

© 2012 Jane S. Tu

CONTINUOUS VOLTAGE-FREQUENCY SCALING (CVFS)

BY

JANE S. TU

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2012

Urbana, Illinois

Adviser:

Professor Naresh Shanbhag

ABSTRACT

Voltage reduction is an effective technique for minimizing energy consumption but suffers from delay penalty. Conventional methodologies require rigorous voltage regulation and workload scheduling to meet timing constraints. In this work, we observe that static CMOS is robust under low supply voltages, operates reliably during voltage transients, and exhibits similar voltage-delay characteristic across logic families. We present a continuous voltage-frequency scaling (CVFS) approach where supply variation is relaxed, and timing violations are avoided through the use of on-chip clock generation. A simple model of the critical path is presented to track circuit behavior in real time. This approach presents small overhead in data transition but enables energy optimization at the system-level. The contribution of this thesis includes the design of the digital blocks for a prototype chip in IBM 130nm technology.

To my parents, for their love and support

TABLE OF CONTENTS

| | |
|---|----|
| LIST OF FIGURES | v |
| LIST OF ABBREVIATIONS | vi |
| CHAPTER 1 INTRODUCTION | 1 |
| 1.1 Low Energy Digital Design | 1 |
| 1.2 Continuous Voltage and Frequency Scaling (CVFS) | 3 |
| 1.3 Thesis Organization | 5 |
| CHAPTER 2 LITERATURE SURVEY | 6 |
| 2.1 Minimum Energy Operation Point (MEOP) | 6 |
| 2.2 Dynamic Voltage Scaling (DVS) | 11 |
| 2.3 Near-Threshold and Sub-Threshold Operation | 20 |
| 2.4 Low-Voltage Design Considerations | 26 |
| CHAPTER 3 CVFS IMPLEMENTATION | 31 |
| 3.1 Basic Cells | 31 |
| 3.2 Critical Path Replica (CPR) | 32 |
| 3.3 Digital Controlled Oscillator (DCO) | 35 |
| CHAPTER 4 DIGITAL CIRCUIT BLOCKS | 38 |
| 4.1 Multiply-Accumulate (MAC) Finite Impulse Response (FIR) Filter | 38 |
| 4.2 Tunable Oscillator | 41 |
| 4.3 Input Controller | 42 |
| CHAPTER 5 SIMULATION RESULTS | 45 |
| 5.1 Design Verification | 45 |
| 5.2 Testing Procedure | 49 |
| CHAPTER 6 CONCLUSION | 51 |
| 6.1 Summary | 51 |
| REFERENCES | 52 |

LIST OF FIGURES

| | | |
|------|--|----|
| 1.1 | Compute VRM top-level block diagram. | 3 |
| 1.2 | Supply voltage in a Compute VRM. | 4 |
| 1.3 | Continuous voltage-frequency scaling (CVFS). | 5 |
| | | |
| 2.1 | Minimum energy operation point (MEOP). | 9 |
| 2.2 | Delay and energy vs. V_{dd} | 10 |
| 2.3 | Time-varying workload. | 12 |
| 2.4 | Normalized energy for fixed and variable supply voltage. . . . | 15 |
| 2.5 | Normalized energy of multiple system undithered and dithered. | 17 |
| 2.6 | Critical Path Monitoring (CPM) circuits | 19 |
| 2.7 | DVS topologies | 21 |
| 2.8 | I_d - V_{gs} curve | 22 |
| 2.9 | Inverted Temperature Dependence (ITD). | 25 |
| 2.10 | Normalized energy vs. sizing factor. | 28 |
| | | |
| 3.1 | Lumped RC models | 34 |
| 3.2 | Delay cell. | 36 |
| 3.3 | Lumped RC model of delay cell | 37 |
| | | |
| 4.1 | MAC FIR data flow graph. | 39 |
| 4.2 | MAC unit schematic. | 40 |
| 4.3 | Switching register. | 40 |
| 4.4 | DCO schematic. | 42 |
| 4.5 | X-register. | 43 |
| 4.6 | H-register. | 44 |
| | | |
| 5.1 | MAC FIR filter computation. | 46 |
| 5.2 | Data switching between voltage domains | 47 |
| 5.3 | DCO clock period vs. MAC critical path delay. | 48 |
| 5.4 | DCO tuning. | 48 |
| 5.5 | Load initial register values. | 49 |
| 5.6 | Rerun test vectors. | 50 |

LIST OF ABBREVIATIONS

| | |
|------|---|
| MEOP | Minimum Energy Operation Point |
| PVT | Process-Voltage-Temperature |
| DVS | Dynamic Voltage Scaling |
| GALS | Globally-Asynchronous Locally-Synchronous |
| MSV | Multiple Supply Voltages |
| LVD | Local Voltage Dithering |
| CVFS | Continuous Voltage-Frequency Scaling |
| DCO | Digitally Controlled Oscillator |
| MAC | Multiply-Accumulate |
| FIR | Finite Impulse Response |
| CPR | Critical Path Replica |
| RO | Ring Oscillator |
| DIBL | Drain-Induced Barrier Lowering |
| ABB | Adaptive Body Bias |
| AVS | Adaptive Voltage Scaling |
| AFS | Adaptive Frequency Scaling |
| RFF | Razor Flip-Flop |
| CPM | Critical Path Monitor |
| VCO | Voltage Controlled Oscillator |
| RDF | Random Dopant Fluctuation |

| | |
|------|---------------------------------|
| ITD | Inverted Temperature Dependence |
| ZTC | Zero Temperature Coefficient |
| LFSR | Linear Feedback Shift Register |

CHAPTER 1

INTRODUCTION

1.1 Low Energy Digital Design

The growing demand for portable, battery-powered electronics translates to the need for devices to be compact and low-power while meeting the functional performance. As device feature size scales, device density is increased and the speed of the devices is enhanced, making large memory storage and sophisticated functionalities possible. Throughput constraints are alleviated through algorithmic and architectural optimization. However, increasing complexity and variation have made power consumption the primary challenge in submicron design [1].

Thus, an attractive approach for energy efficient design is to operate circuits under low supply voltage, since the reduction of supply voltage provides quadratic power and energy saving for digital circuits. However, the reduction of supply voltage comes at the cost of decreased performance. The propagation delay is inversely proportional to the supply voltage and increases exponentially as supply voltage reduces near the threshold voltage. For sequential circuits, the propagation delay sets a lower bound to the clock period and, subsequently, places an upper bound on the maximum operating frequency.

Many device, circuit, and architecture-based scaling strategies have been studied for improvement in speed under voltage scaling [1, 2, 3, 4]. At the

device level, transistor sizing can be optimized for low power, where the low power cells are generally sized smaller than the standard cells to reduce the switching capacitance [3]. Also, threshold voltage can be varied by body bias to compensate for the delay degradation under lower supply voltage [4]. At the circuit-level, the use of complex gates and stacked devices should be minimized to ensure sufficiently large drive current to idle current ratio. At the architecture-level, parallelism can be applied to operate circuitry at a lower frequency while maintaining the same throughput; pipelined architecture can be employed to shorten the critical path [1].

Several system-level design methodologies have been examined. One approach is to design under *minimum energy operation point* (MEOP) [3, 5, 6, 4], where the supply voltage is reduced below the threshold voltage. Although this aggressive voltage scaling offers significant energy saving, it is also highly sensitive to *process-voltage-temperature* (PVT) variation and is constrained to low-performance applications. While promising for distributed sensor networks and medical devices, which operate below tens of megahertz, MEOP is not feasible for moderate and high throughput applications. Alternatively, under *dynamic voltage scaling* (DVS), discrete pairs of voltages and frequencies are dynamically adjusted to meet performance requirements [7, 8, 9, 10, 11]. This approach leverages the varying workload, achieving both high performance and energy minimization; however, its practicality is limited by the energy overhead and complexity of the control circuitry and voltage regulator. Current designs resort to a coarse grain tuning, which deviates from the theoretical maximum achievable energy saving under varying supply voltage.

Other related energy-reduction techniques include the utilization of *globally-asynchronous locally-synchronous* (GALS) clocks [12], *multiple supply volt-*

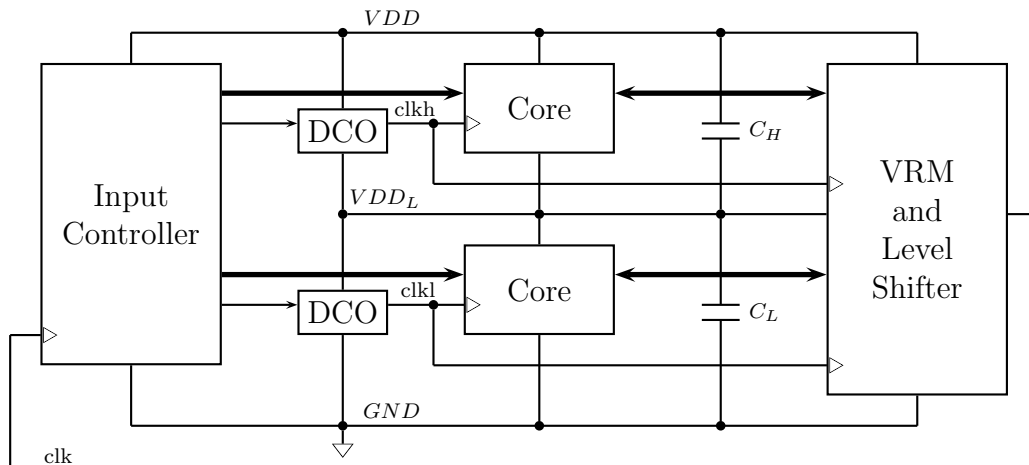


Figure 1.1: Compute VRM top-level block diagram.

ages (MSV) [6], and *local voltage dithering* (LVD) [13]. These methods divide the system-level optimization problem into a block-by-block optimization problem and show that further energy saving can be achieved with fine-grain tuning in multiple partitioned domains. However, these domain-based approaches are subject to latency and additional circuitry overhead.

1.2 Continuous Voltage and Frequency Scaling (CVFS)

Under *Compute Voltage Regulator Module* (Compute VRM) scheme [14], energy reduction is a joint-optimization of the voltage regulator and circuit at the system level. Fig. 1.1 is a top-level block diagram of the scheme. The core is replicated in two stacked voltage domains, where only one core is active at a time. Energy saving is achieved through charge recycling between the two domains and reducing losses associated with standard DC-DC regulators.

