VARIATION-AWARE ADAPTATION OF HETEROGENEOUS VOLTAGE REGULATION SYSTEMS

A Thesis

by

JIANHAO CHEN

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Chair of Committee,	Peng Li
Committee Members,	Eun Jung Kim
	Jiang Hu
Head of Department,	Miroslav M. Begovic

August 2020

Major Subject: Computer Engineering

ABSTRACT

The Power delivery network (PDN) is an electrical system that provides the supply voltage to transistors on a silicon chip. With the development of CMOS technology, the transistor density increases exponentially [1]. However, the supply voltage and the threshold voltage of transistors have limited room to be further reduced, which leads to a huge increment in chip power density. Besides, the areas of chips like server-class CPUs are also increasing in order to integrate more transistors on a chip and fully leverage the benefits of parallelism. Designing PDNs for these power-hungry and large-scale chips has become a challenge.

In recent years, PDNs with distributed on-chip voltage regulation have emerged as a novel design solution to deliver power to transistors with high efficiency as well as ensured quality in many research works and commercial products. With distributed on-chip voltage regulators (VRs), PDNs with distributed on-chip voltage regulation reduce the demands on power pins. They also provide fast suppression of voltage noise thanks to the reduced proximity between the VRs and current loads. Recently, heterogeneous voltage regulation (HVR) architecture was proposed to provide high-efficiency power delivery and high-performance voltage regulation by utilizing heterogeneous on-chip and off-chip VRs [2]. However, the effects of inevitable process and temperature variations on the system efficiency and power integrity were not evaluated.

In this thesis, a variation-aware adaption of the HVR architecture has been proposed to adapt HVR to systems with process and temperature variations using variation tracking circuits and variation-aware control policy. It improves system efficiency by 1.13% on average, while ensuring a given power integrity specification in the presence of process variations. In addition, the HVR architecture is found to possess some degree of robustness to temperature and process variations.

DEDICATION

To my parents

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a thesis committee consisting of Professor Peng Li and Processor Jiang Hu of the Department of Electrical and Computer Engineering and Professor Eun Jung Kim of the Department of Computer Science and Engineering.

Joseph Riad contributes to the work in Chapter.3 by proposing the circuit designs and variation tracking mechanisms and analyzing effects of variations on off-chip converter control in Section.3.4.2

All other work conducted for the thesis was completed by the student independently.

Funding Sources

Graduate study was supported by a fellowship from Texas A&M University.

TABLE OF CONTENTS

Page

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

TABLE Page

1. INTRODUCTION

1.1 Overview

In the past few decades, the number of transistors integrated on a silicon chip has been increasing exponentially with CMOS technology scaling [1]. However, the supply voltage and the threshold voltage of transistors have little room to be further reduced, which leads to a continuous increase in power density. Meanwhile, with the development of parallel computing, die sizes of the chips for parallel computing, such as multi-core CPUs and general-purpose GPUs (GPGPUs), have been growing in order to further exploit the benefits of parallelism. Because of large power consumption, large die sizes and constraints of power pins, powering these large-scale and power-hungry chips and maintaining uniform power integrity across the chip have become a challenge.

The Power Delivery Network (PDN) is an electrical system that provides and regulates the supply voltage for on-chip transistors. Typically a PDN converts the supply voltage provided by the power source on a printed circuit board (PCB), e.g., $12\,\mathrm{V}$, to the supply voltage for transistors, e.g., $1\,\mathrm{V}$, using voltage regulators (VRs). Then it delivers the current to the chip through metal traces on PCB and package and controlled collapse chip connection (C4) bumps, and distributes the supply voltage to transistors using an on-chip power grid.

Except for power delivery, a PDN is also responsible for regulating its output voltage and maintaining uniform power integrity across the chip. Because of supply voltage ripples, load current fluctuations, and PDN parasitics, the output voltage of a PDN, which is also the supply voltage of transistors, will fluctuate. The deviation of the actual supply voltage from its nominal value is called voltage noise. The voltage noise can lead to long-term reliability problems like electromigration [4]. Furthermore, large voltage noise can cause serious timing violations in a circuit [5].

The delay of a circuit is related to its supply voltage. The undesirable voltage noise introduces voltage droops to the supply voltage and increases the circuit delay. A voltage emergency (VE)

occurs when the supply voltage falls below a certain threshold and the circuit delay exceeds the requirement, which will cause the logic's malfunction. An example of VEs is shown in Figure.1.1. The nominal supply voltage is $1\,\mathrm{V}$ as shown by the black dot line. Because of the non-ideal effects of the PDN and fluctuations of the workload, the actual supply voltage fluctuates, as shown by the solid blue curve. Assuming that the minimum required supply voltage is $950\,\mathrm{mV}$, as shown by the red dash line, two VEs occur when the supply voltage drops below $950\,\mathrm{mV}$, as shown in the two red circles.

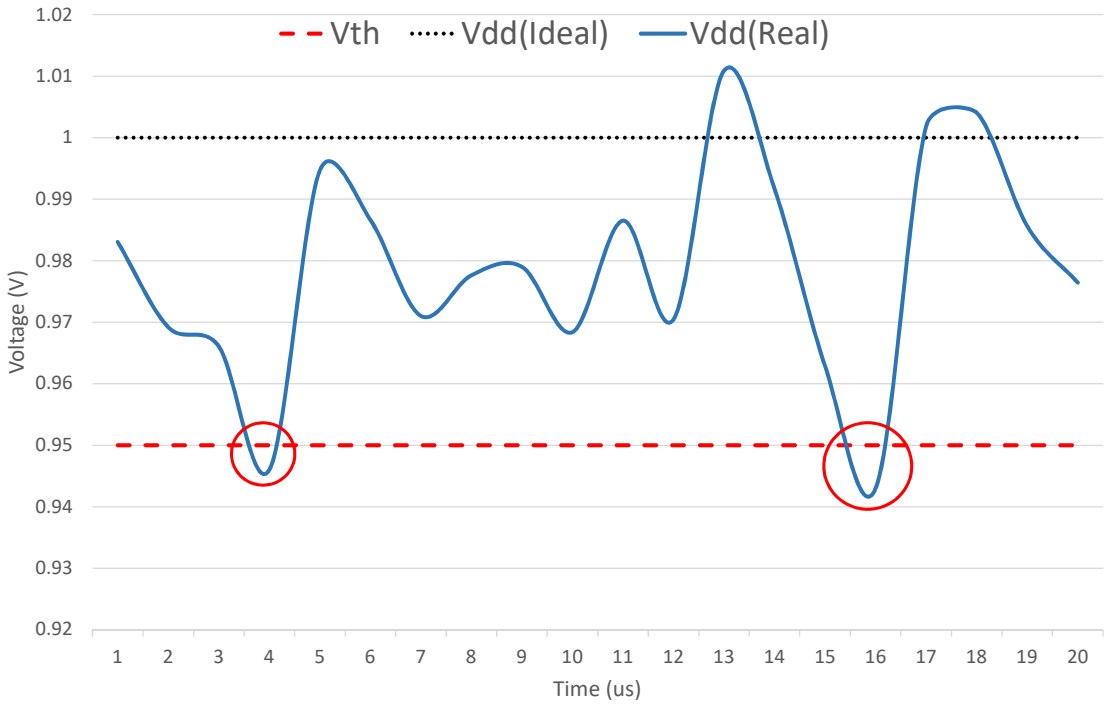

Figure 1.1: An example of voltage emergencies (VEs)

Apart from power integrity, another critical performance metric for a PDN is efficiency. Because of the intrinsic resistance in wires and transistors, power loss exists in metal traces and VRs in a PDN, which increases the total power consumption of a chip and poses challenges to low-power designs. It also generates additional heat and may cause thermal emergencies [6].

Many research works and commercial products devise complex PDN structures and control mechanisms to improve PDN performance. An emerging practice is to exploit on-chip integrated voltage regulators and design a PDN with distributed on-chip voltage regulation. With increased proximity to current loads, on-chip voltage regulators usually have fast response time [2] to workload changes, which makes fine-grain voltage control and regulation possible. Three types of on-chip integrated voltage regulators are commonly used: Low Dropout Voltage Regulators (LDO), Switching Capacitor (SC) voltage regulators, and buck converters.

The LDO is a type of linear voltage regulator, which is area-efficient and can achieve sub-ns response time [7], so it is suitable for fast voltage noise suppression. However, the power loss of an LDO is proportional to its dropout voltage. Thus it is efficient only when the output voltage is close to its input voltage [2]. Buck converters and SC converters usually have large footprints because of their bulky integrated inductors and capacitors. They have a relatively long response time to load current changes compared to LDO, while their conversion efficiencies could remain high in a wide range of conversion ratios with maximum efficiencies up to 90%[8, 9, 10, 11]. Recent years have seen the applications of on-chip integrated voltage regulators in commercial products. For example, Intel Haswell processors use Fully Integrated Voltage Regulator (FIVR) [12] and IBM POWER8 processor employs 1,764 on-chip distributed linear voltage regulators [13].

