# DISTRIBUTED IC POWER DELIVERY: STABILITY-CONSTRAINED DESIGN OPTIMIZATION AND WORKLOAD-AWARE POWER MANAGEMENT

A Dissertation

by

XIN ZHAN

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

DOCTOR OF PHILOSOPHY

| | |
|---|---|
| Chair of Committee, | Peng Li |
| Committee Members, | Edgar Sánchez-Sinencio |
| | Weiping Shi |
| | Duncan M. Walker |
| Head of Department, | Scott L. Miller |

May 2019

Major Subject: Computer Engineering

# ABSTRACT

Power delivery presents key design challenges in today's systems ranging from high performance micro-processors to mobile systems-on-a-chips (SoCs). A robust power delivery system is essential to ensure reliable operation of on-die devices. Nowadays it has become an important design trend to place multiple voltage regulators on-chip in a distributive manner to cope with power supply noise. However, stability concern arises because of the complex interactions between multiple voltage regulators and bulky network of the surrounding passive parasitics. The recently developed hybrid stability theorem (HST) is promising to deal with the stability of such system by efficiently capturing the effects of all interactions, however, large overdesign and hence severe performance degradation are caused by the intrinsic conservativeness of the underlying HST framework. To address such challenge, this dissertation first extends the HST by proposing a frequency-dependent system partitioning technique to substantially reduce the pessimism in stability evaluation. By systematically exploring the theoretical foundation of the HST framework, we recognize all the critical constraints under which the partitioning technique can be performed rigorously to remove conservativeness while maintaining key theoretical properties of the partitioned subsystems. Based on that, we develop an efficient stability-ensuring automatic design flow for large power delivery systems with distributed on-chip regulation. In use of the proposed approach, we further discover new design insights for circuit designers such as how regulator topology, on-chip decoupling capacitance, and the number of integrated voltage regulators can be optimized for improved system tradeoffs between stability and performances.

Besides stability, power efficiency must be improved in every possible way while maintaining high power quality. It can be argued that the ultimate power integrity and efficiency may be best achieved via a heterogeneous chain of voltage processing starting from on-board switching voltage regulators (VRs), to on-chip switching VRs, and finally to networks of distributed on-chip linear VRs. As such, we propose a heterogeneous voltage regulation (HVR) architecture encompassing regulators with complimentary characteristics in response time, size, and efficiency. By exploring

the rich heterogeneity and tunability in HVR, we develop systematic workload-aware control policies to adapt heterogeneous VRs with respect to workload change at multiple temporal scales to significantly improve system power efficiency while providing a guarantee for power integrity. The proposed techniques are further supported by hardware-accelerated machine learning prediction of non-uniform spatial workload distributions for more accurate HVR adaptation at fine time granularity. Our evaluations based on the PARSEC benchmark suite show that the proposed adaptive 3-stage HVR reduces the total system energy dissipation by up to 23.9% and 15.7% on average compared with the conventional static two-stage voltage regulation using off- and on-chip switching VRs. Compared with the 3-stage static HVR, our runtime control reduces system energy by up to 17.9% and 12.2% on average. Furthermore, the proposed machine learning prediction offers up to 4.1% reduction of system energy.

# DEDICATION

To my parents, my brother, and my wife for their great supports.

# ACKNOWLEDGMENTS

First of all, I would like to express my greatest gratitude to my Ph.D. advisor Prof. Peng Li. Over the past four and half years, Prof. Li has been constantly supporting me throughout my Ph.D. study at Texas A&M University. His comprehensive knowledge and expertise in related areas, insightful guidance and advises are the essentials for me to accomplish the work presented in this dissertation. Prof. Li is also a spiritual mentor for me who always encourages me to go beyond my comfort zone, break my own personal limitations, and pursue a higher goal which I have never thought about before. Such experience not only shapes me to become a better researcher during my Ph.D. study, but will also influence me profoundly to overcome life barriers in the future.

Besides, I would like to thank Prof. Edgar Sánchez-Sinencio, who provides a lot of valuable advise based on his expertise in the area of analog circuit design which is an essential to my research work. I also thank Prof. Weiping Shi, Prof. Alexander G. Parlos, and Prof. Duncan M. Walker a lot for their willingness to serve on my committee and offer insightful suggestions and comments for my work.

It has been a great experience to work with all my colleagues in the Department of Electrical and Computer Engineering: Dr. Honghuang Lin, Dr. Qian Wang, Dr. Ya Wang, Dr. Yingyezhe Jin, Yu Liu, Wenrui Zhang, Hanbin Hu and Joseph Riad. We had a lot of interesting and useful discussions which inspire me a lot in both research and life. I would also thank my friends Xin Chang, Dadian Zhou, Juning Jiang, Minxiang Zeng, Lichi Deng who made my Ph.D. life more enjoyable and colorful.