In each domain, the supply voltage can be modeled as the voltage across a discharging capacitor shown in Fig. 1.2. As the supply voltage decreases continuously, a *digitally controlled oscillator* (DCO), which tracks the crit-

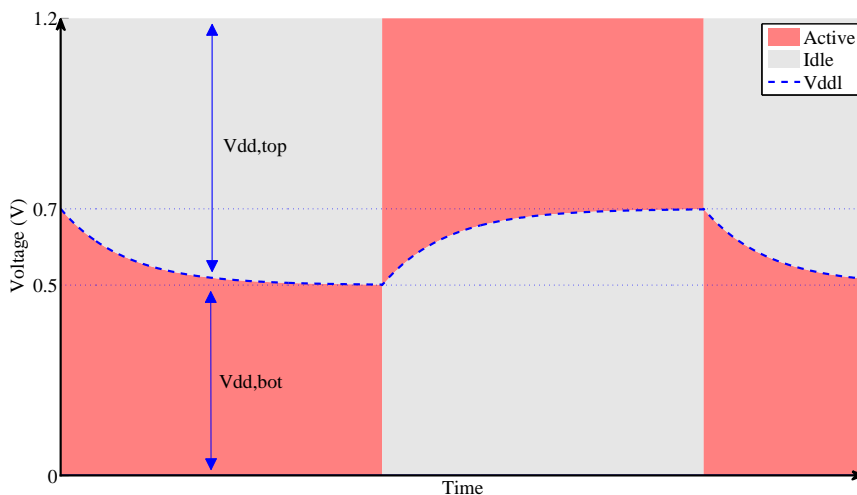


Figure 1.2: Supply voltage in a Compute VRM.

ical path delay of the core, produces a continuously decreasing maximum frequency clock that ensures correct circuit operation. This continuously matching of clock frequency to the instantaneous supply voltage will be referred to as *continuous voltage-frequency scaling (CVFS)* and depicted in Fig. 1.3. This approach tracks the instantaneous voltage-frequency pairs, in contrast to the discrete voltage-frequency pairs in conventional DVS, and achieves a continuous, minimum energy, error-free operation under varying supply voltage.

This thesis describes a prototype IC implementation of a CVFS core in IBM 130nm technology with supply voltage range of 500mV to 700mV, where the nominal supply voltage is 1.2V. The core is an 8-bit, 8-tap, folded *multiply-accumulator (MAC)* based *finite impulse response (FIR)* filter, a fundamental computational kernel employed in digital signal processing systems. Each voltage domain contains a DCO for local clock generation. The DCO incorporates the MAC's *critical path replica (CPR)* in a *ring oscillator (RO)* topology with tunable delay lines to track the circuit delay with

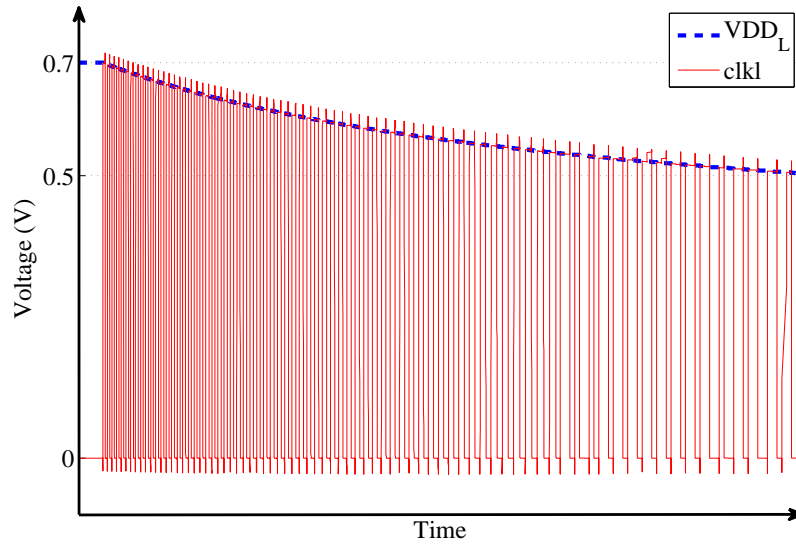


Figure 1.3: Continuous voltage-frequency scaling (CVFS).

adjustable margin for mismatch.

1.3 Thesis Organization

This thesis is organized as follows: Chapter 2 presents two conventional voltage scaling methodologies with analysis of the benefits and challenges in voltage reduction and variable supply. Next, device behavior and design considerations are examined for different operating voltages. Chapter 3 summarizes the basic building blocks of the circuit and describes the lumped RC models. Chapter 4 shows the implementation of each of the digital blocks. Lastly, Chapter 5 presents the simulation results and discusses the test features.

CHAPTER 2

LITERATURE SURVEY

2.1 Minimum Energy Operation Point (MEOP)

The total energy of a system comprises the dynamic and the leakage components. The short circuit component can be ignored for supply voltage smaller than the sum of the NFET threshold voltage and the absolute PFET threshold voltage. Chandrakasan and Brodersen [2] provide a comprehensive modeling of the energy consumption per clock cycle,

$$E_{total} = E_{dynamic} + E_{leakage} \quad (2.1)$$

$$E_{dynamic} = \alpha C V_{dd}^2 \quad (2.2)$$

$$E_{leakage} = V_{dd} I_{OFF} T \quad (2.3)$$

where α is the activity factor, C is the total switching capacitance, V_{dd} is the supply voltage, I_{OFF} is the leakage current, and T is the clock period.

The dynamic energy is dissipated during circuit switching events, where the switching capacitance is charged to the supply voltage for logic high and discharged for logic low. The probability of voltage transitions is modeled by the activity factor. Leakage energy occurs when the circuit is idle and arises from the gate and sub-threshold leakage current.

Consider the case in which the circuit is operating at the highest error-free frequency, where the smallest clock period is bounded by the circuit's critical

path delay t_p , which is modeled by

$$t_p = KC_{cp} \left(\frac{V_{dd}}{I_{ON}} \right) \quad (2.4)$$

where K is the delay fitting parameter that accounts for the slope of the input signal, C_{cp} is the capacitance of the critical path, and I_{ON} is the drive current.

Since the minimum energy operation generally occurs in the sub-threshold region [3, 5], I_{ON} and I_{OFF} can be expressed in terms of the sub-threshold current as a function of gate-to-source voltage V_{gs} and drain-to-source voltage V_{ds} . The sub-threshold current I_{SUB} is modeled by the following expression [4]:

$$\begin{aligned} I_{SUB} &= I_o e^{\frac{V_{gs} - V_{th} + \eta V_{ds}}{nV_T}} \left(1 - e^{-\frac{V_{ds}}{V_T}} \right) \\ &= I_o 10^{\frac{V_{gs} - V_{th} + \eta V_{ds}}{S}} \left(1 - e^{-\frac{V_{ds}}{V_T}} \right) \end{aligned} \quad (2.5)$$

$$I_o = \mu C_{ox} \left(\frac{W}{L} \right) (n-1) V_T^2 \quad (2.6)$$

$$S = nV_T \ln(10) \quad (2.7)$$

where I_o is the reference current, η is the *drain-induced barrier lowering* (DIBL) coefficient, V_T is the thermal voltage, $n \geq 1$ is the sub-threshold swing coefficient, S is the sub-threshold slope, μ is the mobility, C_{ox} is the oxide thickness, and (W/L) is the effective width to length ratio.

Consider an inverter chain of N stages. By substituting the delay and current expression in the energy equation, the total energy can be derived as

$$\begin{aligned} E_{total} &= NC_s V_{dd}^2 \left(\alpha + NK \frac{I_{OFF}}{I_{ON}} \right) \\ &= NC_s V_{dd}^2 \left(\alpha + NK 10^{-\frac{V_{dd}}{S}} \right) \end{aligned} \quad (2.8)$$

where C_s is the average switching capacitance of the node at each stage, which is the equivalent of (C/N) and (C_{cp}/N) for an inverter chain.

The total energy per clock cycle is plotted for a fixed threshold voltage in Fig. 2.1 for a 68-stage inverter chain in IBM 130nm process technology. The curve shows that voltage scaling is effective in energy reduction, and there exists a minimum energy operation point (MEOP) given by the 3-tuple (V_{dd}^*, f^*, E^*) , where V_{dd}^* , f^* , and E^* are the supply voltage, clock frequency, and minimum energy. For $V_{dd} > V_{dd}^*$, energy is dominated by the dynamic energy; for $V_{dd} < V_{dd}^*$, energy is dominated by the leakage energy. Since MEOP is the perfect balance between the dynamic and leakage energy, the minimum point varies with threshold voltage, activity factor, or logic depth.

The increase in propagation delay with respect to the reduction of supply voltage is plotted in Fig. 2.2. The initial reduction in supply voltage reduces the dynamic energy with modest increase in delay. As the supply approaches the threshold voltage, the propagation delay increases significantly, resulting in an increase in the leakage energy at low supply voltage. As a result, MEOP is suitable for applications where energy minimization is the primary concern and processor speed is not critical.

Kao et al. [4] incorporate body bias to explore the effect of varying threshold voltage. They show that scaling of the threshold voltage simultaneously with the supply voltage helps mitigate the delay penalty at low supply voltages. They demonstrate an implementation of an 8-bit MAC unit in 0.14 μm process, which operates at supply voltage down to 175mV at 166kHz, and show the energy benefit of simultaneously scaling the supply and threshold voltage. They proposed an adaptive body bias (ABB) scheme to compensate for the effect of sub-threshold variation on delay.

Wang and Chandrakasan [5] proposed a performance-constraint MEOP

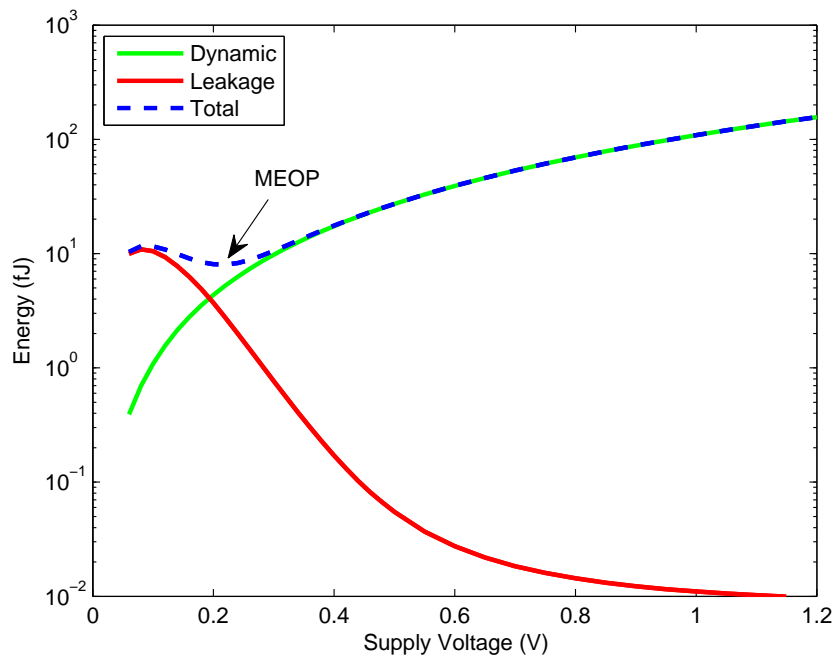


Figure 2.1: Minimum energy operation point (MEOP). This plot shows simulated total energy for a 68-stage inverter chain with minimum sized inverter (P:N size ratio of 2:1) in a 130nm IBM process. The dynamic and leakage energy curves are obtained via curve fitting using (2.9).




Figure 2.2: Delay and energy vs. V_{dd} . This plot shows the simulated delay values for a 68-stage inverter chain in a 130nm IBM process. The delay increases exponentially as the supply voltage approaches the threshold voltage (approximately 300 mV). The energy obtained in Fig. 2.1 is plotted to show the tradeoff between delay and energy.

methodology which binds the optimal V_{dd} - V_{th} combination by the maximum clock frequency requirement. This approach relaxes the speed penalty for applications with moderate throughput but compromises the maximum energy saving. They demonstrated an implementation of a 16-bit fast Fourier transform in 0.18 μm process, which correctly operated at minimum functional voltage of 180mV with clock frequency of 164Hz, while the energy optimal operating point was at 350mV at 10kHz. In their analysis, they assumed the supply regulator and bias circuitry do not exhibit energy overhead.