Several research works propose various PDN architectures and optimization methods to enhance the performance of PDNs with on-chip voltage regulation. [14] formulates the placement of on-chip voltage regulators and decoupling capacitors as an optimization problem and solves it based on the workload distribution. On the contrary, [14] optimizes the sizes and layouts of SC converters according to the workload to reduce power loss. However, a PDN optimized for a particular workload may shift away from its optimal operation point when the workload changes, resulting in low system efficiency. Therefore, some research works [15, 16, 17, 18] propose reconfiguration methods to adapt the PDN to dynamic workloads. For example, [15] proposes reconfigurable interconnections of on-chip voltage regulators to dynamically power each core in a chip multiprocessor(CMP). [16] applies a simple control policy to dynamically adjust the number of active

on-chip and off-chip voltage regulators. These researches propose ideas of the dynamic control of the PDN, while failed to explore the adoption of heterogeneous voltage regulators.

Recently a research work proposes heterogeneous voltage regulation (HVR) architecture with multiple power domains, multi-stage and heterogeneous VRs and distributed on-chip voltage regulation to provide high-performance power delivery [2]. A workload-aware control policy is also proposed to maintain a high system efficiency under different workloads by tuning the intermediate voltages in the conversion chain and gating some of the regulators according to the workload. The HVR architecture reduces system energy cost by up to 23.9% compared with a static 2-stage PDN with the same level of power integrity [2].

However, on-chip VRs are vulnerable to undesirable process and temperature variations, so the system performance of distributed on-chip voltage regulation might be affected. [6] considers the effects of temperature variations on a PDN. As described in [6], the power loss of on-chip voltage regulators may create hotspots and cause thermal emergencies. A collection of thermal-aware runtime policies is proposed to selectively gate the on-chip voltage regulators to reduce voltage noise and temperature gradient and improve efficiency at the same time.

As for the HVR architecture proposed in [2], the effects of process and temperature variations on the system performance were not considered. Its workload-aware control policy depends on the efficiency characteristics of buck converters and LDOs, which will shift from the nominal values under process and temperature variations. The nominal control policy may generate inaccurate control signals, which may compromise power integrity and degrade system efficiency.

Therefore, in this thesis, a variation-aware adaption of HVR architecture has been proposed to adapt HVR architecture to systems with process and temperature variations. The rest of this thesis is organized as follows: Chapter.2 introduces the HVR architecture. Chapter.3 introduces the variation-aware adaption of HVR, demonstrates its performance and analyzes the effects of variations on the original HVR control policy.

2. WORKLOAD-AWARE HETEROGENEOUS VOLTAGE REGULATION

2.1 Overview

This chapter introduces the workload-aware heterogeneous voltage regulation (HVR) architecture proposed in [2]. Figure.2.1 shows the structure of HVR architecture. It has three stages of VRs, including off-chip buck converters, on-chip buck converters, and on-chip LDOs, to capitalize on both the superior efficiency of buck converters and the excellent regulation performance of LDOs.

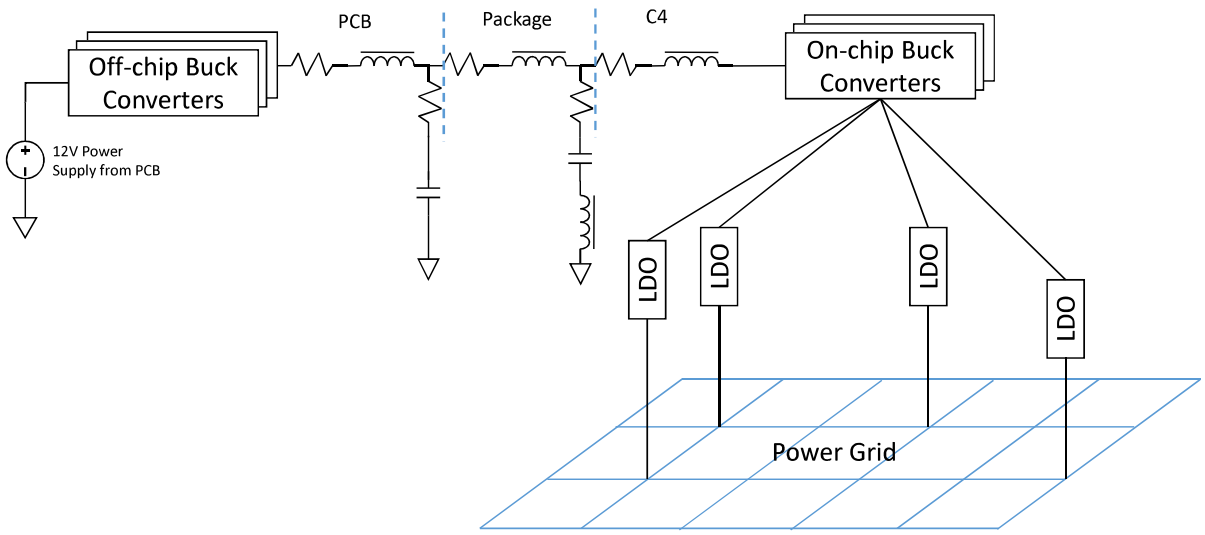

Figure 2.1: The structure of HVR architecture

The off-chip buck converter cluster in the first stage converts the supply voltage on PCB and powers the on-chip buck converters in the second stage through package and C4 connections. The on-chip PDN is divided into multiple power domains, each of which is powered by a cluster of on-chip buck converters and distributed LDOs (Figure.2.1 shows one power domain.). The on-chip buck converters further step down the supply voltage, and the distributed LDOs drive the load

circuits and regulate the supply voltage to transistors.

This chapter first describes the voltage regulator characteristics and PDN modelings for PDN performance evaluation. Then the HVR architecture proposed in [2] and its workload-aware control policy are introduced.

2.2 Power Grid Model

The power grid is a network of metal wires on a silicon chip that distributes the current across the chip to transistors. It contains a VDD network and a GND network connected to the voltage supply and the ground, respectively. A modern processor contains billions of transistors [19], so its power grid has billions of output ports to power the circuit. Building a circuit model directly for all the wires in such a massive power grid, extracting load currents from the hardware for every power grid node, and analyzing its output voltage are computationally expensive. Thus a lumped model for the on-chip power grid is adopted for PDN performance evaluation. On the one hand, a lumped power grid model reduces the number of nodes in circuit simulation, and the block-level workload currents can be easily extracted using architecture simulators, as described in Section 3.3, which reduces the computation cost significantly. On the other hand, a properly-sized lumped model is capable of demonstrating the spatial workload distribution and voltage noise for the purpose of power integrity analysis.

RC and RLC models are usually used for PDN simulation [4, 2, 20], where R, L, and C represent parasitic resistance, inductance, and capacitance in metal stacks. The block-level workload currents of the PDN can be extracted using architectural simulators, e.g., GEM5[21] and McPAT[22], and modeled as current sources draining currents from VDD network to GND network.

In this work, an RC power grid model similar to [20] is used. The details about block-level workload current extraction are described in Section.3.3. Figure.2.2 and Figure.2.3 shows examples of a VDD network and a GND network in a RC power grid model with a size of 5 by 3.