Finally, I own a lot to my parents Jianlin Zhan and Li Sun, and my brother Jie Zhan for their unselfish support and encouragement. My special thank goes to my wife Yan Yu. She always gives me unlimited love and understanding. I will never forget her sacrifice of accompanying me to United States for my pursuit of Ph.D. degree. Without her, it would be much more difficult for me to finish my study.

CONTRIBUTORS AND FUNDING SOURCES

## Contributors

This work was supported by a dissertation committee consisting of Professor Peng Li, Professor Edgar Sánchez-Sinencio, Professor Weiping Shi of the Department of Electrical and Computer Engineering and Professor Duncan M. Walker of the Department of Computer Science and Engineering.

All work conducted for the dissertation was completed by the student independently.

## Funding Sources

This dissertation is based upon work supported by the National Science Foundation (NSF) under Grant No. ECCS-1405774 and No. ECCS-1810125 and the Qatar National Research Fund (a member of Qatar Foundation) under NPRP grant # NPRP 8-274-2-107.

Any opinions, findings, conclusions or recommendations expressed in this dissertation are those of the author and do not necessarily reflect the views of NSF and Qatar National Research Fund and their contractors.

TABLE OF CONTENTS

Page

LIST OF FIGURES

FIGURE Page

# LIST OF TABLES

TABLE Page

# 1. INTRODUCTION*

## 1.1 Preliminaries

### 1.1.1 Power Delivery Network (PDN)

Very Large-Scale-Integration (VLSI) power delivery networks (PDNs) are critical IC subsystems for distributing power from external supply source to on-die devices. A robust power delivery system is essential to ensure reliable operation of circuits on a chip. The main functions of IC PDNs include [4, 5]:

- Converting a high external supply voltage to a required on-chip voltage level for each power domain with high energy efficiency.
- Providing fast line and load regulation during abrupt input voltage or output current changes to satisfy supply noise constraint.
- Supporting advanced power management techniques such as dynamic voltage frequency scaling (DVFS).

A typical PDN consists of off-chip voltage regulators (VRs), parasitics of printed circuit board (PCB) and package, integrated VRs, on-die power grids and fast switching current loads. Besides, there are both off-chip decoupling capacitors (decaps) and on-chip decaps in the power grids to help mitigate voltage fluctuations caused by fast switching activities of load circuits. The diagram of a complete PDN is illustrated in Fig. 1.1.

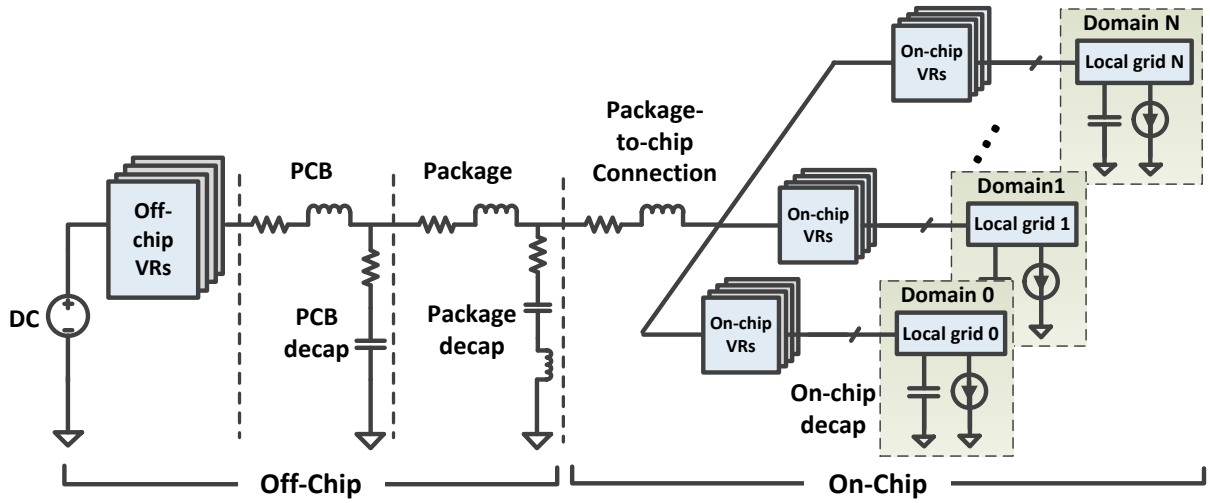


Figure 1.1: A typical structure of IC power delivery network (PDN).