In addition to the throughput constraint, Kwong and Chandrakasan [15] and Fisher et al. [16] observed a functional yield-constraint MEOP, where they examined the effects of technology scaling and process variation for sub-threshold operation. They show that the increasing variability below 90nm technology node had severe impact on logic output swing and sub-threshold current. Subsequently, they observed a critical operating voltage at which minimum size devices can be used to satisfy the yield constraint. To mitigate the effect of variability, device upsizing is required for supply voltage below the critical operating voltage. The energy overhead due to upsizing shifts the energy minimum and raises the optimal operating voltage.

2.2 Dynamic Voltage Scaling (DVS)

The previous section showed a fundamental tradeoff between performance and energy saving under voltage scaling, where the maximum voltage reduction is limited by the worst case critical path of the circuit. Standard designs operate under a fixed supply voltage and a fixed clock frequency, thus energy is wasted when circuit is not operating at its peak performance. Dynamic voltage scaling (DVS) is a methodology that exploits the observation that

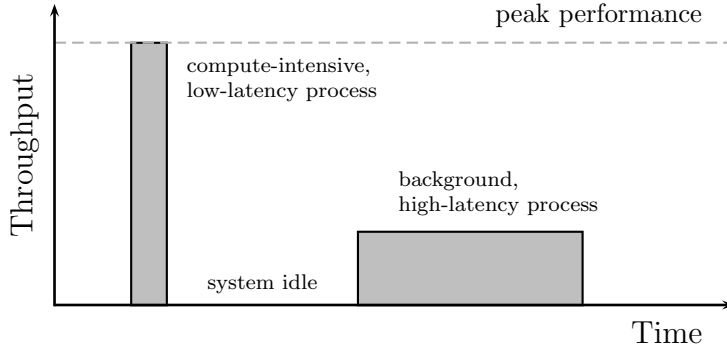


Figure 2.3: Time-varying workload.

the performance requirement is time-varying. An example of processor usage with varying workload is illustrated in Fig. 2.3 [9]. By varying the supply voltage and clock frequency during runtime, energy consumption can be minimized without sacrificing throughput.

Gutnik and Chandrakasan [7] presented an energy model for comparing the energy consumption of a system with fixed and variable supply voltage. The following analysis assumes the circuit operates above the threshold voltage, where the total energy is dominated by the dynamic component. The metric for the comparison is energy per operation, where each operation is required to complete within a sample period. For $V_{dd} > V_{th}$, the total energy E_{total} is given by

$$E_{total} \approx E_{dynamic} = MC_{eff}V_{dd}^2, V_{dd} > V_{th} \quad (2.9)$$

where M is the number of clock cycles per sample period, C_{eff} is the effective switching capacitance, equivalent to αC from (2.2),

Rearranging (2.4) results in an expression for the operating frequency with

respect to the supply voltage,

$$\begin{aligned} f &= \frac{1}{t_p} = \frac{I_{ON}}{KC_{cp}V_{dd}} \propto \frac{(V_{gs} - V_{th})^\beta}{V_{dd}} \\ &= A \frac{(V_{dd} - V_{th})^\beta}{V_{dd}} \end{aligned} \quad (2.10)$$

where β is a technology-dependent constant which accounts for velocity saturation, $\beta \in [1, 2]$. For simplicity, the device and circuit parameters are lumped into A .

With $\beta = 2$, substituting for V_{dd} from (2.9) into (2.10) we obtain an expression of energy in terms of the operating frequency, as follows:

$$E_{total}(f) = MC_{eff} \left(V_{th} + \frac{f}{2A} + \sqrt{f \frac{V_{th}}{A} + \left(\frac{f}{2A} \right)^2} \right)^2 \quad (2.11)$$

For a fixed sample period T_s , the operating frequency f only needs to be sufficiently large to complete M clock cycles, where $M = fT_s$. If a sample can be processed in fewer clock cycles, the operating frequency can be reduced, and the supply voltage can be scaled accordingly to reduce energy consumption. Subsequently, the expression for variable supply voltage can be obtained as

$$E_{var}(r) = fT_s C_{eff} V_o^2 \left(\frac{V_{th}}{V_o} + \frac{f}{2f_{ref}} + \sqrt{\frac{f}{f_{ref}} \left(\frac{V_{th}}{V_o} \right) + \left(\frac{f}{2f_{ref}} \right)^2} \right)^2 \quad (2.12)$$

$$= r f_{ref} T_s C_{eff} V_o^2 \left(\frac{V_{th}}{V_o} + \frac{r}{2} + \sqrt{r \frac{V_{th}}{V_o} + \left(\frac{r}{2} \right)^2} \right)^2$$

$$= E_o r \left(\frac{V_{th}}{V_o} + \frac{r}{2} + \sqrt{r \frac{V_{th}}{V_o} + \left(\frac{r}{2} \right)^2} \right)^2$$

$$r = \frac{f}{f_{ref}} \quad (2.13)$$

$$V_o = \frac{f_{ref}}{A} = \frac{(V_{ref} - V_{th})^2}{V_{ref}} \quad (2.14)$$

$$E_o = f_{ref} T_s C_{eff} V_o^2 \quad (2.15)$$

where V_{ref} is the reference supply voltage of the fixed supply system, f_{ref} is the reference operating frequency at V_{ref} , r is the normalized sample process rates, and V_o and E_o are scale factors to simplify the expression.

For a fixed supply voltage system, the operating frequency is fixed, where $f = f_{ref}$. If a sample is completed in fewer clock cycles, the circuit is idle for the remainder of the sample period, such as using clock gating or power gating. This means the energy per clock cycle is constant; only the number of clock cycles per sample varies. The expression for fixed supply is

$$E_{fixed}(r) = E_{var}(1)r \quad (2.16)$$

The normalized energy for variable and fixed supply voltage system is plotted in Fig. 2.4, as modeled by (2.12) and (2.16). The ratio between the two curves, $E_{var}(r)/E_{fixed}(r)$, is the theoretical energy saving ratio for a given sample. At high rate, the variable voltage is equal to the reference supply voltage; at low rate, the variable voltage approaches V_{th} , and the energy ratio approaches $(V_{th}/V_{ref})^2$. This analysis is for an ideal DVS system, where energy and time overhead is not considered. Nonetheless, this result provides an analytical upper bound on the achievable energy saving under variable supply voltage.

Energy and time overheads arise from the non-idealities that voltage transitions are neither lossless nor instantaneous. Burd and Brodersen [8] introduced two additional performance metrics associated with DVS: transition time and transition energy. For voltage transition between two voltage levels,

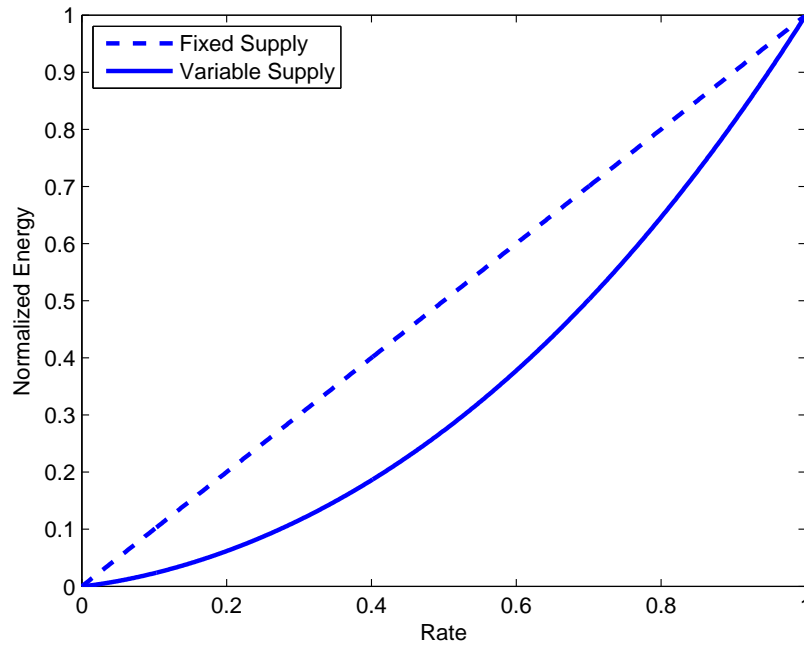


Figure 2.4: Normalized energy for fixed and variable supply voltage. This plot is obtained from the expressions derived in (2.12) and (2.16). It shows the upper bound of achievable energy saving under variable supply voltage from that of fixed supply voltage with respect to normalized sample rate r .

V_{dd1} and V_{dd2} , the two metrics can be modeled by the following:

$$t_{TRAN} = 2 \left(\frac{C_{vdd}}{I_{max}} \right) |V_{dd2} - V_{dd1}| \quad (2.17)$$

$$E_{TRAN} = (1 - \eta)C_{vdd} |V_{dd2}^2 - V_{dd1}^2| \quad (2.18)$$

where C_{vdd} is the supply capacitor, I_{max} is the maximum output current of the voltage regulator, η is the conversion efficiency of the DC-DC converter.

These metrics exhibit tradeoffs with the performance metrics of a standard voltage regulator, namely the supply ripple and conversion efficiency. For a standard regulator, it is advantageous to increase C_{vdd} to reduce supply ripple and enhance low-voltage conversion efficiency [8, 9]. On the contrary, as shown in (2.17) and (2.18), a large C_{vdd} increases both the transition time and the transition energy, which negatively impacts the potential benefit of DVS. This shows that voltage regulator for variable supply design is fundamentally different from that for standard fixed supply design.

From this observation, Yuan and Qu [11] classified DVS systems into three categories: ideal system, feasible system, and multiple system. An ideal system, as previously analyzed, assumes that voltage transitions can be an arbitrary function of time, with infinite allowable voltage levels such that the energy reduction is maximized. A feasible system is one in which voltage transition is continuous and is sub-categorized into optimistic and pessimistic models. The optimistic model allows the circuit to continue operating during voltage transient; the pessimistic model stops the operation until the new voltage level reaches steady state. A multiple system, also called voltage quantization, is one which has only a discrete number of voltage levels available simultaneously.