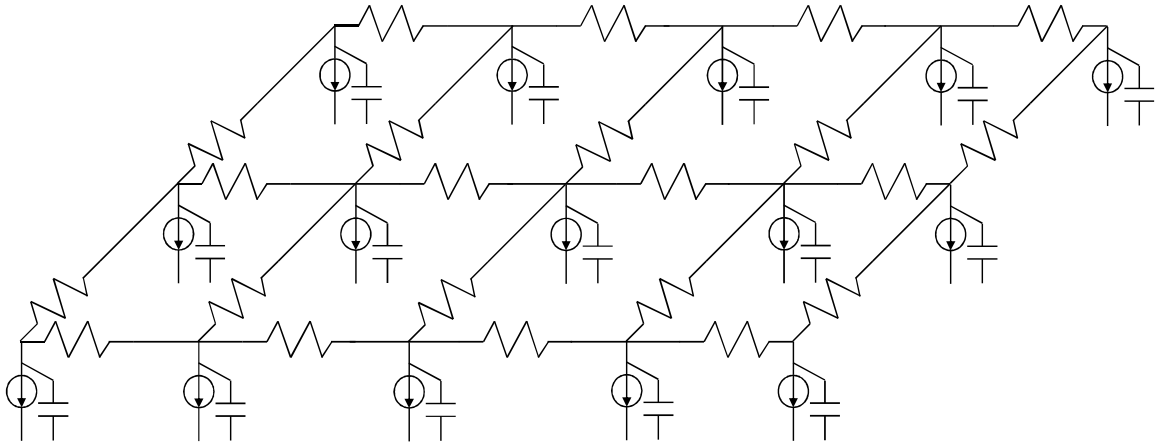

Figure 2.2: The VDD network of a power grid

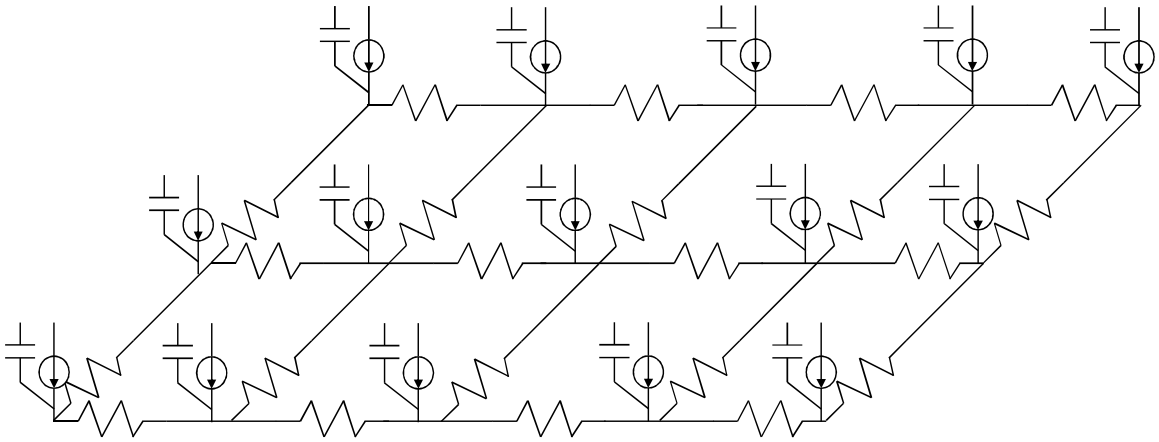

Figure 2.3: The GND network of a power grid

2.3 Voltage Regulators

Voltage regulators are critical components in a PDN because they determine the power integrity and system efficiency of the PDN. In this section, voltage regulators used in HVR architecture [2] are introduced, including off-chip buck converters, on-chip buck converters, and on-chip LDOs. The buck converters can maintain high conversion efficiency within a broad conversion ratio range,

but they have relatively long response time, and their on-chip integration of can be expensive. In contrast, the efficiency of an LDO is relatively low, but they are area-efficient and have sub-ns response time to the load current change. The comparison of these voltage regulators is shown in the table below [2].

	Response Time	Area	Efficiency
On-chip LDO	Sub-ns	Small	Low with large step down ratio
On-chip Buck	10's of ns	Medium	Medium
Off-chip Buck	10's of us	Large	High

Table 2.1: Comparison of LDOs, on-chip buck converters and off-chip buck converters.

2.3.1 Buck Converters

The buck converter is a type of DC-to-DC voltage converter. The Figure.2.4 shows the basic structure of a buck converter. Its output voltage is controlled through pulse width modulation

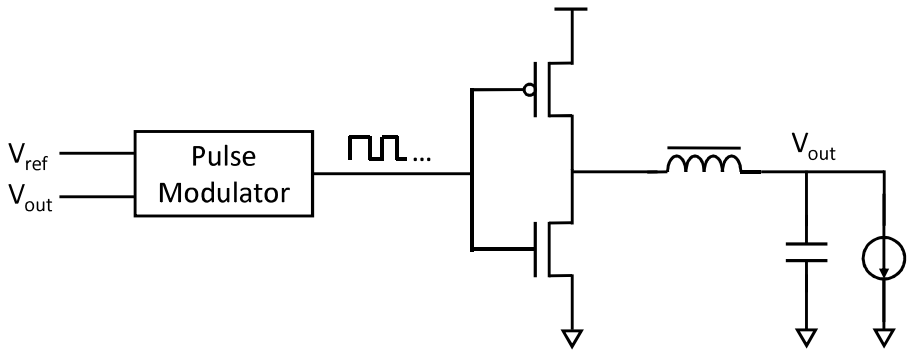

Figure 2.4: Structure of a buck converter

(PWM). As shown in Figure.2.4, the pulse modulator compares the output voltage V_{out} and the reference voltage V_{ref} and modulates the width of its output pulse. The PWM pulse controls the

output voltage with switching transistors and an inductor, and the output voltage is given by

$$V_{out} = D \times V_{DD}$$

where D is the duty cycle of the PWM pulse [23]. The transistor switching introduces voltage ripples in its output. Increasing the pulse frequency to hundreds of MHz and using multi-phase buck converters can reduce the amplitude of ripples[12]. The power loss of a buck converter mainly comes from switching power and resistive power. The switching power is the power dissipated on switching transistors, so it's independent of load current; the resistive power is caused by the intrinsic resistance of the inductance and the output resistance, so the resistant power becomes significant with high load current.

Off-chip buck converters are usually deployed on a motherboard to provide the supply voltage for the on-board chips. With little limitations on area and cost, high-quality capacitors and inductors can be employed to provide high efficiency and low noise voltage conversion. Off-chip buck converters usually operate at KHz to MHz [24, 25], which is usually lower than the frequency of on-chip buck converters. Also, because of the off-chip deployment, its response time to on-chip load current changes is longer.

On the other side, there have been great progress on on-chip buck converters thanks to the development of on-die/in-package inductors and novel magnetic materials [12]. With a switching frequency of tens to hundreds of MHz and increased proximity to the current loads, on-chip buck converters have faster response time compared to off-chip buck converters. Hence they are used in many dynamic voltage scaling (DVS) applications. That been said, integrating inductors with high-Q value is still a challenge. Thus the efficiency of on-chip buck converters is lower than the off-chip ones, but they are still much more efficient than LDOs in terms of voltage conversion.

2.3.2 On-chip Low Dropout Voltage Regulator

The low dropout voltage regulator (LDO) is a type of linear voltage regulator whose main components are a pass transistor and an error amplifier, as shown in Figure.2.5. It's very area-

efficient and easy to be integrated on silicon, which makes distributed deployment of LDOs across a power domain possible. The error amplifier compares the reference voltage V_{ref} and the feedback voltageV_f generated by resistors R_1 and R_2 and controls the gate voltage of the pass transistor. Therefore, the output voltage of the LDO is given by

$$V_{out} = V_{ref} \times \frac{R_1 + R_2}{R_2}$$

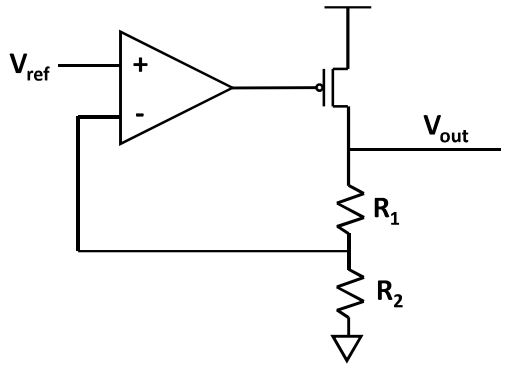

Figure 2.5: Structure of an LDO

Compared with load current I_L, the current flowing through the resistors, R_1 and R_2, and the amplifier is small and can be ignored. Hence the input current $I_{in} \approx I_L$, and the efficiency of the LDO is given by

$$\eta_{LDO} = \frac{P_{out}}{P_{in}} = \frac{I_L \times V_{out}}{I_{in} \times V_{DD}} \approx \frac{V_{out}}{V_{DD}} = \frac{V_{out}}{V_{out} + \Delta V}$$

where ΔV is the dropout voltage.

It's evident that reducing dropout voltage ΔV can improve the efficiency of LDO, but ΔV cannot be infinitely small. Figure.2.6 shows the output voltages of an LDO design [26] with varying input voltage values where the nominal output voltage is $1\,\mathrm{V}$. When ΔV is smaller than

