 4.5: TIMR based race to sleep with individual voltage supplies and "sub-sleep."

# Chapter 5 – Conclusion

While a number of energy saving techniques currently exist, trade offs exist with each. We have shown a technique that provides large energy savings without timing restrictions or static supplies. Time Interleaved Multi Rail offers a dynamic supply rail as seen in traditional DVFS without the time penalty of switching the supply. TIMR also offers multiple supply rails as seen in typical designs with multi-rail supplies, but does so without a static supply voltage. We have demonstrated a number of possible control schemes for TIMR as well as the design considerations. We have shown that there are time and energy overheads of the design that must be taken into account with the control scheme. Finally, we compared TIMR with a few common energy saving techniques. Our approach to energy savings offers a number of benefits with manageable overhead.

# Bibliography

[1] K. Agarwal and K. Nowka. Dynamic power management by combination of dual static supply voltages. In *Quality Electronic Design, 2007. ISQED '07. 8th International Symposium on*, pages 85 –92, march 2007.

[2] M. Anis, S. Areibi, M. Mahmoud, and M. Elmasry. Dynamic and leakage power reduction in mtcmos circuits using an automated efficient gate clustering technique. In *Design Automation Conference, 2002. Proceedings. 39th*, pages 480 – 485, 2002.

[3] ARM. The cortex a8 microprocessor. Technical Report SPRY112A, Mar. 2008.

[4] M.A. Awan and S.M. Petters. Enhanced race-to-halt: A leakage-aware energy management approach for dynamic priority systems. In *Real-Time Systems (ECRTS), 2011 23rd Euromicro Conference on*, pages 92 –101, july 2011.

[5] B.H. Calhoun and A.P. Chandrakasan. Ultra-dynamic voltage scaling (udvs) using sub-threshold operation and local voltage dithering. *Solid-State Circuits, IEEE Journal of*, 41(1):238 – 245, jan. 2006.

[6] Juanjuan Chen, Xing Wei, Yunjian Jiang, and Qiang Zhou. Improve clock gating through power-optimal enable function selection. In *Design and Diagnostics of Electronic Circuits Systems, 2009. DDECS '09. 12th International Symposium on*, pages 30 –33, april 2009.

[7] Intel Corporation. Voltage regulator-down (vrd) 11.0. Technical report, Nov. 2006.

[8] Linear Technology Corporation. 5-bit vid voltage programmer for amd opteron cpus. Technical Report 170661f, 2002.

[9] Linear Technology Corporation. 15v, 5a synchronous step-down regulator. Technical Report LT 1109 REV B, 2009.

[10] Linear Technology Corporation. Ltc3605 demo circuit - 15v, 5a, 4mhz, synchronous step-down regulator (12v to 1.8v @ 5a). Technical report, 2009.

[11] Linear Technology Corporation. 5-bit vid voltage programmer for amd opteron cpus. Technical Report 170661f, 2012.

[12] Tom Gross Linear Technology Corporation. $15v_{IN}$, 4mhz monolithic synchronous buck regulator delivers 5a in 4mm 4mm qfn. Technical Report Design Note 467, 2009.

[13] Joseph Crop, Evgeni Krimer, Nariman Moezzi-Madani, Robert Pawlowski, Thomas Ruggeri, Patrick Chiang, and Mattan Erez. Error detection and recovery techniques for variation-aware cmos computing: A comprehensive review. *Journal of Low Power Electronics and Applications*, 1(3):334–356, 2011.

[14] J. De Boeck. Game-changing opportunities for wireless personal healthcare and lifestyle. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 15 –21, feb. 2011.

[15] Liang Di, M. Putic, J. Lach, and B.H. Calhoun. Power switch characterization for fine-grained dynamic voltage scaling. In *Computer Design, 2008. ICCD 2008. IEEE International Conference on*, pages 605 –611, oct. 2008.

[16] D. Ernst, Nam Sung Kim, S. Das, S. Pant, R. Rao, Toan Pham, C. Ziesler, D. Blaauw, T. Austin, K. Flautner, and T. Mudge. Razor: a low-power pipeline based on circuit-level timing speculation. In *Microarchitecture, 2003. MICRO-36. Proceedings. 36th Annual IEEE/ACM International Symposium on*, pages 7 – 18, dec. 2003.

[17] A.H. Farrahi, Chunhong Chen, A. Srivastava, G. Tellez, and M. Sarrafzadeh. Activity-driven clock design. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 20(6):705 –714, jun 2001.

[18] S. Hanson, B. Zhai, K. Bernstein, D. Blaauw, A. Bryant, L. Chang, K. K. Das, W. Haensch, E. J. Nowak, and D. M. Sylvester. Ultralow-voltage, minimum-energy cmos. *IBM Journal of Research and Development*, 50(4.5):469 –490, july 2006.

[19] Steve Heath. *Embedded Systems Design, Second Edition*. Newnes, 2002.

[20] Jingcao Hu, Youngsoo Shin, N. Dhanwada, and R. Marculescu. Architecting voltage islands in core-based system-on-a-chip designs. In *Low Power Electronics and Design, 2004. ISLPED '04. Proceedings of the 2004 International Symposium on*, pages 180 – 185, aug. 2004.

[21] Texas Instruments. Two-phase, synchronous buck controller with integrate mosfet drivers. Technical Report SLUS776B, Dec. 2007.

[22] Texas Instruments. 3.3-v/5-v input, d-cap+tm mode synchronous step-down integrated fets converter with 2-bit vid. Technical Report SLUSAX2, Feb. 2012.