A general multiple system selects the lowest available voltage level which

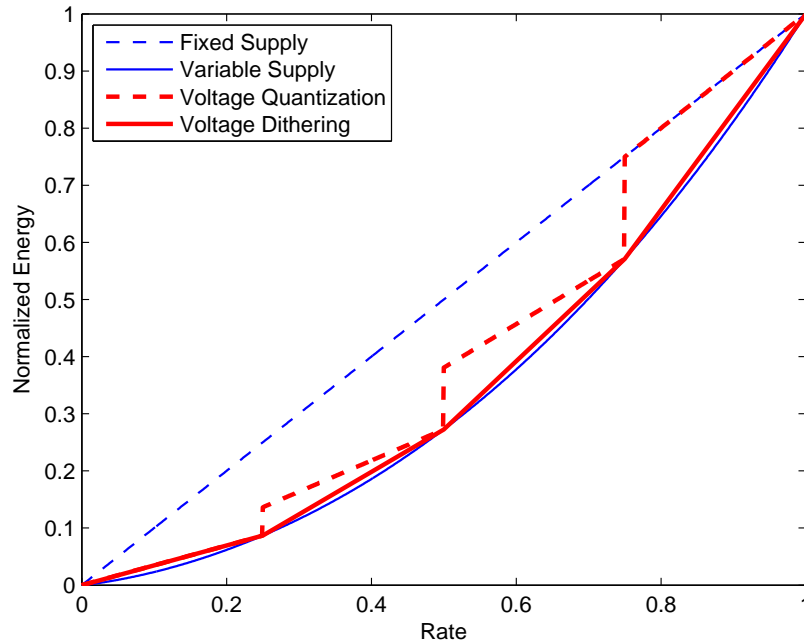


Figure 2.5: Normalized energy of multiple system undithered and dithered. This plot shows an example of multiple system with four available voltage levels. Voltage quantization is plotted using (2.16), where V_{ref} is allowed to switch to one of the lower voltage levels when frequency requirement is reduced. Voltage dithering is plotted using the average achievable rates among the available voltage levels. Figure 2.4 is plotted again to show that voltage dithering can achieve energy saving similar to that of an ideal DVS system.


satisfies the computation and latency requirements, such as workload (the number of clock cycles of a task), data rate (operating frequency), and completion deadlines (latency). A slight modification is called voltage dithering, which allows more than one voltage-frequency transition within one sample period, such that the average rate is matched to the desirable rate. Calhoun and Chandrakasan [13] showed that voltage dithering provides a near optimal energy saving, close to that of the ideal DVS system. Figure 2.5 shows the theoretical energy profile for multiple system of four available voltage levels with and without voltage dithering.

The optimal selection of frequency and voltage is implemented as part of


the real-time operating system, referred to as the voltage scheduler [13]. The basic scheduler algorithm predicts the workload requirement of a given task, calculates the optimal clock frequency or voltage, and updates the workload at the end of task. The optimization problem of the scheduler is not the focus of this discussion, but rather the underlying software which determines the order of operations and provides the corresponding workload requirements to the hardware.

A conventional DVS scheme uses pre-stored voltage-frequency pairs which are hardcoded in a lookup table. A more effective approach is to explore the timing margin by monitoring the circuit behavior using hardware. Two of the methods available are *adaptive voltage scaling* (AVS) and *adaptive frequency scaling* (AFS). The former requires a generation or regulation of the minimum supply voltage which satisfies a desired operating frequency. The latter requires the generation and adaptation of a maximum frequency clock to a variable supply voltage.

The effectiveness of these approaches relies on the accuracy of the critical path monitoring circuit. Two popular monitoring configurations are the Razor flip-flop (RFF) and TEAtime [17]; the block diagrams are shown in Fig. 2.6. RFF is a direct form of the *critical path monitor* (CPM). It uses double flip-flops to double sample the data in the critical path. If the normal sample and the delayed sample disagree, either voltage or frequency is adjusted and the data is recomputed. This method requires the modification of existing architecture and an additional pipeline stage, thus resulting in area overhead and performance penalty. TEAtime is an indirect form of CPM. It requires a non-inverting *critical path replica* (CPR) which emulates the longest delay path of the core. The CPR can be synthesized [17] or programmable [18]. It toggles one bit input and checks the output of the CPR



(a)



(b)

Figure 2.6: Critical Path Monitoring (CPM) circuits: (a) direct method: Razor Flip-Flop, and (b) indirect method: TEAtime.

per clock cycle. This approach requires additional safety margin to account for local variation.

Several generalized DVS topologies are shown in Fig. 2.7. The first one is a feasible system proposed by Kuroda et al. [10]. In their scheme, a clock of desired frequency is available externally. A CPM senses the supply voltage and indicates the slack time of the circuit. This information is fed to the feedback controller to adjust the supply voltage using a buck converter. The second one is proposed by Burd et al. [9], where the coarse voltage-frequency behavior of the core is pre-characterized and programmed to the *voltage*

controlled oscillator (VCO) that generates an internal clock. The discrepancy between the generated clock frequency and the desired frequency is fed to the feedback controller to vary the supply voltage. The last one is proposed by Calhoun and Chandrakasan [13], which employed voltage dithering. The controller switches between two available voltage levels based on the latency requirement. A *ring oscillator* (RO) which contains the CPR of the circuit generates an internal clock that matches the selected voltage level. The clock is gated during voltage transient.

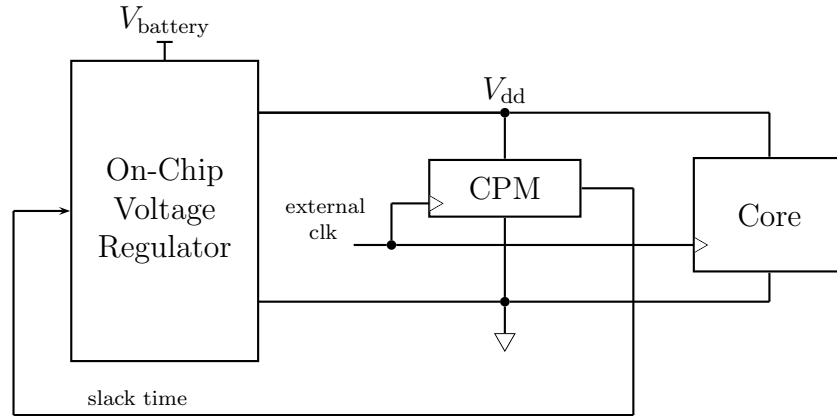
2.3 Near-Threshold and Sub-Threshold Operation

An important consideration when designing under voltage scaling is that a device behaves differently under different supply voltages. This aspect is especially a challenge for variable supply design since the circuit needs to operate under a wide range of supply voltages. Figure 2.8 shows the I_d - V_{gs} curve of a MOSFET in logarithmic and linear scale. The three regions denote the super-threshold, near-threshold, and sub-threshold operation.

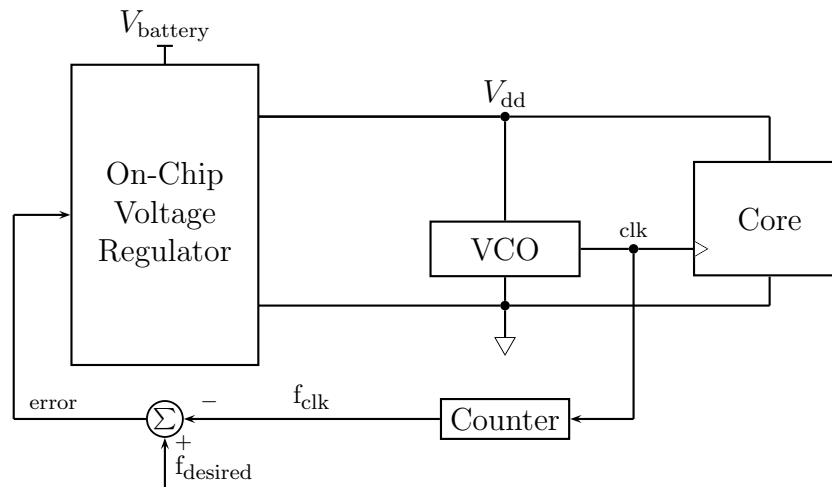
Commercial applications mostly operate in the super-threshold region [19], where the gate-to-source potential is above the threshold voltage, and the channel depletion region is in strong inversion. The first-order device model in saturation indicates that the drive current has a square dependence on the gate voltage as

$$I_d = \frac{1}{2} \mu C_{ox} \left(\frac{W}{L} \right) (V_{gs} - V_{th})^2 (1 + \lambda V_{ds}) \quad (2.19)$$

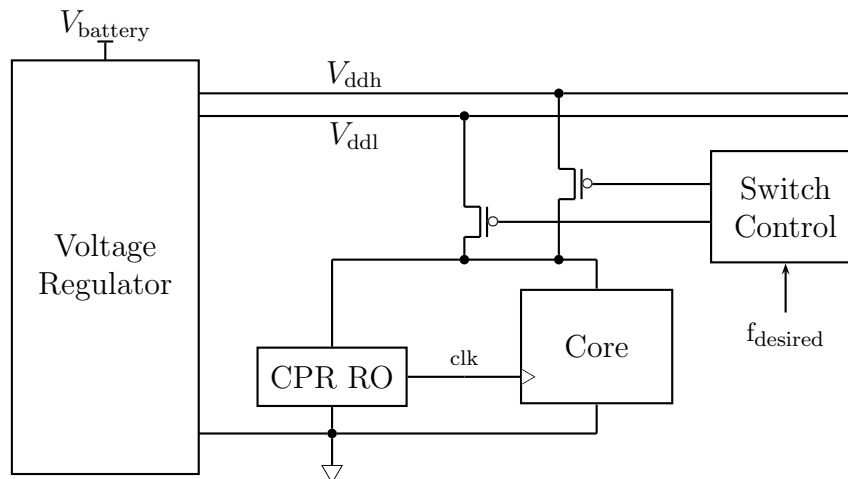
where λ is the channel-length modulation parameter. This is the case for long-channel devices; however, for short-channel device in submicron pro-



(a)

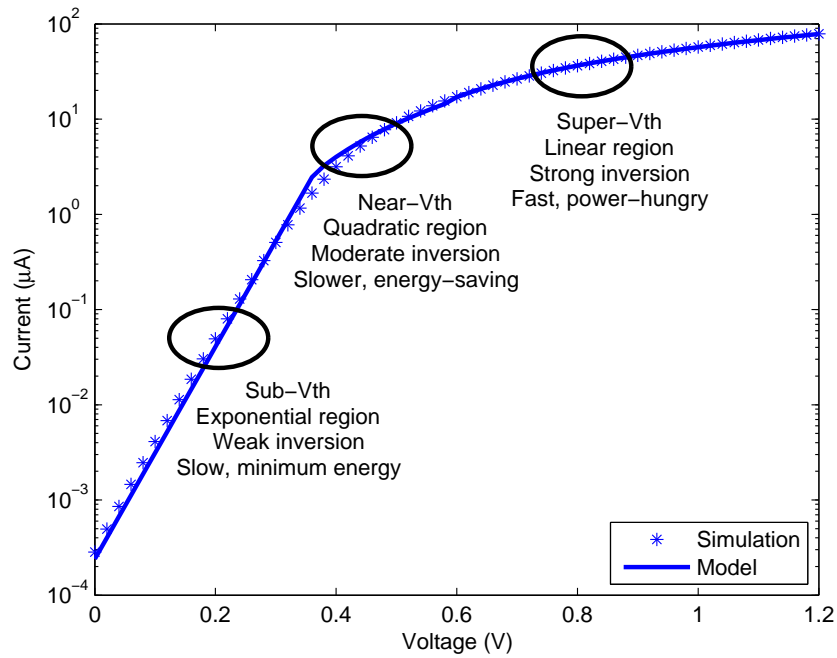


(b)




(c)

Figure 2.7: DVS topologies: (a) feasible system with AVS [10], (b) feasible system with AFS [9], and (c) multiple system with voltage dithering [13].



(a)



(b)

Figure 2.8: I_d - V_{gs} curve: The plot shows simulation data points for a NFET in a IBM 130nm process where $V_{gs} = V_{ds}$. The solid line is a fitting curve using (2.19) and (2.21) (a) logarithmic plot and (b) linear plot.

cesses, velocity saturation under high electric field causes this dependence to approach linear as

$$I_d \approx v_{sat} C_{ox} W (V_{gs} - V_{th}) \quad (2.20)$$

where v_{sat} is the saturation velocity. A practical modeling of this effect is to modify the square dependence in (2.19) to β -power law dependence, where $\beta \in [1, 2]$. β is the velocity saturation index and models the behavior amid linear and square dependence. Circuits that operate in the super-threshold region are fast but also power-hungry.