### 1.1.2 Voltage Regulators (VRs)

Voltage regulators (VRs) are key components of a power delivery network with the main function of providing stable and correct voltages in response of dynamics in both input voltage and load currents. The characteristics of VRs have critical impact on power efficiency and regulation response of the entire PDN system. Buck switching VR and linear VR such as low dropout voltage regulator (LDO) are commonly used for voltage conversion and regulation, respectively. To sufficiently supply power, a higher voltage from external supply is usually stepped down by off-chip buck VRs switching at a rate of hundreds of KHz to tens of MHz [6, 7, 8]. They can achieve excellent power efficiency over a wide output range at the expense of bulky and costly off-chip passive components. However, off-chip VRs have slow response times, and hence cannot support fine-grained dynamic voltage scaling (DVS) which changes the voltage supply levels according to the runtime computation needs of processor [9, 10, 11, 12]. There has been a great deal of progress on fully-integrated buck switching VRs thanks to on-die/in-package inductors and new magnetic materials [13, 14, 15]. Operating at a much higher frequency of tens or hundreds of MHz, fully integrated switching VRs come with fast response times and promises for efficient local power

delivery and fine-grain DVS. However, Integrating high-Q power inductors to support high current density with low loss is still a significant challenge [13, 14, 15]. Compared to their off-chip counterparts, on-chip switching VRs incur more conduction and switching losses, leading to lower efficiency, especially at light loads. As an alternative of on-chip switching buck VRs, on-chip linear voltage regulators (e.g. LDOs) are area efficient and can achieve sub-$ns$ response times [16]. However, their efficiency drops with increasing dropout voltage, making them inefficient for wide-range voltage conversion. Clearly, those VRs have their own pros and cons and the VR topology should be carefully selected according to the specific application requirements.

### 1.1.3 Distributed Voltage Regulation

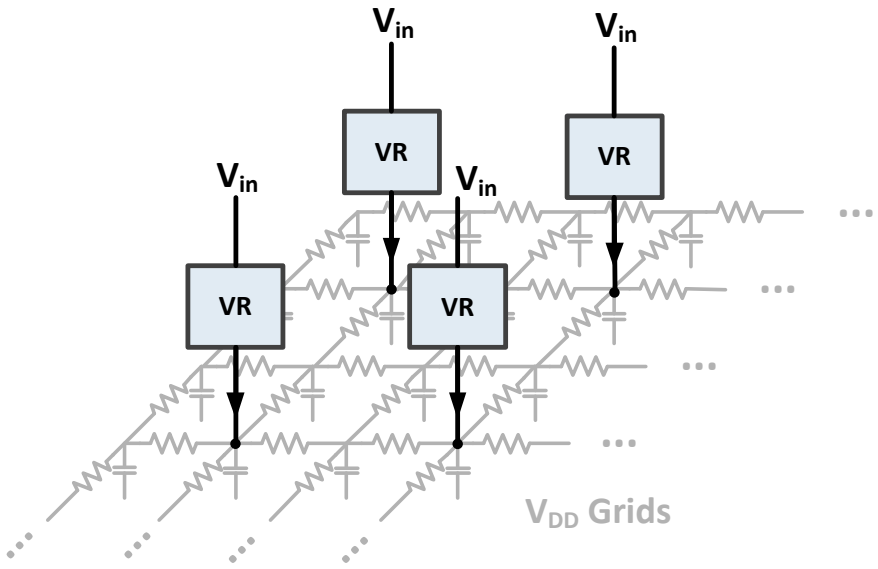


Figure 1.2: Demonstration of distributed voltage regulation.

Recently, it has become an important trend to place multiple on-chip VRs close to heavy-noise circuitries in a distributive manner [17, 18, 19, 20, 16, 21, 22]. The distributed voltage regulation system, as demonstrated in Fig. 1.2, has several important advantages. First, it forms an interconnected active regulation network and shortens the distance between any load location on the power grids and the nearest VR, resulting in reduced $di/dt$ noises and IR drops [16]. The improved sup-

ply noise can be further translated into a lower safe-operating voltage guardband, achieving energy reduction of the entire system. Second, it enables fine-grained power management system where the voltage supply levels can be independently controlled for different circuit blocks. In addition, the use of distributed voltage regulation contributes to a better distribution of heat [23]. Due to all the benefits as discussed above, distributed voltage regulation is receiving a large amount of interests in the design of power delivery network for modern high-performance SoCs and processors. For instance, the recent POWER8 microprocessor integrates over a thousand on-chip micro-VRs to support 48 power domains [21].