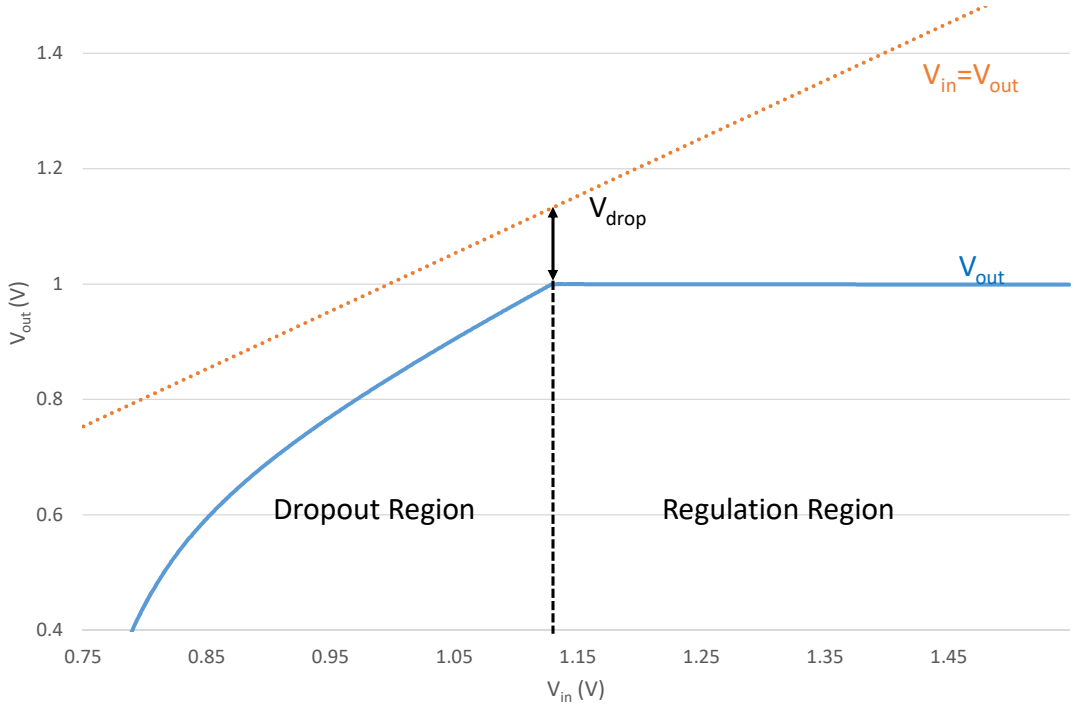

Figure 2.6: The dropout voltage of LDO

the minimum dropout voltage V_{drop}, the LDO ceased to regulate the output voltage and the output voltage falls below $1\,\mathrm{V}$, which may jeopardize the power integrity. The minimum dropout voltage V_{drop} is approximately proportional to the load current I_L for a realistic LDO design [2, 26]. Thus the minimum dropout voltage is given by

$$V_{drop} = \frac{I_L}{I_{L,max}} \Delta V_{drop,max}$$

where $I_{L,max}$ is the maximum load current of the LDO, and $\Delta V_{drop,max}$ is the dropout voltage when the load current is maximum [2].

LDOs can also be used to filter input voltage ripples. As shown in Figure.2.6, when the dropout voltage is higher than V_{drop}, the LDO maintains its output voltage at $1\,\mathrm{V}$ regardless of the input voltage values. In practice, LDOs can be designed with high Power Supply Ripple Rejection (PSRR)[26] to suppress the voltage noise from input.

2.4 Workload-Aware HVR Control Policy

A workload-aware control policy was proposed in [2] to maintain a high system efficiency under different workloads based on the characteristics of the voltage regulators. The system efficiency is determined by the efficiencies of VRs in the PDN, which are related to their input and output voltages and load currents. So the control policy tunes these parameters of VRs to adjust their efficiency. Meanwhile, the control policy also searches for the best efficiency trade-offs between three stages to achieve the best system efficiency when tuning the VRs in the HVR architecture.

The applications of hybrid voltage regulators and the multi-stage structure provide rich tunability in the HVR architecture. The control variables proposed in [2] and their descriptions are shown in the table below.

Variables	Description
$V_{out,off}$	Output voltage of off-chip buck converter cluster
N_{off}	Number of active off-chip buck converters
$V_{out,on}$	Output voltage of on-chip buck converter cluster
N_{on}	Number of active on-chip buck converters

Table 2.2: Control variables and the descriptions. [2]

N_{on} and N_{off} are numbers of active on-chip and off-chip buck converters. With the same workload current, changing N_{on} and N_{off} is equivalent to adjusting the load current of each buck converter, which can be accomplished by gating some of the converters. Two intermediate voltages in the conversion chain, the output voltage of off-chip buck converters, $V_{out,off}$, and the output voltage of on-chip buck converters, $V_{out,on}$, can also be tuned. Note that $V_{out,off}$ approximately equals the input voltage of on-chip buck converters and $V_{out,on}$ approximately equals the input voltage of on-chip LDOs. The reason why they are not exactly equal is because of the existence of parasitic resistance between the output of the VRs in the previous stage and the input of the VRs in the next stage.

Based on the different response time of on-chip and off-chip VRs, the control policy utilizes two cycle periods, T_{on} and T_{off}, to control the on-chip PDN and off-chip PDN separately as shown in Figure.2.7. The T_{on} and T_{off} are set at $1\,\mu s$ and $100\,\mu s$ respectively in [2].

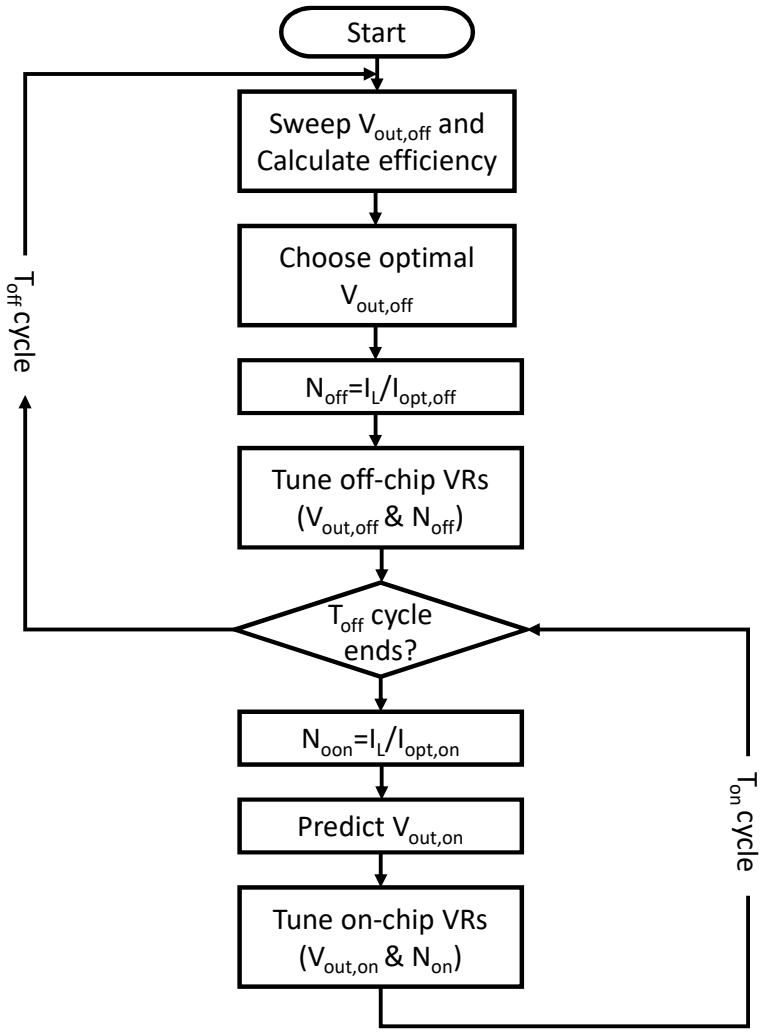

Figure 2.7: The control policy of HVR[2]

During the T_{on} cycle, the control variables for the on-chip converters, $V_{out,on}$ and N_{on}, are tuned. $V_{out,on}$ should be set close to the output voltage of LDOs to reduce the power loss on LDOs, while it should also keep the LDOs in regulation, especially when there is non-uniform workload

current distribution across the chip. Since the spatial workload distribution for a multi-core CPU is difficult to predict, an offline-trained Sparse Relevance Kernel Machine (SRKM) [27, 28] model is trained to learn the spatial workload distribution and predict the $V_{out,on}$ online using an SRKM accelerator to improve the LDO regulation efficiency without jeopardizing the power integrity. The SRKM model takes the readouts of 10 voltage sensors distributed within the power domain and the LDO input voltage in the previous control cycle as input features. These features contain the information on the workload currents and their spatial distribution so that the model can predict the best input voltage for LDOs. With $V_{out,on}$ determined, the number of active on-chip buck converters is given by

$$N_{on} = \frac{I_L}{I_{opt,on}(V_{in,on}, V_{out,on})}$$

where I_L is the total load current measured in the previous T_{on} cycle and $I_{opt,on}(V_{in,on}, V_{out,on})$ is the optimal load current of a on-chip buck converter given its input and output voltages, which is stored in on-chip Lookup Tables (LUTs).

During the T_{off} cycle, the control variables related to off-chip buck converters, $V_{out,off}$ and N_{off}, are tuned. First, $V_{out,off}$ is determined by sweeping its possible values in a given range. In each sweep step, average $V_{out,on}$ in the previous T_{off} cycle is used for estimation and estimated N_{on} is calculated with the same method mentioned above. The estimated value for N_{off} can be calculated using the same approach for N_{on}, which is given by

$$N_{off} = \frac{I_L}{I_{opt,off}(V_{in,off}, V_{out,off})}$$

where I_L is the total load current measure in the previous T_{off} cycle and $I_{opt,off}(V_{in,off}, V_{out,off})$ is the optimal load current given the input and output voltages of a off-chip buck converter, which is also stored in LUTs. Given all the control variables, the system efficiency can be evaluated using the efficiency values of on-chip and off-chip buck converters stored in LUTs and indexed by input and output voltages and load currents. After sweeping all the candidate $V_{out,off}$ values, the one with the maximum estimated system efficiency value is selected. The N_{off} is calculated again as a

system control variable, given the actual input and output voltages of the off-chip buck converters.