Near-threshold operation refers to the region in which the gate voltage is in the proximity of the threshold voltage. The reduction of supply voltage reduces the lateral electric field of the channel, thus reducing the carrier velocity below v_{sat} , and the drive current exhibits square dependence on the gate voltage. This region is slower but provides moderate energy saving.

Further reduction of the gate voltage below the threshold voltage enters the sub-threshold region, where device operates in weak inversion, and the drive current is dominated by the subthreshold current. As analyzed in Section 2.1, the drive current has an exponential dependency on the gate voltage. (2.5) and (2.6) are combined and repeated here:

$$I_{SUB} = \mu C_{ox} \left(\frac{W}{L} \right) (n - 1) V_T^2 e^{\frac{V_{gs} - V_{th} + \eta V_{ds}}{n V_T}} \left(1 - e^{-\frac{V_{ds}}{V_T}} \right) \quad (2.21)$$

Typical values of the sub-threshold slope $S = n V_T \ln(10)$ are around 60 to 90mV/decade current [2, 1]. This region is ideal for minimum energy designs but suffers large delay penalty.

In summary, as voltage scales, current dependence on the bias changes from linear to quadratic, and is exponential below the threshold voltage. Consequently, the circuit becomes extremely sensitive to voltage noise and

threshold voltage variation, which indirectly corresponds to process and temperature variation due to their impact on the threshold voltage. Note that while the current dependence on bias and V_{th} approaches exponential, dependence on other device and geometric parameters, namely the terms in $\mu C_{ox}(W/L)$, remains linear; as a result, supply and threshold variation becomes the dominant source of variation under low supply voltage.

The supply voltage fluctuation is highly correlated with system activities; however, the changes in the threshold voltage are somewhat unpredictable. The threshold variation is largely affected by drain-induced barrier lowering (DIBL) and random dopant fluctuation (RDF) [20]. DIBL is a short channel effect, that occurs when the drain depletion region near the channel surface lowers the source potential barrier. This effectively decreases the threshold voltage with increasing drain voltage. DIBL worsens with shorter channel length but improves with lower drain-to-source voltage. RDF is a process variation result from impurities in the channel, which varies the threshold voltage with no spatial correlation. RDF worsens with technology scaling and is the dominant source of variation in sub-threshold [21].

Temperature variations also have different impact at low voltages. The effect of temperature variation on the mobility and the threshold voltage can be modeled by the following [13]:

$$\mu(T) = \mu(T_o) \left(\frac{T}{T_o} \right)^{-M} \quad (2.22)$$

$$V_{th}(T) = V_{th}(T_o) - K(T - T_o) \quad (2.23)$$

where T_o is the reference temperature, M and K are positive fitting parameters. These expressions show that both mobility and the threshold voltage decrease with increasing temperature. In strong inversion, as temperature




Figure 2.9: Inverted Temperature Dependence (ITD). The plot shows the delay values from simulation of a 68-stage inverter chain in a IBM 130nm process for standard temperature corners. It shows that temperature dependence of performance inverts below the zero-temperature coefficient (ZTC).

increases, the degradation in mobility dominates and the circuit slows down. However, in weak inversion, the reduction in threshold voltage overtakes the reduction in mobility, and the circuit speeds up with increase in temperature. Thus, the effect of temperature on delay is dependent on the operating voltage. This phenomenon is referred to as inverted temperature dependence (ITD) [22]. Figure 2.9 illustrates an example of ITD; the crossover point is referred to as the zero-temperature coefficient (ZTC).

In addition to sensitivity to variations, as voltage scales, the drive current decreases tremendously, causing the circuit to be slower and more susceptible to error. At low voltages, the leakage current becomes comparable to the drive current, which is generally described by the I_{ON}/I_{OFF} ratio. To

illustrate the orders of magnitude reduction: in strong inversion, the ratio is approximately 10^7 ; whereas near MEOP, this ratio has been observed to be around 10^3 - 10^4 [23]. This ratio is further degraded based on circuit style and operating condition.

As discussed, circuit behavior and sensitivity to PVT variation vary significantly at different operating voltages. Thus, design under low supply voltage requires adequate margin to tolerate greater mismatch and careful choice of circuit topologies to guarantee sufficient drive current. For variable supply design, optimization will need to account for circuit robustness over the range of operating voltages.

2.4 Low-Voltage Design Considerations

While a circuit operating at low supply voltage provides significant energy saving, it has tradeoffs in speed and robustness. The key challenge arises from the reduction in I_{ON}/I_{OFF} ratio and the increasing sensitivity to process and temperature variation. Hence, techniques to minimize or compensate for these drawbacks are studied, while keeping energy reduction as the primary objective. Several low-power techniques at the device, circuit, and architecture-level are outlined in [6, 2, 1].

At the device-level, triple-well structure can be employed to scale V_{th} using body bias. Reverse body bias raises V_{th} and reduces leakage, while forward body bias lowers V_{th} and enhances performance. Thus, variable threshold voltage can be used to assist power gating or mitigate delay penalty. Wang et al. [24] performed energy and delay simulations to characterize the tradeoff between low-power and high-performance for a given circuit. Subsequently, an optimal throughput-constrained V_{dd} - V_{th} operating point can be obtained

from the resulting energy contours. Alternatively, Calhoun et al. [25] derived an analytical expression for the optimal V_{th} for performance-constraint MEOP. Threshold voltage tuning can also be used to compensate for current mismatch. Roy et al. [26] proposed a process-voltage monitoring logic which actively applies body bias to balance the strength of PFET and NFET.

Another design variable is device size. The traditional method of transistor sizing is to optimize for speed and balanced noise margins. From this standpoint, the optimal sizing of a cell is the equivalent of a symmetric inverter, where the ratio of PFET and NFET drive current ($I_{ON,P}/I_{ON,N}$) is close to unity and the switching threshold is roughly $V_{dd}/2$. However, near minimum functional voltage, the optimal ratio of PFET width to NFET width (W_P/W_N) is observed to be around 11-13 [25]. This ratio is area expensive and results in significant energy overhead with little performance gain. Keane et al. [27] calculated a closed form solution for sizing in sub-threshold, and their results showed that adequate performance can be achieved for sizing ratio of 2.5-3.

Other than relative sizing for performance, Chandrakasan and Brodersen [2] derived a first-order model to evaluate the need of uniformly upsizing (W/L) with respect to energy reduction. Figure 2.10 shows the normalized energy with respect to the sizing factor, where r is the ratio of parasitic capacitance to the gate capacitance of minimum sized devices. The plot shows that if the total load capacitance is not dominated by the interconnect, minimum sized devices are the theoretical optimal for minimum energy consumption. Kwong and Chandrakasan [15] considered the impact of variation and proposed a yield-constraint sizing scheme. From Monte Carlo simulations of a 32-bit adder synthesized in 90nm, it was shown that minimum sized devices satisfy yield constraint for supply above 340mV. Bol et al. [21] investigated




Figure 2.10: Normalized energy vs. sizing factor. This plot is reproduced from the first-order model derived in [2], where r is the ratio of parasitic capacitance to the gate capacitance of minimum sized devices. It shows that minimum sized devices are theoretical optimal for when the total load capacitance is not dominated by the interconnect.

the effect of RDF in sub-threshold regime, where the standard deviation of V_{th} variation is proportional to $(WL)^{-1/2}$. From this observation, enlarging the width or length helps mitigate RDF variation but at the expense of energy. A study was done for sub-threshold operation under technology scaling from 0.25 μm to 32nm. The results show that due to the severe impact of variations in subthreshold, upsizing of channel length would provide energy gain for advance technology nodes below 65nm.

At the circuit-level, under low voltages, long stacked devices and complex logic are undesirable. Long stacks of transistors reduce the effective drive current and suffer from body effect. Complex logic often contains stacked

devices and are also likely to suffer from parallel leakages, where idle current of parallel devices worsens I_{OFF} . Circuit styles using transmission gate or pass gate also contribute to long stacks and degraded output drive; however buffering can be added to eliminate long leakage path and guarantee logic regenerative property.

Many studies have shown that static CMOS is robust at low supply voltages [3, 5, 13] and demonstrated that various logic circuits exhibit the same supply dependence on delay [1]. For multiple or variable supply voltages operation, Burd and Brodersen [8] have shown that CMOS logic can tolerate large voltage transitions and that various logic circuits are functional during voltage transient.

At the architecture-level, parallel and pipelined architectures are common practice to retain throughput at reduced supply voltage. Chandrakasan and Brodersen [2] performed a comparison of the two in terms of power saving. The effectiveness of energy reduction can be analyzed from the following expressions for N -parallel structures or N -stage pipelining, considering only the dynamic energy component:

$$E_{ref} = C_{ref}V_{ref}^2 \quad (2.24)$$

$$E_{par} = C_{par}V_{par}^2 = (NC_{ref} + \Delta C) \left(\frac{V_{ref} + \Delta V}{N} \right)^2 \quad (2.25)$$

$$E_{pip} = C_{pip}V_{pip}^2 = (C_{ref} + \Delta C) \left(\frac{V_{ref} + \Delta V}{N} \right)^2 \quad (2.26)$$

where the subscript *ref* refers to the variables associated with the original architecture, ΔC and ΔV are discrepancies due to overhead of the additional circuitry. For parallel architecture, duplicated logic and routing can have significant area overhead; for pipelined architecture, registers can contribute substantial energy overhead. For both cases, supply voltage has limited scal-

ability depending on circuit style.

Furthermore, from Section 2.3, it was discussed that circuits become very sensitive to variations at very low voltages. These variations can be categorized in two forms. The variations due to device dimensions and system activities are systematic variations, which have global behavior and are correlated. These can be relaxed using adaptive techniques such as ABB or DVS schemes described in previous sections. Other forms of variation, particularly RDF, are random and uncorrelated. As shown earlier, one method is to increase sizing of the devices. Alternatively, the randomness can be treated as independent normally distributed variables. Thus, variability on delay can be reduced through averaging using stacked transistors, long logic depth, or multiple delay paths.

Despite many design considerations and variability, it was shown that for static CMOS digital circuits, design verification is only needed at the maximum operating voltage and timing verification is needed at the extrema voltages [8]. Hanson et al. [6] show that although the absolute noise margin degrades with reduced voltage, the fractional noise margin stays fairly constant for supply above 100mV. This is because at reduced voltages, the increase in gain improves device voltage transfer characteristic.

Lastly, despite the optimistic outlook for low voltage operation, many studies have not incorporated the overhead due to supporting circuitry in terms of overall system efficiency. These elements include level-shifter, voltage regulator, and interprocessor communication. Also, there are area and complexity overheads due to the usage of triple-well devices, extra supply rails, and additional signal routing.

CHAPTER 3

CVFS IMPLEMENTATION

Traditional energy reduction techniques are focused at the circuit level, where the reduction in supply voltage seems trivial. However, the reduction in supply voltage poses challenges in voltage regulation and interprocessor communication, which are commonly neglected and results in overly optimistic measurements. Therefore, VRM and the associated circuitry must be accounted for. Compute VRM is a system-level energy optimization methodology that addresses these tradeoffs and overheads, which involves the utilization of a CVFS core. This chapter describes the building blocks of the CVFS core.