## 1.2 Present Challenges in PDN Design

We summarize several important design challenges for modern PDN designs as below:

- **Stability.** Stability is the first-order design consideration for power delivery system. An unstable PDN can induce sustained supply voltage oscillations in the on-chip power grids and cause severe circuit performance degradation or even chip function failure. Although distributed voltage regulation has many appealing benefits as aforementioned, designing such a PDN system with guaranteed stability is very challenging. As illustrated in Fig. 1.3(a), a large number of feedback loops are formed in a distributed PDN system due to the complex interactions between a large number of integrated active VRs, bulky on-chip passive power grids, and on-board/package parasitics, giving rise to pressing stability design challenges. It has been observed that integrating multiple stable LDO VRs with large gain/phase margin can render the entire network unstable (Fig. 1.3(b)), invalidating the conventional phase/gain margin based stability metric targeting single-loop systems. On the other hand, the brute force pole analysis has a cubic computational complexity, and is hence not practical for real-life PDNs.

- **Power efficiency.** Power efficiency has become one of the greatest bottlenecks for modern systems-on-a-chips (SoCs). While transistor sizes continue scaling down, power limitations diminish the fraction of usable chip area [24, 25]. The arrival of the dark silicon age, i.e.,

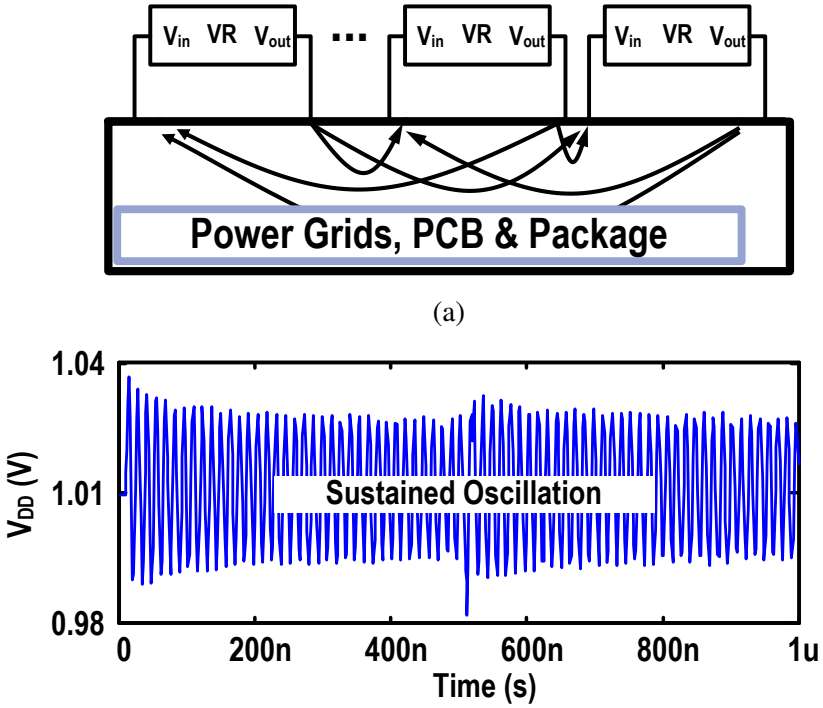


Figure 1.3: (a) A large number of feedback loops formed by complex interactions among active voltage regulators, PCB/package parasitics and bulky power grids, and (b) resultant continuous oscillation in a PDN integrating four LDOs with phase margin $> 100°$.

cores in many-core processors cannot be powered on at the time is precisely due to the unacceptable power and thermal limits. Therefore, power management must be leveraged to improve system power efficiency in every possible way. During runtime, modern processors must adapt to the very dynamic workloads with diverse compute, memory and temporal characteristics with efficiency. At system and architecture level, various power management techniques such as dynamic voltage frequency scaling (DVFS) [9, 10, 11, 12, 26] and clock gating [27, 28] have been proposed to save power and improve the overall processor's performance. At circuit level, dynamic workloads and power management can push the VRs away from their optimal operating points, degrading the efficiency of the entire PDN system. Therefore, modern PDNs have to be workload-aware and re-configurable to minimize the power loss in power delivery.

- **Power quality.** Power must be delivered to on-die devices with high quality to prevent operational failures such as timing violations. In addition to the stability issue highlighted already, several other design challenges arise for maintaining a desirable power delivery quality while maximizing the power efficiency. The ever-growing stringent workload demands in today's high-performance SoCs cause significant $di/dt$ dynamics and IR drops across the power grids, increasing the risk of voltage emergencies (i.e. supply voltage noise exceeding a safe-operating margin). On the other hand, the use of advanced power management techniques further exacerbates the supply noise. For example, power gating [27, 28] incurs large voltage fluctuations in the power supply, while adaptive frequency and voltage control [26, 29, 30, 31] reduces the amount of voltage guardband to save power and thus causes smaller safe-operating margin for noise tolerance. The overall effect is that the PDN is facing more and more design challenges to ensure power quality.