The HVR architecture and the workload-aware control policy provide high-quality and high-efficiency power delivery, which reduces system energy by up to 23.9% compared with static 2-stage PDN while ensuring the same power integrity specification[2].

3. VARIATION-AWARE HETEROGENEROUS VOLTAGE REGULATION *

3.1 Overview

As described in the previous chapter, the HVR control policy requires efficiency characteristics of buck converters to estimate the system efficiency and determine the system control variables. Besides, the SRKM module for $V_{out,on}$ prediction is trained with the data generated from the PDN without any temperature and process variation. Therefore, when inevitable variations exist in the system, the efficiency characteristics and the SRKM module predictions become inaccurate, which may compromise the power integrity and degrade the system efficiency. In this chapter, a variation-aware adaption of HVR is proposed based on the control policy proposed in [2]. Figure.3.1 shows the proposed variation-aware control policy. By exploiting mechanisms of variation tracking, the variation-aware control policy takes the chip's variations into account when generating control variables for the system. Finally, the performance of the proposed variation-aware adaption of HVR is evaluated and analyzed. The effects of temperature and process variations on the original HVR control policy are also analyzed.

3.2 Tracking Variations in PDN

This section explains the tracking circuits shown in Figure.3.1 and how they are used in the variation-aware control policy.

First, the off-chip buck converters are assumed to be free from variations because it's possible to employ high-accuracy inductors and capacitors on PCB to reduce variations.

The on-chip buck converters are assumed to be auto-tuned so that their regulation performances are not affected by temperature and process variations. A measurement circuit is proposed to measure the efficiencies of buck converters with variations online at $-40\,^{\circ}\mathrm{C}$ and $125\,^{\circ}\mathrm{C}$. The efficiency measurements are stored in on-chip LUTs during the setup stage, and the efficiency

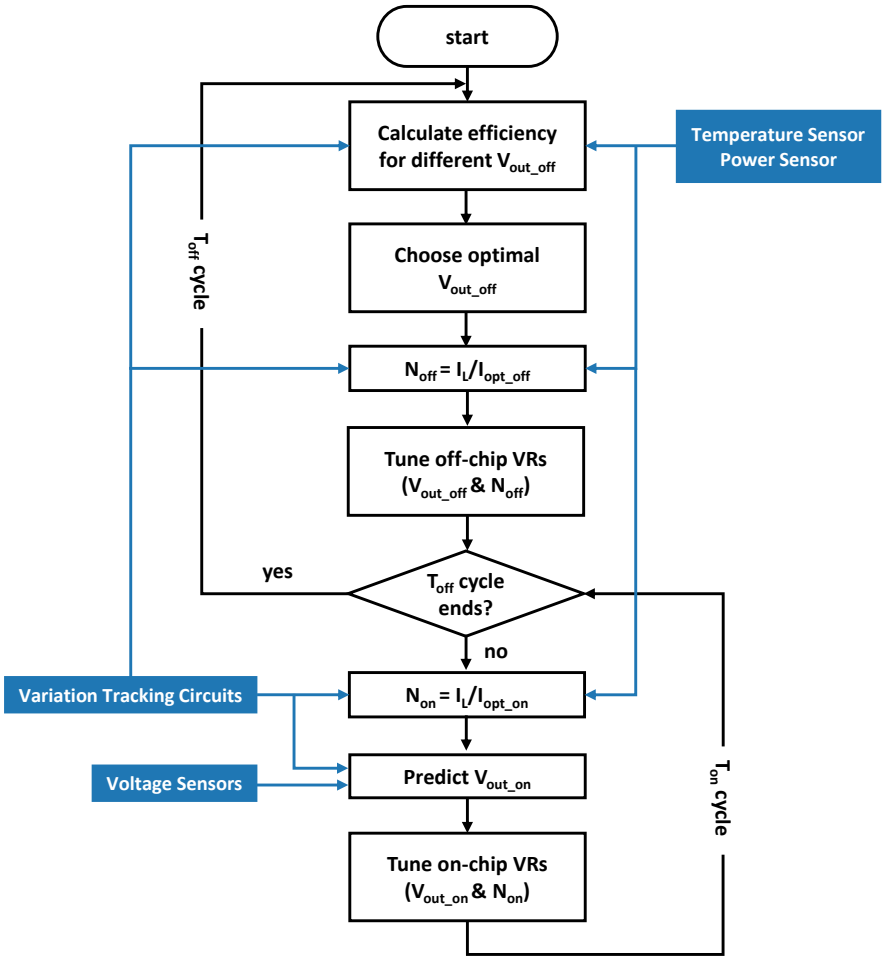

Figure 3.1: Variation-aware control policy.
Reprinted, with permission, from [3] © 2020 IEEE

values used by the variation-aware control policy during runtime are interpolated from these stored values with the temperature measured by integrated temperature sensors.

A variation tracking circuit for LDOs is proposed to measure two variation-related parameters: average loop gain and bandwidth of LDOs. These parameters are used as additional features of a variation-aware SRKM model to predict the LDOs' input voltage based on both the workload and the variations of the LDOs. The variation-aware SRKM model is trained offline using 3000 training samples collected from different variation cases. It is loaded to the on-chip SRKM accelerator during the setup stage, and make predictions during runtime.

3.2.1 Tracking Variations of On-chip Buck Converters

The regulation performance of a buck converter is determined by it's loop dynamics. [29] proposed a technique for adapting buck converter's loop dynamics to process variations using low-cost digital circuits. It ensures fixed values for loop crossover frequency and its phase margin regardless of load current or process variations. In this thesis, we assume that this auto-tuning technique is adopted so that the regulation performance of on-chip buck converters is not affected by process variations.

However, their efficiencies will still be affected by variations. The efficiency of an on-chip buck converter is related to the switching transistors and the inductor which are vulnerable to the process and temperature variations. In this thesis, the efficiency values of on-chip buck converters are measured using a efficiency measurement circuit.

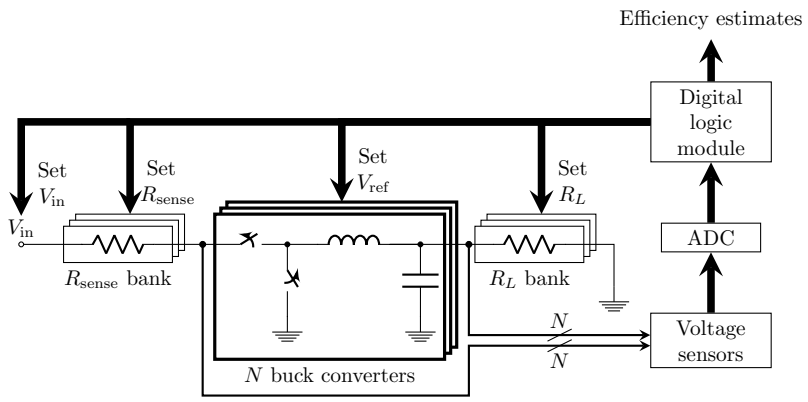

Figure 3.2: Proposed variation tracking circuit for on-chip buck converters.
Reprinted, with permission, from [3] © 2020 IEEE

Figure.3.2 shows the proposed measurement circuit for on-chip buck converters. The control module sweeps the load current and input and output voltages of a on-chip buck converter by selecting R_{sense} and R_L and setting V_{in} and V_{ref}. The voltage drop across R_{sense} and R_L can be measured by voltage sensor and digitized using an analog-to-digital converter (ADC). Then the efficiency can be calculated using corresponding resistance values of R_{sense} and R_L and voltage

measurements. Resistors, voltage sensors and the required digital logics can be easily integrated on the die. Since ADCs are essential components of large SoCs, so they can be reused to prevent additional hardware costs.

3.2.2 Tracking Variations of on-chip LDOs

Control loop gain and loop bandwidth are two metrics directly related to the regulation performance of the LDOs. The loop gain determines the accuracy of the output voltage while the loop bandwidth decides the response time. Since the effects of process and temperature variations on LDO performance can be reflected in the changes of these two metrics, a circuit used to measure the loop gain and bandwidth of LDOs online is proposed to track the variations in LDOs.