3.1 Basic Cells

For the implementation of the digital circuitry, static CMOS logic and transmission gate logic cells are used. The specification of the design is IBM 130nm technology for supply voltage range from 500mV to 700mV. The cells are minimum size for minimal switching capacitance. The basic cell layout uses standardized cell height and is adaptable for either regular or triple-well design.

The only static CMOS cells used are the inverter and 2-input NAND gate. The sizing is equivalent to that of a minimally sized inverter with PFET width to NFET width ratio of 2:1. This ratio provides sufficient I_{ON}/I_{OFF} ratio for near-threshold operation because complex gates with long transistor

stacks are avoided. To implement complex gates efficiently and robustly under low voltage, transmission gate cells are employed. The cells used are 2-input XOR and 2-input multiplexer (MUX), where PFET and NFET are both minimum width. Each stage of a transmission gate is buffered by an inverter to avoid long transistor stacks and to provide output drive strength.

The derived cells based on the combination of static CMOS and transmission gate are the full-adder (FA) and the register. The transmission gate FA is based on the propagate-generate model and has the advantage of providing symmetric delays for the carry and sum outputs. The register is a positive-edge triggered master-slave D-flip-flop (DFF). For each latch, feedback transmission gate keepers are used instead of ratioed inverters for better noise immunity under low supply voltages. Asynchronous reset is added to clear the registers whenever the reset signal is logic low.

3.2 Critical Path Replica (CPR)

As briefly mentioned in Section 2.2, the critical path replica is a circuit which mimics the longest delay path of the core. The matching of its delay to that of the core is critical in achieving the optimal energy-delay-product under a workload constraint. Methods for tracking the critical path are: synthesized CPR based on empirical delay-matching of the worst case PVT corner simulations, direct or indirect CPM circuits shown in Fig. 2.6, and programmable lookup table for a pre-designed array of delay-cells based on post-layout measurements.

In this work, the CPR is implemented by an inverter chain with tunable delay cells. A chain of inverters are used instead of a direct mapping of core's critical path components to provide a near-50% duty cycle clock. The

number of inverters is calculated based on a first-order approximation using the Elmore delay formula for a resistor-capacitor (RC) network. The analysis assumes that each transistor can be modeled by an equivalent resistor and the load capacitance at each node can be lumped to a single capacitor coupled to ground. Furthermore, it is assumed that the intrinsic capacitance is approximately equal to the extrinsic capacitance for an inverter.


A simple illustration of a static CMOS logic example is shown in Fig. 3.1. Since all the gates are assumed to be the same size as that of a minimally sized inverter, the equivalent resistance can be represented by R , and the equivalent capacitance can be modeled by C or a multiple of C based on transistor sizing and the number of fanouts. For a chain of three inverters, assuming the input is driven by logic high, the equivalent lumped-RC model is shown in Fig. 3.1(a). The delay from the input to output can be approximated by the sum of the product of each resistor and its downstream capacitance. In this example, the propagation delay is proportional to the time constant τ ,

$$\tau = RC + RC + RC \approx 3RC \quad (3.1)$$


Similarly, a transmission gate buffered by inverters is shown in Fig. 3.1(b). The only difference between a transmission gate and an inverter are that the resistance of PFET and NFET are in parallel and that the downstream capacitance includes the capacitance of the subsequent node.

$$\tau = R(2C) + (2R \parallel R)C + RC \approx \left(3 + \frac{2}{3}\right)RC \quad (3.2)$$

This method can be extended to more complicated circuit paths in a similar manner. By dividing the resulting time constant of a critical path by the



(a)



(b)

Figure 3.1: Lumped RC models: (a) inverters and (b) transmission gate buffered by inverters.

time constant of a single inverter, $\tau_{INV} = RC$, an estimated number of inverters needed to represent the CPR can be derived. Note that the Elmore delay model is a simple yet pessimistic model, thus the resulting replica is slower than the critical path delay. Also note that the wire delay is not yet considered. The interconnect delay has different sensitivity to PVT variation and should be modeled by matching the wire dimensions rather than adding extra devices.

3.3 Digital Controlled Oscillator (DCO)

The core in each voltage domain is paired with an oscillator for supplying the varying frequency clock signal under varying supply voltage. By folding the CPR in a RO configuration, the clock period can track closely with the delay of the core. There are two considerations when incorporating the CPR in an oscillator. First, the oscillator requires an odd number of inverting stages; second, the clock period generated by a RO is twice of the oscillator delay. Thus, the estimated number of inverters of the CPR needs to be halved and rounded up to the nearest odd number.

The analysis in Section 3.2 provides a coarse indirect representation of the critical path, but does not consider mismatch due to process and temperature variation. Thus, a coarse tuning is needed to provide adjustable delay margin and ensure that the clock period is always greater than the critical path delay. A digital control is chosen over voltage control for the simplicity of matching the CPR and reliable bias under voltage transient and different voltage domains.

The main building block of the DCO is a delay cell, as shown in Fig. 3.2. Each delay cell consists of a long path and a short path. The selection of




Figure 3.2: Delay cell.

the paths is controlled by the tri-state buffers, which can also be viewed as a multiplexer. Each delay cell uses only a single-bit control to minimize the capacitive loading at the output and to simplify delay analysis.

Figure 3.3 shows the equivalent RC-network for the delay cell when the control signal is logic low or logic high. When control signal c is ‘0’, the short/fast path is selected; when control signal is ‘1’, the long/slow path is selected. Assuming that the output is buffered by an inverter, the time constant of the fast and slow paths can be approximated by the following:

$$\tau = \begin{cases} R(3C) + \frac{2}{3}R\left(\frac{3}{2}C\right) = 4RC \approx 4\tau_{INV} & c = \text{‘0’} \\ R\left(\frac{3}{2}C\right) + 3RC + R\left(1 + \frac{3}{2}\right)C + \frac{2}{3}R\left(\frac{3}{2}C\right) = 8RC \approx 8\tau_{INV} & c = \text{‘1’} \end{cases}$$

This shows that the difference between the fast and slow paths is four inverter delays and that the addition of each delay cell is equivalent to the addition of four or eight inverters. This approximation is useful when incorporating the delay cells in the CPR inverter chain.

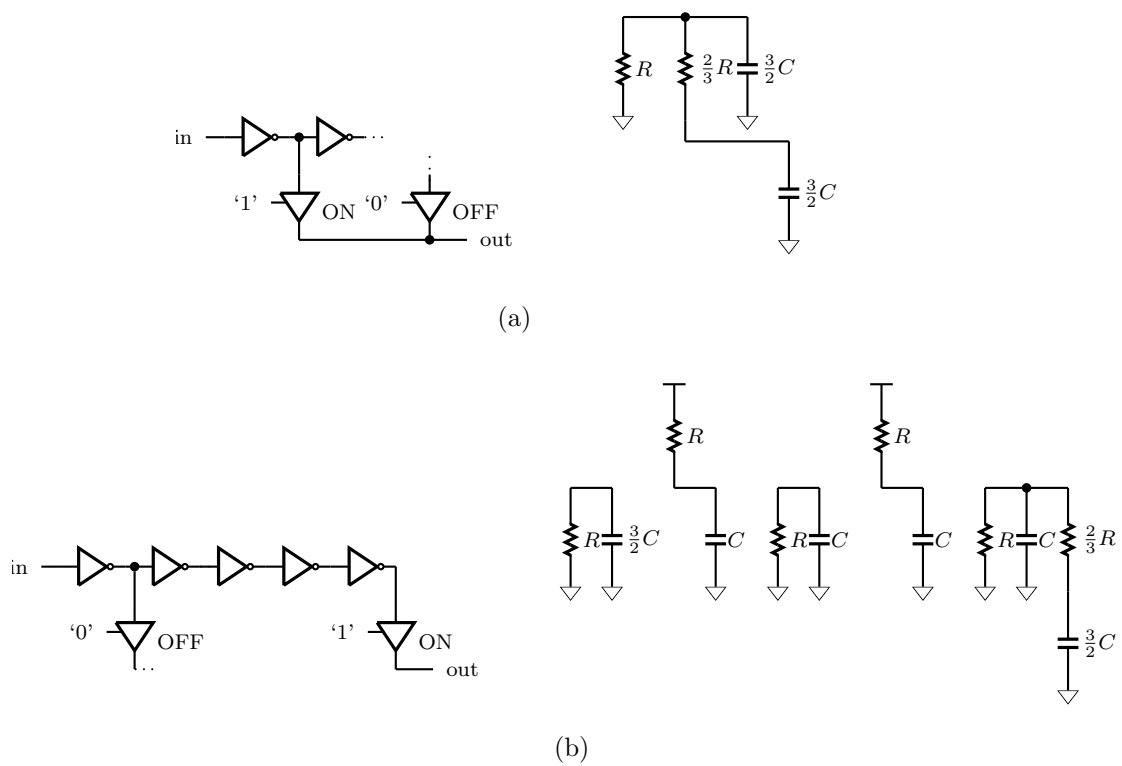


Figure 3.3: Lumped RC model of delay cell: (a) fast path: $c = '0'$ and (b) slow path $c = '1'$.

CHAPTER 4

DIGITAL CIRCUIT BLOCKS

This chapter describes the implementation of the digital blocks shown in Fig. 1.1. The CVFS core consists of two replicated FIR filters, each paired by a tunable oscillator. The input controller provides the input vectors and the control signals to the core. The following sections describes each of the blocks in details.

4.1 Multiply-Accumulate (MAC) Finite Impulse Response (FIR) Filter

The core is an 8-bit 8-tap MAC FIR filter. The hardware data flow graph is shown in Fig. 4.1. The input vector $x[n]$ and filter coefficient $h[n]$ are unsigned 8-bits and are generated by the input controller. The feedback registers and the output registers are both unsigned 19-bits to avoid overflow. The output registers update the value of $y[n]$ every eight clock cycle. The following expression describes an 8-tap FIR filter:

$$y[n] = \sum_{k=0}^7 h[k]x[n - k] \quad (4.1)$$

The MAC unit is implemented using the standard carry-save multiplier. The feedback term is combined in the vector merging adder of the multiplier using a carry-save adder. The MAC unit schematic is shown in Fig. 4.2,




Figure 4.1: MAC FIR data flow graph.

where x_i and h_i represent the i -th bit of $x[n-k]$ and $h[k]$ respectively, fb_j is the j -th bit of the feedback sum, and sum_j is the j -th bit of the MAC result where $i \in [0 : 7]$ and $j \in [0 : 18]$. Each square block represents an FA. The shaded blocks can be replaced by half-adders, and the blocks with thick outlines each contain an AND gate. The arrow shows the critical path of the MAC unit.

The 8-tap FIR filter is an 8-folded architecture of the MAC unit. Each sample period has eight clock cycles. At the end of each sample period, the output register is updated and the feedback register is reset. The output registers are gated by a MUX to ensure that the output updates synchronously with the clearing of the feedback registers. Since the inputs are shared by the cores in the two voltage domains, AND gates are added to block the inputs so that no switching occurs when the core is idle. Also, a MUX is added to every register that requires data transfer between the cores when transitioning from one voltage domain to another. The switching registers are shown in Fig. 4.3, where sw is the switch indication signal, and Dsw is the level-shifted Q from the corresponding registers in the other voltage domain.




Figure 4.2: MAC unit schematic.




Figure 4.3: Switching register.