## 1.3 Survey of Previous Work

In the past few decades, there have been a lot of efforts and progresses on the stability-oriented PDN design approaches and adaptive power management for improved tradeoffs among power efficiency and quality.

### 1.3.1 Survey of Stability-Ensuring PDN Design

There are prior works such as [32] exploring the interaction between a single voltage regulator and a passive input network by examining the open-loop characteristic for the coupled system. However those approaches cannot be applied to our PDN system since deriving the transfer function for the bulky passive network is impractical and it cannot find a major loop in the multi-loop PDN system to open for the stability analysis. Targeting at a distributed PDN with multiple on-chip LDOs, [20] employs an adaptive RC compensation network to maintain a high phase margin for each single VR over a wide range of load currents and process, voltage, and temperature (PVT) variations. However, as discussed already, the classical single-loop based stability metrics such as phase/gain margin fail to capture the complex interaction between active VRs and passive net-

works and can lead to misleading stability result. Besides, [20, 33] study the stability of distributed power system through an equivalent single-loop small-signal model, assuming equally shared load current and symmetry VR placement. Such simplified model compromises the accuracy for stability analysis. Therefore, it is highly desirable to conduct more systematic and rigorous analysis towards the stability of complex PDN systems.

Recently, several new design methodologies are proposed by leveraging the control theorems targeting at multiple-input-multiple-output (MIMO) systems. A stability criterion based on the passivity theory has been proposed in [34] and the effectiveness of this approach has been demonstrated on practical power delivery systems. Earlier, [22, 35, 36] have presented a more desirable localized stability-ensuring design methodology based on hybrid stability theory (HST) [37] for large PDNs with distributed voltage regulation. This approach partitions the PDN system into a feedback interconnection of the passive network and the active voltage regulators, and then the HST is applied to check the system stability efficiently. Since HST leverages the passivity and gain of the PDN system simultaneously, more design freedom is provided. Nonetheless, both the passivity theory and HST only provide sufficient conditions for stability, and hence can lead to pessimism in stability analysis and severe performance degradation due to large over-design [38, 39]. Besides, the associated new design approaches based on complicated control theorems are unfamiliar to circuit designers and, as a result, cannot offer much intuition to designers [40].

### 1.3.2 Survey of Adaptive Power Management

To save power, various power management techniques [9, 10, 11, 12, 26] have been proposed recently at the system and architecture level. For example, [9] explores the benefits of fast DVFS at sub-us time scale using on-chip switching regulators. And [26] proposes an adaptive guardbanding approach to dynamically adapt the operating voltage and frequency of the chip based on timing-margin measurements at runtime. These DVFS techniques are more focus on the optimization of processor's power and performance, without exploring the energy reduction opportunity in the PDN which delivers energy to the processor.

At the circuit level, several works have investigated the benefits of workload-aware PDN de-

signs [41, 42, 43, 44, 45]. They attempted to optimize the power efficiency of the system by reconfiguring the PDN according to the practical workload of the processor at runtime. In [41], multiple power domains with the same supply voltage level can be consolidated to share a single off-chip inductor-based buck VR to avoid low conversion efficiency at light load condition. In the same spirit, [42] proposes a reconfigurable PDN in which the output of multiple VRs can be combined when the amount of workload exceeds the maximum workload for a single VR. These PDN reconfiguration techniques are all based on the core- or chip-level workload estimations without considering on-chip distributed LDOs and finer grained spatial workload distribution which can significantly impact power delivery quality. Targeting a 2-stage PDN using both off-chip and on-chip switching VRs, a workload-aware Quantized Power Management (QPM) scheme is proposed in [43] to dynamically adjust the number of active on-chip and off-chip switching VRs at multiple granularities according to the chip-level runtime workload. However, they do not consider the interdependencies among different power stages during power efficiency optimization and lack the demonstration of guaranteed power integrity while tuning the PDN.

## 1.4 Proposed Solutions

This dissertation is motivated by the foregoing problems in the design of modern PDNs. New methodology and approach are proposed for stability-ensured PDN design and workload-aware power management.

### 1.4.1 Proposed Stability-Ensuring PDN Design Methodology

As discussed earlier, the recently developed hybrid stability theory (HST) provides an efficient stability checking and design approach, giving rise to highly desirable localized design of PDNs. However, the inherent conservativenesss of the hybrid stability criteria can lead to pessimism in stability evaluation and hence large overdesign. In this dissertation, we present several new contributions towards establishing a theoretically complete framework for HST-based PDN design and pessimism reduction.