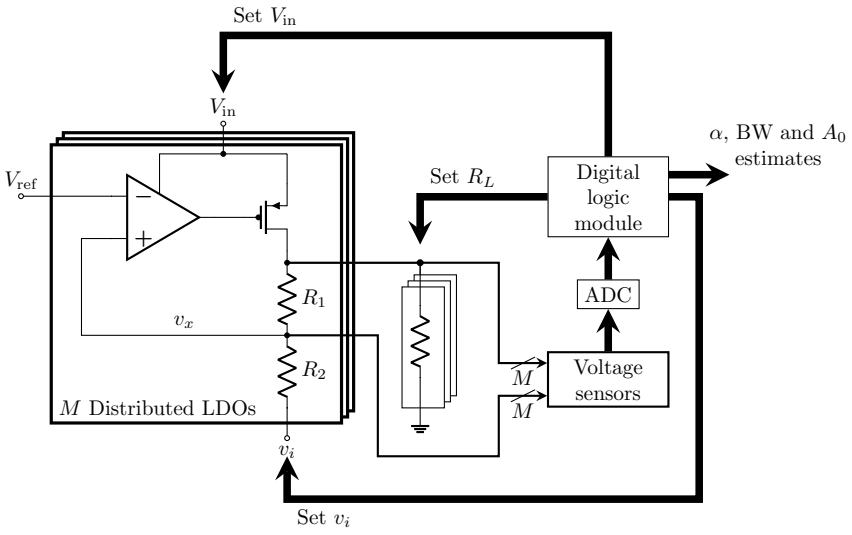

Figure 3.3: Proposed variation tracking circuit for LDOs.
Reprinted, with permission, from [3] © 2020 IEEE

Figure.3.3 shows the proposed circuit to measure LDO's loop gain and bandwidth. The loop gain can be estimated by sweeping v_i and measuring v_x, which gain can be given by

$$L(s) = \frac{R_1}{R_1 + R_2} \cdot \frac{v_i(s)}{v_x(s)} - 1$$

Using a sinusoidal test signal with varying frequency $v_i(j\omega)$, the value of $L(j\omega)$ can be measured at different frequencies and the loop bandwidth and low frequency gain can be estimated.

3.2.3 Temperature Variations

Apart from the process variation, the temperature change is also important to be tracked because it also affects the characteristics of the on-chip devices. In this thesis, a temperature sensor is employed on chip to provide the temperature information.

For buck converters, their efficiency values at a certain temperature is interpolated using efficiency values measured at two temperature extremes ($-40\,^{\circ}\mathrm{C}$ and $125\,^{\circ}\mathrm{C}$). Since the efficiency change caused by temperature is quite linear, the two-point measurement and interpolation are sufficient for efficiency estimations.

As for LDOs, temperature variations can also change their loop gain and bandwidth, so the proposed variation tracking circuit can capture the temperature change of the silicon. The temporal temperature change can be captured by measuring the parameters periodically. It can be accomplished using replica LDOs and will not cause significant overheads because the temporal temperature is relatively slow and the measurement can be repeated in a low frequency.

3.3 Experiment Setups

In this section, the details about experiment setups are introduced. The experiments in this thesis include circuit simulations for power integrity assessments and system efficiency evaluations.

In circuit simulations, a circuit of a PDN containing a RC power grid model, LDOs, behavior models of on-chip buck converters, and package/C4 connections is simulated to observe the voltage noise in the power grid.

The RC power grid model described in Chapter.2 is adopted, and current traces in the power grid are extracted using architectural simulators. In this thesis, we use the full-system multi-core simulator, GEM5[21], to generate the runtime statistics of a processor. Then they are fed to McPAT[22], an integrated power, area and timing modeling framework for multi-core processors, to generate the runtime power traces for all components of the processor. Finally, the current

traces, which are the workloads in our PDN, can be given by $I = \frac{P}{V_{DD}}$, where I is the dynamic current trace, P is the dynamic power trace generated by the GEM5 and McPAT and V_{DD} is the supply voltage which is assumed to be $1\,\mathrm{V}$. We run the PARSEC[30] benchmark suite for a $45\,\mathrm{nm}$ four-core processor with configurations shown in the table below. Eventually, the current traces for 11 different components in each core are generated with a granularity of $100\,\mathrm{ns}$.

# of cores	4
Frequency	1.8GHz
Vdd	1V
Imax (per core)	25A
Core area (per core)	40.4mm^2

Table 3.1: Configurations of simulated processor

Besides, a transistor-level implementation of LDO proposed in [26] with a maximum load current of $100\,\mathrm{mA}$ and the same PCB and package connection model proposed in [20], which is derived from Intel Pentium 4 processors, are used in circuit simulations.

Because behaviors and efficiency characteristics of on-chip and off-chip buck converters are more complicated than LDOs, we adopt different models for different evaluation purposes: circuit simulations and system efficiency evaluations. A behavior model of on-chip buck converters with ideal switches and passive devices is adopted to accelerate circuit simulations. The off-chip buck converters are modeled as ideal voltage sources in circuit simulations because they are assumed to be unaffected by variations and the voltage noise introduced by them will be further filtered by on-chip buck converters and LDOs. On the other side, PowerSoC [31] is used to find the best design parameters for on-chip and off-chip buck converters, such as switching frequency and filter inductance, and generate analytical efficiency models for system efficiency evaluations.

In addition to nominal models, VRs with process and temperature variations are also modeled in both circuit simulations and efficiency evaluations to assess the performance of different control policies under temperature and process variations.

In system efficiency evaluations, on-chip buck converters with process and temperature variations in their transistors and passive devices are simulated using a standard $90\,\mathrm{nm}$ technology and their efficiencies are measured. The passive devices are assumed to have a standard deviation of 20% because their characteristics tend to have a large variance [32]. The Monte Carlo simulation is used to sample the random process variations and generate multiple instances of on-chip buck converters with variations for efficiency measurement. However, variations in on-chip buck converters are not modeled in circuit simulations since the auto-tuning technique is expected to adjust the converters and ensure their regulation performance as described in Section.3.2.1

For the LDOs, the Monte Carlo sampling is used to generate multiple LDO instances with variations using a standard $90\,\mathrm{nm}$ technology. These LDO instances are used in circuit simulations to assess the power integrity of different control policies. Besides, the simulations of PDNs consisting of LDOs with process and temperature variations and the measurements of their loop gain and bandwidth are used to build a dataset and train the variation-aware SRKM model mentioned in the previous section. 3000 training samples are collected from simulations with 200 different process variation cases at temperatures of $27\,^{\circ}\mathrm{C}$, $-40\,^{\circ}\mathrm{C}$, and $125\,^{\circ}\mathrm{C}$. The variation-aware SRKM model achieved a normalized mean square error (NMSE) of 0.0343.

3.4 Performance Evaluations and Discussions

In this section, we compare the efficiency and power integrity of the original HVR control policy and the proposed variation-aware control policy in systems with process and temperature variations.

For a fair comparison of performance of both control policies, we first adjust both policies so that they provide the same level of power integrity. Since designing a PDN based on worst-case scenarios to guarantee the power integrity would be costly and degrade the system efficiency significantly, rare occurrences of VEs are allowed. Fail-safe mechanisms like rolling-back recovery [33] and adaptive frequency tuning [34] are assumed to equip the target hardware and protect it from these rare VEs. A guardband voltage was added to the prediction of the SRKM model to compensate for prediction errors. It is adjusted for both nominal and variation-aware control

policies to provide the same level of power integrity measured by the frequency of VE occurrences.

3.4.1 Efficiency evaluation

With the same level of power integrity, we can evaluate and compare the efficiency of both control policies. The results of efficiency evaluations are shown in Figure.3.4.

Four different optimization schemes are evaluated with nine workloads extracted form nine benchmark programs in PARSEC benchmark suite [30]. The original HVR control policy is denoted by `Nominal` in Figure.3.4. Control policies with variation-aware LUTs for on-chip buck converters or variation-aware SRKM model are denoted by `LUT_OPT` and `ML_OPT`, respectively. Finally, the proposed variation-aware HVR control policy, which utilizes both optimization mechanisms, is denoted by `LUT_ML_OPT`. The efficiency of each scheme is evaluated under 4 different process cases and 3 different temperatures, $-40\,^{\circ}\mathrm{C}$, $27\,^{\circ}\mathrm{C}$, and $125\,^{\circ}\mathrm{C}$. The average system efficiency under these 12 different cases are shown by the bars, and the maximum and the minimum efficiencies are shown by the error bars in Figure.3.4.

As shown in Figure.3.4, the variation aware control policy outperforms other schemes in every benchmark. The maximum efficiency improvement is 4.28%, and the average efficiency improvement is 1.13%.

The efficiencies of the original HVR control policy in a system without any variations, denoted by `Nominal System`, are also measured. Comparing the efficiency values of `Nominal` and `Nominal System`, it's easy to find out that the process and temperature variations can increase or decrease system efficiency, but the average effect is efficiency degradation. Therefore, the significance of the variation-aware HVR control policy is to compensate for the negative effects of variations on the system by tracking the variations and adjust the system accordingly. It can also be seen that the disturbance of variations on the system efficiency is not significant. The maximum efficiency deviation from the nominal system is 3.28%.

To sum up, the variation-aware control policy improves the system efficiency under process and temperature variations. However, the efficiency improvement is not substantial. Meanwhile, the efficiency disturbance of variations on the original HVR control policy is also small. These

results show that the original HVR control policy possesses certain robustness against process and temperature variations, which will be explained in the next two subsections.