4.2 Tunable Oscillator

The oscillator accompanying each core is a DCO, which includes the CPR of the MAC unit. The critical path delay of the MAC unit is highlighted in Fig. 4.2 and is approximated by

$$T_{MAC} = T_{NAND} + T_{INV} + 19T_{FA} \quad (4.2)$$

The RC lumped model of a FA can be derived using analysis similar to that of Section 3.2. Since the critical path is the propagation of the carry signal, the analysis only considers the path from carry-in to carry-out. Thus, assuming each inverter and transmission gate has the same input and output capacitances of approximated $C/2$, the FA delay in terms of the delay of an equivalent inverters can be derived as

$$\tau_{FA} = R \left(\frac{3}{2} + \frac{3}{2} \right) C + R \left(\frac{3}{2} + \frac{3}{2} \right) C + \frac{2}{3} R \left(\frac{3}{2} C \right) = 7RC \approx 7\tau_{INV} \quad (4.3)$$

Substituting (4.3) in (4.2), the critical path delay can be represented by

$$\tau_{MAC} = \tau_{NAND} + (1 + 7 \cdot 19)\tau_{INV} \approx \tau_{NAND} + 134\tau_{INV} \quad (4.4)$$

The result shows that the critical path can be approximated by a NAND gate and a chain of 134 inverters. The NAND gate is set aside for implementing the enable signal and also serves as the inverting stage of the ring oscillator. As discussed previously, in order to match the delay of an inverter chain to the clock period of a ring oscillator, the number of inverters needs to be halved and rounded up to the nearest inverting stages. Here, 68 inverters are employed in the inverter chain.




Figure 4.4: DCO schematic.

The tuning function is implemented using the delay cells described in Section 3.3. The control bits are arbitrarily chosen to be 8-bits, where the most significant 4-bits are by default logic low and the least significant 4-bits are by default logic high. This allows the inverter chain to be scaled up or down by 16 inverters. The tunable range is not only beneficial for ensuring adequate delay margin, but also enables testing for frequency overscaling. From previous analysis, the delay of eight delay cells represents the delay of 48 inverters, thus reducing the inverter chain to 20 inverters.

The top-level schematic of the DCO is shown in Fig. 4.4, which consists of an inverter chain of 20 inverters, 4 default-high delay cells, 4 default-low delay cells, a NAND gate for the enable signal, and an output buffer stage to enhance clock driving strength.

4.3 Input Controller

The input controller generates the input vectors and control signals for the MAC FIR and DCO. It also handles the peripheral signals for initializing registers and testing. The detection of supply voltage and the regulation of switching events are managed by the VRM and level-shifter block and are not part of the input controller. This block operates at 1.2V nominal supply voltage and its outputs are shared by the core in both voltage domains. The input controller is organized to four sub-blocks: x-register, h-register,

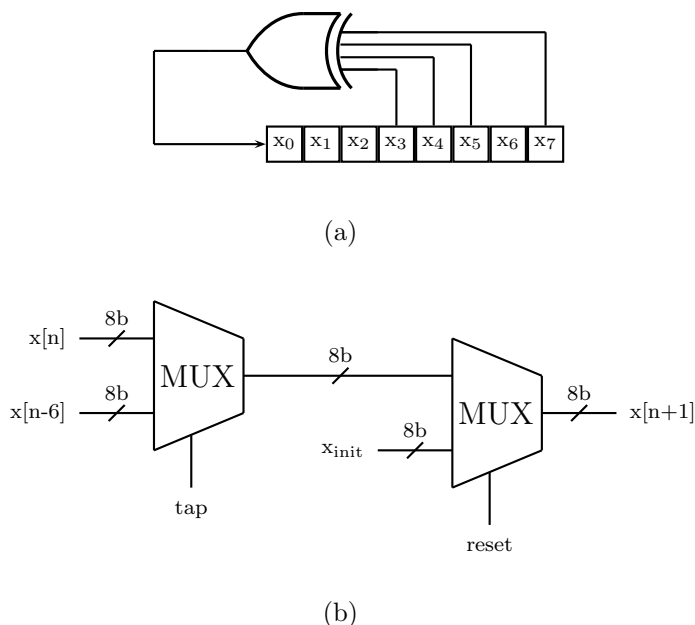


Figure 4.5: X-register: (a) LFSR and (b) simplified schematic.

control-register, and oscillator-register.

The x-register is primarily an 8-bit *linear feedback shift register* (LFSR) with feedback taps shown in Fig. 4.5(a), where $x[n] = (x_7x_6x_5\dots x_0)_2$. The LFSR is a pseudorandom generator that cycles through all $(2^8 - 1) = 255$ non-zero values; each subsequent value is determined from the previous value. In other words, $x[n] = x[n + 255]$ and $x[n]$ is determined from $x[n - 1]$, where n is any integer index. The x-register also contains two additional 8-bit registers; one stores the user-defined initial value x_{init} and another one stores the value of $x[n - 6]$ for retrieval in the subsequent computation. Figure 4.5(b) illustrates the x-register. The signal *tap* is generated from the control-register and is set at the end of every computation period.

Figure 4.6 shows that the h-register is an 8-bit circular shift register, where $h[n] = (h_7h_6h_5\dots h_0)_2$. The MSB is shifted to the LSB every clock cycle, and the value repeats itself every eight clock cycles. The h-register also contains another 8-bit register h_{init} for storing the initial value, similar to that of the




Figure 4.6: H-register: (a) circular shift register and (b) simplified schematic.

x-register.

The oscillator-register is an 8-bit register that stores the control bits of the DCO. It is identical to the initial value registers in the x-register and h-register. The four MSB are internally inverted for default logic high and the four LSB are for default logic low.

The control-register translates all the peripheral signals to internal control signals. It consists primarily of two counters. The first counter keeps count of the number of clock cycles within each computation period and generates control signals. The control signal indicates when to update the output register, clear the feedback register, and save $x[n - 6]$. The second counter is used for testing. The counter counts 256 computations (2048 clock cycles), disables all digital circuitry, and reinitializes all the registers in the input controller. A monostable circuit is used to generate a short pulse to reset the second counter and re-enables the circuit.

CHAPTER 5

SIMULATION RESULTS

5.1 Design Verification

The core was verified under discrete and continuous voltage scaling. Under discrete voltages, the FIR filter was simulated at the supply voltage extremes of 1.2V and 460mV. For continuous voltage scaling operation, the supply was modeled by a capacitor discharging from 700mV to 500mV. The computation results are logged and compared to the expected results using MATLAB. Figure 5.1 shows an example set of the vectors generated by the input controller and the result of the FIR filter. The output $y[n]$ is zero for the first eight clock cycles and its value is updated every eight clock cycles. The vertical line marks the 2048th (256 x 8) clock cycle, where the sequence repeats itself.

A 3-bit counter is employed to easily demonstrate that the switching register can correctly exchange data between the two voltage domains. Figure 5.2 shows the transition of clock between stacked voltage domains. Fig. 5.2(b) and Fig. 5.2(c) represent the simulated data corresponding to the high and low voltage domains, respectively, where the output of each counter is denoted by $(q_2q_1q_0)_2$. In each voltage domain, data switching occurs when the sw signal is logic high. The counter resumes normal operation when the sw signal is logic low.

The matching of the clock periods to the critical path delay is important in ensuring correct circuit operation and minimizing energy consumption under

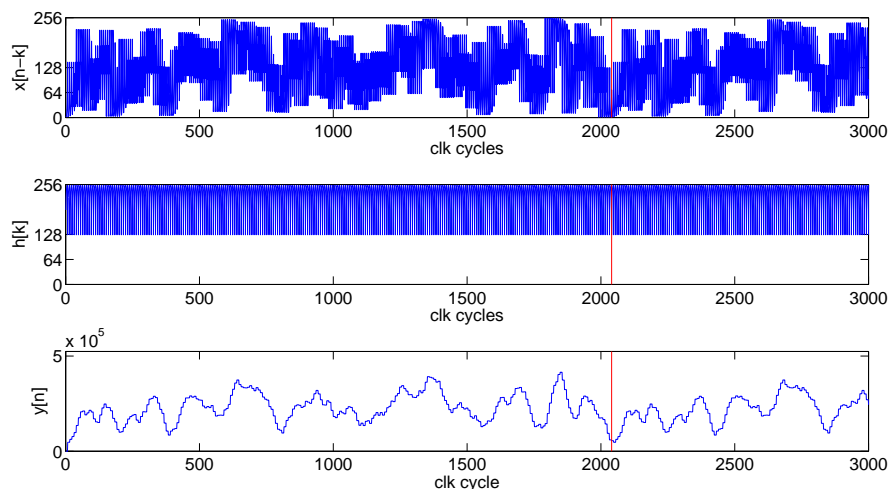



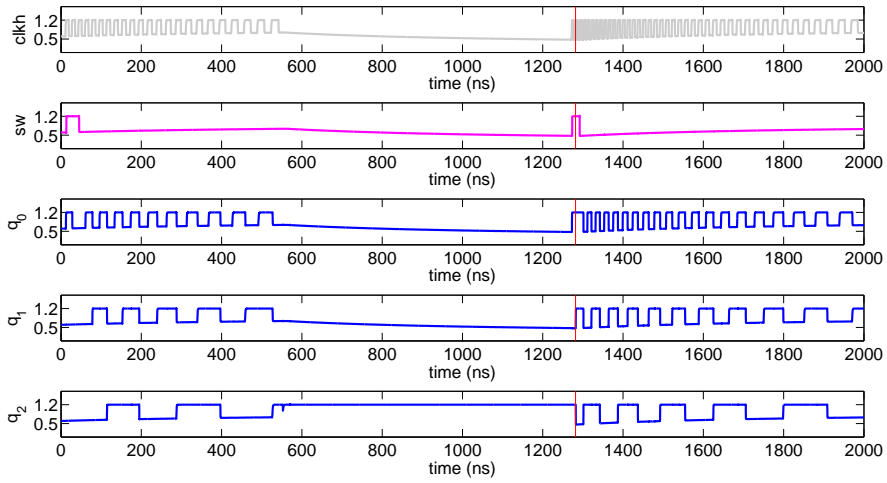
Figure 5.1: MAC FIR filter computation.

variable supply voltage. The DCO and MAC unit are simulated over five process corners. The measured results are shown in Fig. 5.3, where the clock frequency tracks closely with the MAC’s delay. Also, the clock period is always greater than the critical path delay, which is desirable for error-free operation.

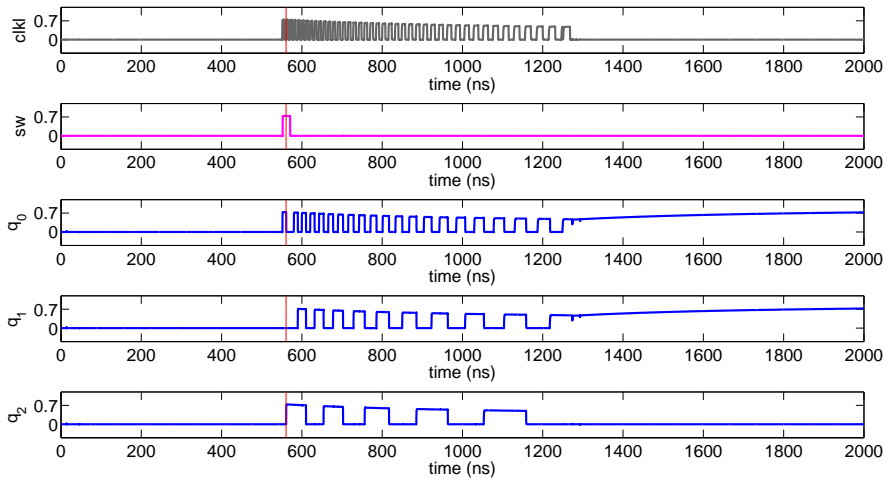
Additionally, the clock period is adjustable by an 8-bit control to compensate for unpredictable variations during fabrication. The DCO tuning of the delay for the TT corner is shown in Fig. 5.4. Each control bit provides approximately 5% increase or reduction of the nominal delay. Setting any of the four most significant bits (MSB) reduces the clock frequency, while setting any of the four least significant bits (LSB) increases the clock frequency. The design of DCO is implemented using triple-well devices, thus the same tuning behavior applies to both voltage domains.