- We first identify the great overdesign reduction opportunities brought by smart PDN system

partitioning. Correspondingly, a frequency-dependent partitioning technique is proposed to do so optimally. By the virtue of the fact that HST is only a sufficient stability condition, the presented partitioning technique allows for bidirectional frequency-dependent admittance splitting between the passive power grids and active voltage regulators in such a way to minimize pessimism in PDN stability analysis. Importantly, we prove that the proposed bidirectional admittance splitting only changes how the system-level building blocks are constructed and internal feedback loops are formed, but does not alter the overall physical PDN system.

- We further show that appropriate application of HST requires that the two partitioned PDN system blocks fall into the $\mathcal{H}_\infty$ space. This leads to the recognition of two important constraints for performing admittance splitting: 1) the frequency-dependent splitting functions need to be well-behaving, i.e. they shall be in the $\mathcal{H}_\infty$ space to ensure that the re-partitioned admittance matrices of both the passive block and regulator block are in the $\mathcal{H}_\infty$ space; 2) the impedance matrix of the re-partitioned passive block, which is the inversion of its admittance matrix, shall not have singularities in the right half complex plane. The first constraint can be satisfied by defining the frequency-dependent splitting functions with a set of well-behaving filters. Although the second constraint is a non-problem for the original RLC passive block whose admittance and impedance matrices are always stable, the artificially partitioned passive block can violate such constraint and thus invalidate the application of HST. To meet the second constraint, we propose to adopt the generalized Nyquist stability criterion for multiple-input-multiple-output (MIMO) systems to check the stability of the inversion of the admittance matrix (i.e. impedance matrix) of the passive block. Correspondingly, we append the stability examination of the impedance matrix of the passive block to an optimal hybrid stability margin (HSM) evaluator to rigorously assess the whole-system stability. This leads to a theoretically rigorous and robust automated PDN design flow for joint optimization of key system specifications such as stability, regulation performance, and power efficiency via localized design of the voltage regulators.

- Using a comprehensive set of design studies, we demonstrate that our stability evaluator significantly reduces pessimism in stability analysis, which translates into large performance gains produced by our automated design flow. When compared to the classical phase margin design approach which provides no guarantee of stability, the proposed approach ensures stability while improving system performance by up to 53%, measured by a figure of merit (FOM). Furthermore, on average our approach boosts the FOM by 113% while consuming 11% less power compared to a reference hybrid stability design approach.

- As helpful as the new HST-based PDN design methodology is, the new stability margin is unfamiliar to circuit designers and, as a result, the above design approach cannot offer much intuition to designers, and therefore stability-constrained design intuitions are derivable. In this dissertation, our systematic analysis further reveals unique design considerations which can significantly impact the system-level stability and performances. Within a large design space, a comprehensive set of design studies are conducted to shed light on the tradeoffs between the HST-based stability margin and other PDN design specifications such as the quiescent current consumption, maximum switching noise, and area overhead. Useful design insights like how regulator topology, passive decoupling capacitance, and the number of on-chip regulators may be optimized for improved tradeoffs between stability and system performance are discussed. These useful insights can aid circuit designers to make appropriate design choices at the beginning of the design process for improved system tradeoffs.

### 1.4.2 Proposed Adaptive Power Management

In terms of workload-aware power management, improvements on both the PDN architecture and control policies are explored in this dissertation. The proposed work is based on the vision that the ultimate power quality and efficiency may be best achieved via a heterogeneous chain of voltage processing starting from on-board switching VRs, to on-chip switching VRs, and finally to networks of distributed on-chip linear voltage regulators. As depicted in Fig. 1.4, we propose a heterogeneous voltage regulation (HVR) architecture encompassing regulators with complimen-

tary characteristics in response time, size, and efficiency. This dissertation aims to answer the following key question for the first time. Given a desired power supply voltage set by a higher-level power management policy, e.g. one based on DVFS, for each power domain, how shall the voltage regulators in the HVR system be adapted autonomously with respect to workload change at multiple temporal scales to significantly improve system power efficiency while providing a guarantee for power integrity? The contributions of the proposed work are severalfold:

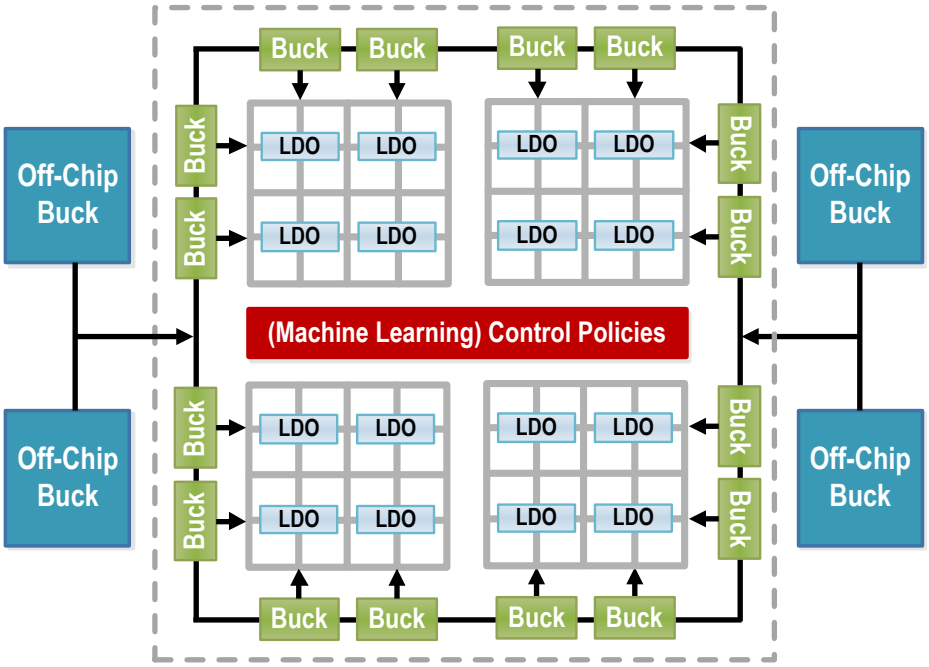


Figure 1.4: Illustration of proposed heterogeneous voltage regulation (HVR).

- We systematically explores the potential benefits of HVR. The most general form of HVR consists of voltage regulators (VRs) with complimentary characteristics across three processing stages. In the first two stages, off- and on-chip switching (DC-DC) buck converters are employed to achieve high efficiency over a wide output voltage range, serving the major role of *voltage conversion*. Compared with single stage DC-DC conversion, two stages DC-DC conversion allows for area reduction, improved power efficiency, and fine-gained DVFS,

which is supported by the fast response time of on-chip DC-DC converters. Unlike conventional one or two-stage power delivery networks, HVR largely decouples *voltage conversion* from *voltage regulation*, the latter of which is optimally achieved by placing a large number of compact LDOs with sub-ns response time in a distributed manner within each power domain, forming an interconnected active regulation network.

- We propose systematic workload-aware control policies to jointly optimize power efficiencies of all voltage processing stages to maximize the overall system power efficiency. To best exploit the potential of energy efficiency of HVR, our control policies minimize system power losses by considering interdependencies across the entire voltage processing chain and adapt HVR at multiple time scales given the significantly different response times of the considered VRs.

- Uncertainties caused by unknown non-uniform spatial distribution of the workload are hard to predict but can jeopardize power integrity. To minimize the extra voltage margin, hence power loss, we propose a novel machine learning (ML) solution for the first time that accurately sets the output voltage of the on-chip switching VRs to maximize the system power efficiency while effectively tracking the worst-case voltage drop in each power domain to safeguard power integrity. Our ML solution consists of a few on-chip voltage-noise sensors that provides inputs to a low-overhead hardware-accelerated ML predictor, which fine tunes the output voltage of the on-chip switching VRs. This provides an autonomous end-to-end integrated ML solution whose low latency allows for fine-grained adaptation of HVR.

### 1.4.3 Organization of Dissertation

The rest of this dissertation is organized as follows. In Chapter II, we propose an efficient and theoretically rigorous technique to assess the hybrid stability of large PDNs with integrated LDOs. The achieved pessimism reduction in stability checking has been leveraged in a new automated design flow that delivers significant performance improvements while ensuring the network-wide stability of the PDN. The presented HST-based PDN design methodology is leveraged in Chap-

ter III to systematically explore the large design space of distributed power delivery, and several new design insights in terms of the network-wide PDN stability are summarized to facilitate improved PDN design tradeoffs. In Chapter IV, we propose a heterogeneous voltage regulation (HVR) architecture encompassing regulators with complimentary characteristics in response time, size and efficiency. Systematic workload-aware power management policies are then developed to adapt heterogeneous VRs with respect to workload change at multiple temporal scales to significantly improve system power efficiency while providing a guarantee for power integrity. The proposed techniques are further supported by hardware-accelerated machine learning prediction of non-uniform spatial workload distributions for more accurate HVR adaption at fine time granularity. Chapter V concludes this dissertation and discusses the future work.