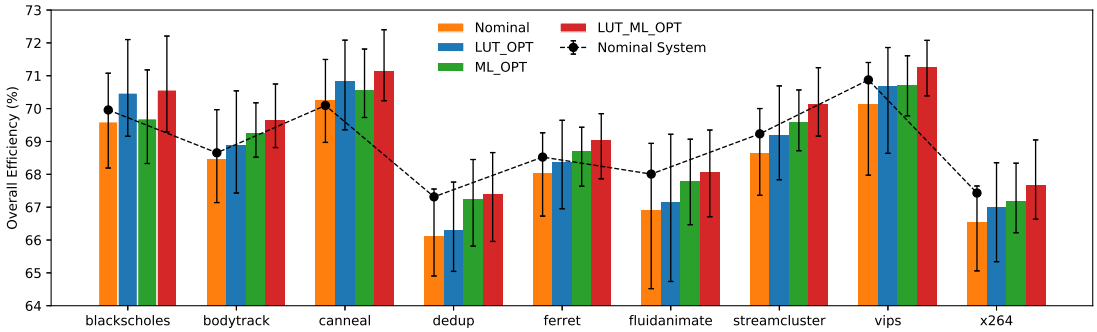

Figure 3.4: Efficiency evaluation of improved schemes for different benchmarks. Reprinted, with permission, from [3] © 2020 IEEE

3.4.2 Effect of variations on off-chip buck converter control

First, we compare control signals generated for off-chip buck converters, $V_{out,off}$ and N_{off}. As described in Chapter.2, $V_{out,off}$ is determined by sweeping possible values of $V_{out,off}$ and select the one with highest estimated system efficiency. Then N_{off} is decided accordingly. So we analyze the searching process of $V_{out,off}$ to assess the effects of variations on off-chip buck converter control.

The curves in Figure.3.5 show six searching processes under different system configurations. The solid red curve shows the searching process of the original HVR control policy in a nominal system. The curve shows a peak efficiency value at $1.3\,\mathrm{V}$, so $1.3\,\mathrm{V}$ is selected by the control policy. The dashed curves show the proposed variation-aware control policy's searching processes in five systems with different variations. All curves shown in the figure have the same workload current, and the only difference is the system configurations and control policies. The original HVR control policy will select the same value for $V_{out,off}$ regardless of system variations because it only relies on the workload. In contrast, the variation-aware control policy will select different $V_{out,off}$ values based on system variations, as shown by the dashed curves.

For 3 out of 5 variation cases, the original HVR control policy selects the optimal $V_{out,off}$ value that the variation-aware control policy selects, which is $1.3\,\mathrm{V}$. For the rest two variation cases, the original HVR control policy failed to select the best $V_{out,off}$, but the variations only shift the optimal value by $200\,\mathrm{mV}$.

Meanwhile, Figure.3.5 also shows the estimated system efficiency. In the cases when the original control policy failed to select the optimal operation point, the degradation of the estimated efficiency is minimal and less than 1%. Even though this estimation is not accurate considering the rough estimates of control variables for on-chip buck converters, it can still shed some light on how little system variations affect system efficiency. In conclusion, the original control policy manages to select the best $V_{out,off}$ or a value that is close to the optimal value despite the existence of variations. Since the nominal control policy performs sub-optimally in some cases, it can be expected that adding variation awareness can improve system efficiency.

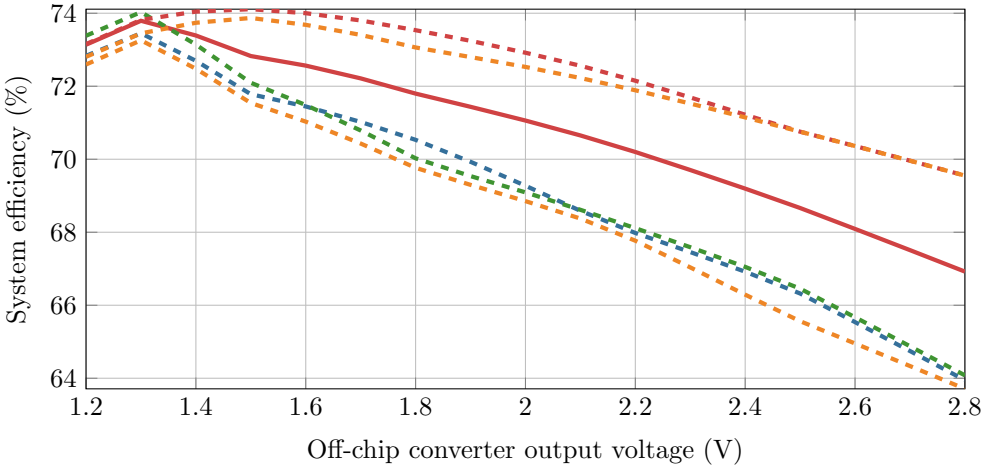

Figure 3.5: Estimated efficiency as a function of $V_{out,off}$.
Reprinted, with permission, from [3] © 2020 IEEE

3.4.3 Effect of variations on on-chip buck converter control

Then we consider the other two control variables, $V_{out,on}$ and N_{on}. $V_{out,on}$ is determined by an SRKM model, and N_{on} is decided based on input and output voltages and load current of on-chip

buck converters. So we analyze SRKM models to estimate the effects of variations on on-chip buck converter control.

SRKM models are responsible for $V_{out,on}$ prediction, which should maintain power integrity and reduce LDO power loss at the same time. The prediction of SRKM models, $V_{out,on}$, is related to LDO input voltage, and the input features of SRKM models, voltage sensor readouts, are related to LDO output voltage. Hence we observe the waveform of input an output voltages of LDOs to analyze the performance of SRKM models. The waveforms generated by the nominal and variation-aware SRKM models are shown in Figure.3.6.

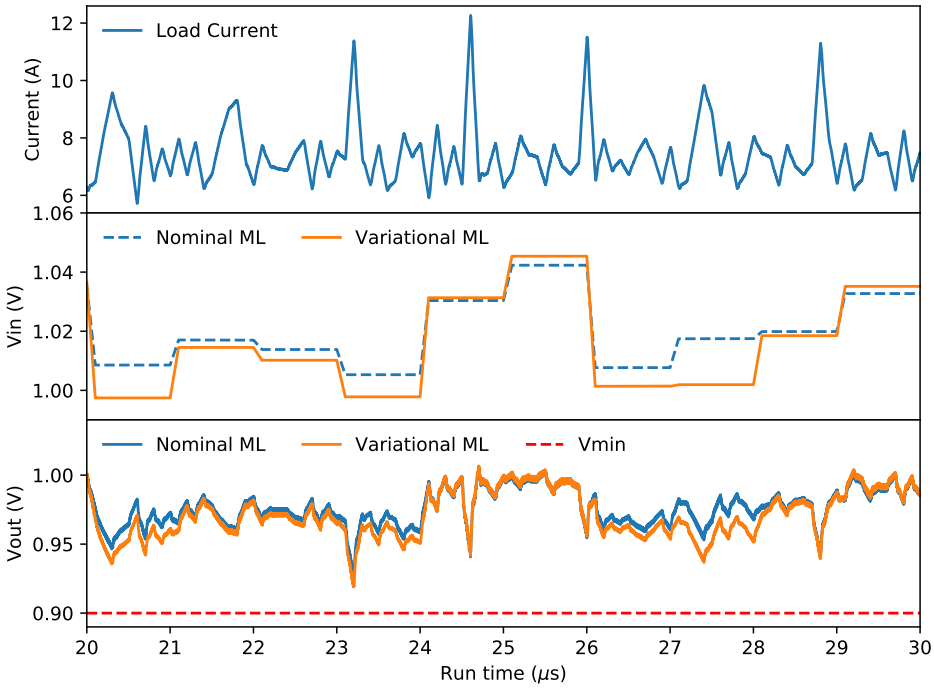

Figure 3.6: Predicitons in variation-aware and nominal SRKM models.
Reprinted, with permission, from [3] © 2020 IEEE

The waveforms are simulated with a segment of the Blackscholes benchmark. The guardband of both nominal and variation-aware SRKM models have been adjusted to provide the same level

of power integrity.

As shown in the figure, At $23\,\mu s$, a current surge causes output voltage drops in both waveforms, and both SRKM models raised the input voltage for LDO to prevent potential VEs as expected. Another current surge between $24\,\mu s$ and $25\,\mu s$ also triggers the SRKM models to raise LDO input voltage between $25\,\mu s$ and $26\,\mu s$. It indicates that both nominal and variation-aware control policies can function correctly in the presence of variation; i.e., they can react to voltage drops in voltage sensor readouts and raise LDO input voltage to prevent potential occurrences of VEs accordingly. However, the voltage increment applied might not be optimal due to variations. As shown in the figure, the variation-aware SRKM model could predict lower voltages without jeopardizing power integrity to reduce LDO power loss.