(a)



(b)



(c)

Figure 5.2: Data switching between voltage domains: (a) voltage domain transitions, (b) high voltage domain, and (c) low voltage domain.

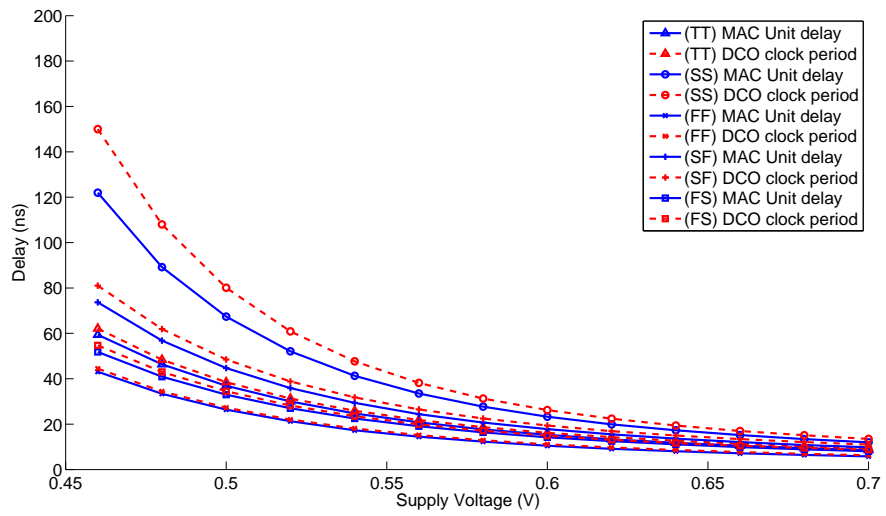


Figure 5.3: DCO clock period vs. MAC critical path delay.




Figure 5.4: DCO tuning.

5.2 Testing Procedure

Figure 5.5 shows the test procedure for initializing the DCO control bits as well as the x_{init} and h_{init} register values. The top three signals are inputs and the remaining two are outputs. Upon startup, *reset* clears all register values, including the registers in the controller. This pin is asynchronous and enabled low. Initialization is needed by pulling the *load* signal high. The *load clk* is an external clock that loads one bit per positive edge; it is only effective when *load* is high. The waveform shows eight pulses since all the initial registers are 8-bits. The internal clock is enabled immediately after the load signal pulls low.




Figure 5.5: Load initial register values.

For testing, the circuit is automatically disabled after 256 computation periods. As shown earlier, this refers to the point at which the data sequence repeats itself. For simplicity of demonstration, the input controller is modified to disable after four computations and illustrated in Fig. 5.6. The tap signal indicates the end of each computation period, where each computation period is eight clock cycles. The circuit remains idle until the *rerun* signal

is asserted by the user. The falling edge of the *rerun* signal triggers a pulse to reactivate the circuit. This allows the user to rerun the same test without the need to reinitialize the registers. The output signal *active* indicates the length of time which the circuit is active. This can be used to computing energy per operation.




Figure 5.6: Rerun test vectors.

Due to the limitation of the number of pins available for the chosen chip package, the output vectors, namely the 19-bit $y[n]$ from each voltage domain, are truncated to 8-bit.

CHAPTER 6

CONCLUSION

6.1 Summary

This thesis presents an overview of energy saving methodologies based on supply voltage reduction and introduces the concept of CVFS. Unlike conventional schemes, which rely heavily on stringent voltage regulation, CVFS explores circuit operation under continuous voltage variation. A simple first-order RC-model for constructing the CPR was proposed. The incorporation of the CPR in a DCO enables an indirect monitoring of the circuit delay under instantaneous supply voltages. The results demonstrate a close tracking of the clock period to the critical path delay across process corners, while maintaining sufficient delay margin to avoid timing violation.

This work involves the implementation of a prototype chip using IBM 130nm technology for near-threshold operation. The contribution of this thesis is the design of the FIR filter, DCO, and the logics for on-chip vector generation and testing. The operation of the digital components has been verified under continuous voltage scaling. The digital blocks serve as the computational core for a system-level energy optimization scheme.

REFERENCES

- [1] A. Chandrakasan, S. Sheng, and R. Brodersen, “Low-power cmos digital design,” *Solid-State Circuits, IEEE Journal of*, vol. 27, no. 4, pp. 473–484, Apr. 1992.
- [2] A. Chandrakasan and R. Brodersen, “Minimizing power consumption in digital cmos circuits,” *Proceedings of the IEEE*, vol. 83, no. 4, pp. 498–523, Apr. 1995.
- [3] B. Calhoun, A. Wang, and A. Chandrakasan, “Device sizing for minimum energy operation in subthreshold circuits,” in *Custom Integrated Circuits Conference, 2004. Proceedings of the IEEE 2004*, Oct. 2004, pp. 95–98.
- [4] J. Kao, M. Miyazaki, and A. Chandrakasan, “A 175-mv multiply-accumulate unit using an adaptive supply voltage and body bias architecture,” *Solid-State Circuits, IEEE Journal of*, vol. 37, no. 11, pp. 1545–1554, Nov. 2002.
- [5] A. Wang and A. Chandrakasan, “A 180-mv subthreshold fft processor using a minimum energy design methodology,” *Solid-State Circuits, IEEE Journal of*, vol. 40, no. 1, pp. 310–319, Jan. 2005.
- [6] 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*, vol. 50, no. 4.5, pp. 469–490, July 2006.
- [7] V. Gutnik and A. Chandrakasan, “Embedded power supply for low-power dsp,” *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 5, no. 4, pp. 425–435, Dec. 1997.
- [8] T. Burd and R. Brodersen, “Design issues for dynamic voltage scaling,” in *Low Power Electronics and Design, 2000. ISLPED '00. Proceedings of the 2000 International Symposium on*, 2000, pp. 9–14.
- [9] T. Burd, T. Pering, A. Stratakos, and R. Brodersen, “A dynamic voltage scaled microprocessor system,” *Solid-State Circuits, IEEE Journal of*, vol. 35, no. 11, pp. 1571–1580, Nov. 2000.

- [10] T. Kuroda, K. Suzuki, S. Mita, T. Fujita, F. Yamane, F. Sano, A. Chiba, Y. Watanabe, K. Matsuda, T. Maeda, T. Sakurai, and T. Furuyama, "Variable supply-voltage scheme for low-power high-speed cmos digital design," *Solid-State Circuits, IEEE Journal of*, vol. 33, no. 3, pp. 454–462, Mar. 1998.
- [11] L. Yuan and G. Qu, "Analysis of energy reduction on dynamic voltage scaling-enabled systems," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 24, no. 12, pp. 1827–1837, Dec. 2005.
- [12] R. Mullins and S. Moore, "Demystifying data-driven and pausable clocking schemes," in *Asynchronous Circuits and Systems, 2007. ASYNC 2007. 13th IEEE International Symposium on*, March 2007, pp. 175–185.
- [13] B. Calhoun and A. Chandrakasan, "Ultra-dynamic voltage scaling using sub-threshold operation and local voltage dithering in 90nm cmos," in *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, Feb. 2005, pp. 300–599 Vol. 1.
- [14] S. Zhang, "Compute VRM: Integrating power converter and compute core for energy-efficient computing," unpublished.
- [15] J. Kwong and A. Chandrakasan, "Variation-driven device sizing for minimum energy sub-threshold circuits," in *Low Power Electronics and Design, 2006. ISLPED'06. Proceedings of the 2006 International Symposium on*, Oct. 2006, pp. 8–13.
- [16] S. Fisher, A. Teman, D. Vaysman, A. Gertsman, O. Yadid-Pecht, and A. Fish, "Digital subthreshold logic design - motivation and challenges," in *Electrical and Electronics Engineers in Israel, 2008. IEEEI 2008. IEEE 25th Convention of*, Dec. 2008, pp. 702–706.
- [17] J. Park and J. Abraham, "A fast, accurate and simple critical path monitor for improving energy-delay product in dvs systems," in *Low Power Electronics and Design (ISLPED) 2011 International Symposium on*, Aug. 2011, pp. 391–396.
- [18] M. Elgebaly and M. Sachdev, "Variation-aware adaptive voltage scaling system," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 15, no. 5, pp. 560–571, May 2007.
- [19] B. Calhoun and D. Brooks, "Can subthreshold and near-threshold circuits go mainstream?" *Micro, IEEE*, vol. 30, no. 4, pp. 80–85, July-Aug. 2010.

- [20] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, “Analysis and mitigation of variability in subthreshold design,” in *Low Power Electronics and Design, 2005. ISLPED '05. Proceedings of the 2005 International Symposium on*, Aug. 2005, pp. 20 – 25.
- [21] D. Bol, R. Ambroise, D. Flandre, and J.-D. Legat, “Interests and limitations of technology scaling for subthreshold logic,” *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 17, no. 10, pp. 1508 –1519, Oct. 2009.
- [22] A. Sassone, A. Calimera, A. Macii, E. Macii, M. Poncino, R. Goldman, V. Melikyan, E. Babayan, and S. Rinaudo, “Investigating the effects of inverted temperature dependence (itd) on clock distribution networks,” in *Design, Automation Test in Europe Conference Exhibition (DATE), 2012*, March 2012, pp. 165 –166.
- [23] N. Verma, J. Kwong, and A. Chandrakasan, “Nanometer mosfet variation in minimum energy subthreshold circuits,” *Electron Devices, IEEE Transactions on*, vol. 55, no. 1, pp. 163 –174, Jan. 2008.
- [24] A. Wang, A. Chandrakasan, and S. Kosonocky, “Optimal supply and threshold scaling for subthreshold cmos circuits,” in *VLSI, 2002. Proceedings. IEEE Computer Society Annual Symposium on*, 2002, pp. 5 –9.
- [25] B. Calhoun, A. Wang, and A. Chandrakasan, “Modeling and sizing for minimum energy operation in subthreshold circuits,” *Solid-State Circuits, IEEE Journal of*, vol. 40, no. 9, pp. 1778 – 1786, Sept. 2005.
- [26] K. Roy, J. Kulkarni, and M.-E. Hwang, “Process-tolerant ultralow voltage digital subthreshold design,” in *Silicon Monolithic Integrated Circuits in RF Systems, 2008. SiRF 2008. IEEE Topical Meeting on*, Jan. 2008, pp. 42 –45.
- [27] J. Keane, H. Eom, T.-H. Kim, S. Sapatnekar, and C. Kim, “Subthreshold logical effort: a systematic framework for optimal subthreshold device sizing,” in *Design Automation Conference, 2006 43rd ACM/IEEE*, 0-0 2006, pp. 425 –428.