[23] R. Jejurikar, C. Pereira, and R. Gupta. Leakage aware dynamic voltage scaling for real-time embedded systems. In *Design Automation Conference, 2004. Proceedings. 41st*, pages 275 –280, july 2004.

[24] Hailin Jiang, M. Marek-Sadowska, and S.R. Nassif. Benefits and costs of power-gating technique. In *Computer Design: VLSI in Computers and Processors, 2005. ICCD 2005. Proceedings. 2005 IEEE International Conference on*, pages 559 – 566, oct. 2005.

[25] S. Khanna, K. Craig, Y. Shakhsheer, S. Arrabi, J. Lach, and B.H. Calhoun. Stepped supply voltage switching for energy constrained systems. In *Quality Electronic Design (ISQED), 2011 12th International Symposium on*, pages 1 –6, march 2011.

[26] Wonyoung Kim, D.M. Brooks, and Gu-Yeon Wei. A fully-integrated 3-level dc/dc converter for nanosecond-scale dvs with fast shunt regulation. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 268 –270, feb. 2011.

[27] Chien-Wei Kuan and Hung-Chih Lin. Near-independently regulated 5-output single-inductor dc-dc buck converter delivering 1.2w/mm2 in 65nm cmos. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pages 274–275, feb. 2012.

[28] F. Kuttner, H. Habibovic, T. Hartig, M. Fulde, G. Babin, A. Santner, P. Bogner, C. Kropf, H. Riesslegger, and U. Hodel. A digitally controlled dc-dc converter for soc in 28nm cmos. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 384 –385, feb. 2011.

[29] Yongpan Liu, Huazhong Yang, R.P. Dick, H. Wang, and Li Shang. Thermal vs energy optimization for dvfs-enabled processors in embedded systems. In *Quality Electronic Design, 2007. ISQED '07. 8th International Symposium on*, pages 204 –209, march 2007.

[30] S.M. Martin, K. Flautner, T. Mudge, and D. Blaauw. Combined dynamic voltage scaling and adaptive body biasing for lower power microprocessors under dynamic

workloads. In *Computer Aided Design, 2002. ICCAD 2002. IEEE/ACM International Conference on*, pages 721 – 725, nov. 2002.

[31] J.D. Meindl and J.A. Davis. The fundamental limit on binary switching energy for terascale integration (tsi). *Solid-State Circuits, IEEE Journal of*, 35(10):1515 –1516, oct 2000.

[32] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, and J. Yamada. 1v high-speed digital circuit technology with 0.5 mu;m multi-threshold cmos. In *ASIC Conference and Exhibit, 1993. Proceedings., Sixth Annual IEEE International*, pages 186 –189, sep-1 oct 1993.

[33] S. Oesterle, P. Gerrish, and Peng Cong. New interfaces to the body through implantable-system integration. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 9 –14, feb. 2011.

[34] Robert Pawlowski, Evgeni Krimer, Joseph Crop, Jacob Postman, Nariman Moezzi-Madani, Mattan Erez, and Patrick Chiang. A 530mv 10-lane simd processor with variation resiliency in 45nm soi. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pages 492–493, feb. 2012.

[35] R. Puri, D. Kung, and L. Stok. Minimizing power with flexible voltage islands. In *Circuits and Systems, 2005. ISCAS 2005. IEEE International Symposium on*, pages 21 – 24 Vol. 1, may 2005.

[36] M. Putic, Liang Di, B.H. Calhoun, and J. Lach. Panoptic dvs: A fine-grained dynamic voltage scaling framework for energy scalable cmos design. In *Computer Design, 2009. ICCD 2009. IEEE International Conference on*, pages 491 –497, oct. 2009.

[37] J. Rabaey, H. DeMan, M. Horowitz, T. Sakurai, J. Sun, D. Dobberpuhl, K. Itoh, P. Magarshack, A. Abidi, and H. Eul. Beyond the horizon: The next 10x reduction in power - challenges and solutions. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, page 31, feb. 2011.

[38] Jan M. Rabaey, Anantha Chandrakasan, and Borivoje Nikolic. *Digital Integrated Circuits (2nd Edition)*. Prentice Hall, 2003.

[39] Sven Rosinger, Domenik Helms, and Wolfgang Nebel. Rtl power modeling and estimation of sleep transistor based power gating. In Nadine Azmard and Lars

Svensson, editors, *Integrated Circuit and System Design. Power and Timing Modeling, Optimization and Simulation*, volume 4644 of *Lecture Notes in Computer Science*, pages 278–287. Springer Berlin / Heidelberg, 2007.

[40] D. Sengupta and R.A. Saleh. Application-driven voltage-island partitioning for low-power system-on-chip design. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 28(3):316 –326, march 2009.

[41] Y. Shakhsheer, S. Khanna, K. Craig, S. Arrabi, J. Lach, and B.H. Calhoun. A 90nm data flow processor demonstrating fine grained dvs for energy efficient operation from 0.25v to 1.2v. In *Custom Integrated Circuits Conference (CICC), 2011 IEEE*, pages 1 –4, sept. 2011.

[42] Application Note 556 Introduction to Power Supplies. National semiconductor. Technical Report AN010061, Sep. 2002.

[43] Qi Wang and S. Roy. Power minimization by clock root gating. In *Design Automation Conference, 2003. Proceedings of the ASP-DAC 2003. Asia and South Pacific*, pages 249 – 254, jan. 2003.