In conclusion, the nominal SRKM model can function appropriately in systems with variations, but they might not be as efficient as the variation-aware SRKM model because of the inaccurate predictions. Both on-chip and off-chip buck converter controls in the original HVR control policy exhibits a certain degree of robustness in systems with variations. Thus the efficiency disturbance caused by variation is small, and the efficiency improvement provided by the proposed variation-aware control policy is limited.

3.5 Conclusions and Future works

A variation-aware adaption of HVR architecture is proposed to adapt the HVR to systems with process and temperature variations. Two variation tracking circuits were proposed to enable the tracking of process and temperature variation effects on VRs in HVR architecture with minimal hardware overhead. Using the variation tracking circuit, a variation-aware control policy is implemented to improve HVR performance in systems with variations, resulting in an efficiency increase of up to 4.28% and 1.13% on average with the same level of power integrity. A detailed study of the effects of process and temperature variations on the original HVR architecture was presented, and the HVR architecture was found to be robust against the process and temperature variations.

However, one drawback for this project is that the stability of the control policy is not guaranteed. The proposed control policy uses the workload information gathered in the previous control

cycle to determine the control variables in the next control cycle, forming a feedback loop with delay. Even though there were no issues during the evaluation and simulations, the stability of this feedback control loop is not guaranteed. Thus more researches should be conducted with regard to the stability of the control policy.

REFERENCES

[1] C. A. Mack, "Fifty years of moore's law," *IEEE Transactions on Semiconductor Manufacturing*, vol. 24, no. 2, pp. 202–207, 2011.

[2] X. Zhan, J. Chen, E. Sánchez-Sinencio, and P. Li, "Power Management for Multicore Processors via Heterogeneous Voltage Regulation and Machine Learning Enabled Adaptation," *TVLSI*, pp. 1–14, 2019.

[3] J. Riad, J. Chen, E. Sánchez-Sinencio, and P. Li, "Variation-aware heterogeneous voltage regulation for multi-core systems-on-a-chip with on-chip machine learning," in *2020 21st International Symposium on Quality Electronic Design (ISQED)*, pp. 190–194, 2020.

[4] R. Zhang, *Pre-RTL On-Chip Power Delivery Modeling and Analysis*. PhD thesis, University of Virginia, 2015.

[5] M. Saint-Laurent and M. Swaminathan, "Impact of power-supply noise on timing in high-frequency microprocessors," in *2002 IEEE 11th Topical Meeting on Electrical Performance of Electronic Packaging*, pp. 261–264, 2002.

[6] S. K. Khatamifard, L. Wang, W. Yu, S. Köse, and U. R. Karpuzcu, "Thermogater: Thermally-aware on-chip voltage regulation," in *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, pp. 120–132, 2017.

[7] J. F. Bulzacchelli *et al.*, "Dual-Loop System of Distributed Microregulators with High DC Accuracy, Load Response Time Below 500 ps, and 85-mV Dropout Voltage," *JSSC*, vol. 47, no. 4, pp. 863–874, 2012.

[8] H. Krishnamurthy *et al.*, "A Digitally Controlled Fully Integrated Voltage Regulator with On-Die Solenoid Inductor with Planar Magnetic Core in 14nm Tri-Gate CMOS," in *ISSCC*, pp. 336–337, IEEE, 2017.

[9] C. Huang and P. K. Mok, "An 84.7% efficiency 100-mhz package bondwire-based fully integrated buck converter with precise dcm operation and enhanced light-load efficiency," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 11, pp. 2595–2607, 2013.

[10] W. Kim, D. M. Brooks, and G.-Y. 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*, pp. 268–270, IEEE, 2011.

[11] G. Sizikov, A. Kolodny, E. G. Fridman, and M. Zelikson, "Efficiency optimization of integrated dc-dc buck converters," in *Electronics, Circuits, and Systems (ICECS), 2010 17th IEEE International Conference on*, pp. 1208–1211, IEEE, 2010.

[12] E. A. Burton *et al.*, "FIVR- Fully Integrated Voltage Regulators on 4th Generation Intel® Core SoCs," in *APEC*, pp. 432–439, IEEE, 2014.

[13] V. Zyuban *et al.*, "IBM POWER8 Circuit Design and Energy Optimization," *IBM Journal of Research and Development*, vol. 59, no. 1, pp. 9–1, 2015.

[14] S. Kose and E. G. Friedman, "Distributed on-chip power delivery," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, no. 4, pp. 704–713, 2012.

[15] D. Pathak, H. Homayoun, and I. Savidis, "Smart grid on chip: Work load-balanced on-chip power delivery," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 9, pp. 2538–2551, 2017.

[16] H. Li, J. Xu, Z. Wang, P. Yang, R. K. V. Maeda, and Z. Tian, "Adaptive power delivery system management for many-core processors with on/off-chip voltage regulators," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*, pp. 1265–1268, 2017.

[17] W. Lee, Y. Wang, and M. Pedram, "Optimizing a reconfigurable power distribution network in a multicore platform," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 7, pp. 1110–1123, 2015.

[18] A. A. Sinkar, H. R. Ghasemi, M. J. Schulte, U. R. Karpuzcu, and N. S. Kim, "Low-cost per-core voltage domain support for power-constrained high-performance processors," *IEEE*

Transactions on Very Large Scale Integration (VLSI) Systems, vol. 22, no. 4, pp. 747–758, 2014.

[19] R. Courtland, *Intel Now Packs 100 Million Transistors in Each Square Millimeter*, 2017. https://spectrum.ieee.org/nanoclast/semiconductors/processors/intel-now-packs-100-million-transistors-in-each-square-millimeter.

[20] M. S. Gupta, J. L. Oatley, R. Joseph, G. Wei, and D. M. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *2007 Design, Automation Test in Europe Conference Exhibition*, pp. 1–6, 2007.

[21] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, *et al.*, "The gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol. 39, no. 2, pp. 1–7, 2011.

[22] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "Mcpat: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pp. 469–480, ACM, 2009.

[23] Wikipedia contributors, "Buck converter — Wikipedia, the free encyclopedia," 2020. [Online; accessed 22-June-2020].

[24] Y. L. L. Cheng *et al.*, "A 10/30MHz Wide-Duty-Cycle-Range Buck Converter with DDA-Based Type-III Compensator and Fast Reference-Tracking Responses for DVS Applications," in *ISSCC*, pp. 84–85, IEEE, 2014.

[25] P. Y. Wu, S. Y. Tsui, and P. K. Mok, "Area-and power-efficient monolithic buck converters with pseudo-type iii compensation," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 8, pp. 1446–1455, 2010.

[26] S. Lai and P. Li, "A fully on-chip area-efficient cmos low-dropout regulator with fast load regulation," *Analog Integrated Circuits and Signal Processing*, vol. 72, pp. 433–450, Jul. 2012.

[27] H. Lin and P. Li, "Relevance Vector and Feature Machine for Statistical Analog Circuit Characterization and Built-In Self-Test Optimization," in *DAC*, p. 11, ACM, 2016.

[28] H. Lin, A. M. Khan, and P. Li, "Statistical circuit performance dependency analysis via sparse relevance kernel machine," in *IC Design and Technology (ICICDT), 2017 IEEE International Conference on*, pp. 1–4, IEEE, 2017.

[29] J. Morroni, R. Zane, and D. Maksimovic, "Adaptive tuning of digitally controlled switched mode power supplies based on desired phase margin," in *2008 IEEE Power Electronics Specialists Conference*, pp. 1250–1256, 2008.

[30] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The parsec benchmark suite: Characterization and architectural implications," in *2008 International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pp. 72–81, 2008.

[31] X. Wang, J. Xu, Z. Wang, K. J. Chen, X. Wu, Z. Wang, P. Yang, and L. H. K. Duong, "An analytical study of power delivery systems for many-core processors using on-chip and off-chip voltage regulators," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 9, pp. 1401–1414, 2015.

[32] P. Napolitano, M. Poncik, and A. Bonelli, "Effect of Parasitics and Process Variations in Buck Converters Operated in Voltage-Mode," in *2018 29th Irish Signals and Systems Conference (ISSC)*, pp. 1–6, June 2018.

[33] H. Akkary, R. Rajwar, and S. T. Srinivasan, "Checkpoint processing and recovery: towards scalable large instruction window processors," in *Proceedings. 36th Annual IEEE/ACM International Symposium on Microarchitecture, 2003. MICRO-36.*, pp. 423–434, 2003.

[34] K. A. Bowman, S. Raina, J. T. Bridges, D. J. Yingling, H. H. Nguyen, B. R. Appel, Y. N. Kolla, J. Jeong, F. I. Atallah, and D. W. Hansquine, "A 16 nm all-digital auto-calibrating

adaptive clock distribution for supply voltage droop tolerance across a wide operating range," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 1, pp. 8–17, 2016.