# 2. STABILITY-ENSURING DESIGN OPTIMIZATION FOR LARGE PDNS WITH DISTRIBUTED ON-CHIP VOLTAGE REGULATION *

As stated in Chapter I, the recently developed hybrid stability theory (HST) provides an efficient stability checking and design approach, giving rise to highly desirable localized design of PDNs. However, the inherent conservativenesss of the hybrid stability criteria can lead to pessimism in stability evaluation and hence large overdesign. In this chapter, we address this challenge by proposing an optimal frequency-dependent system partitioning technique to significantly reduce the amount of pessimism in stability analysis. With theoretical rigor, we then show how to partition a PDN system by employing optimal frequency-dependent admittance splitting between the passive network and voltage regulators while maintaining the desired theoretical propoerties of the partitioned system blocks upon which the hybrid stability principle is anchored. At last, we demonstrate a new stability-ensuring PDN design approach with the proposed over-design reduction technique using an automated optimization flow which significantly boosts regulation performance and power efficiency.

## 2.1 Background

Firstly, the theoretical background and framework of the hybrid stability theory (HST) are laid out. The HST is built on the $L_2[0,\infty)$-space of square integrable functions defined as $L_2[0,\infty) = \{\mathbf{v} : \mathbb{R}^+ \rightarrow \mathbb{R}^m | \int_0^\infty \mathbf{v}^{\mathrm{T}}(t)\mathbf{v}(t)dt < \infty\}$, where $\mathbf{v}$ is a vector function of time and $\mathbf{v}^{\mathrm{T}}$ is its transpose. A system is said to be $L_2$-stable if for any input $\mathbf{e} \in L_2[0,\infty)$ the resultant output $\mathbf{y} \in L_2[0,\infty)$. We consider the negative feedback interconnection of two multiple-input-multiple-output (MIMO) blocks $\mathbf{G}$ and $\mathbf{H}$ as in Fig. 2.1, where the $\mathbf{G}$ block has input $\mathbf{e_G}$ and output $\mathbf{y_G}$, and the $\mathbf{H}$ block has input $\mathbf{e_H}$ and output $\mathbf{y_H}$. To assess the stability of the entire feedback system

with input $\mathbf{e}$ and output $\mathbf{y}$, two classical theorems are introduced first, and HST which combines both of them follows.

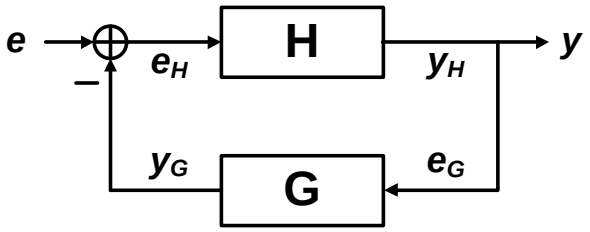


Figure 2.1: General negative feedback interconnection of two MIMO blocks.

### 2.1.1 Classical Stability Theorems

#### *2.1.1.1 Small Gain Theorem*

**Definition 1.** *(Finite gain) A general square system with input* $\mathbf{e} \in L_2[0,\infty)$ *and output* $\mathbf{y} \in L_2[0,\infty)$ *mapped through the operator* $\mathbf{M} : L_2[0,\infty) \to L_2[0,\infty)$ *possesses a 'finite gain' if there exists* $0 < \gamma < \infty$ *such that*

$$\gamma^{-1} \int_0^T \mathbf{y}^T(t)\mathbf{y}(t)dt \leq \gamma \int_0^T \mathbf{e}^T(t)\mathbf{e}(t)dt \quad \forall \mathbf{e} \in L_2[0,\infty), \forall T \geq 0. \tag{2.1}$$

For any LTI system the system gain $\gamma$ can be computed in frequency domain

$$\gamma = \sup_{\forall\omega} \gamma(\omega), \tag{2.2}$$

where $\gamma(\omega)$ is the local system gain at a single frequency $\omega$

$$\gamma(\omega) = \overline{\sigma}\{\mathbf{M}(j\omega)\}, \tag{2.3}$$

where $\overline{\sigma}\{\cdot\}$ is the maximum singular value.

Intuitively, if the loop gain of a feedback system is less than 1, the system shall remain stable as any oscillation through the loop will be finally attenuated. This understanding leads to the following small-gain theorem.

**Theorem 1.** *(Small-gain theorem) The negative feedback interconnection in Fig. 2.2 is $L_2$-stable if the product of the system gains is less than one, i.e. $\gamma_{\mathbf{G}}\gamma_{\mathbf{H}} < 1$, where $\gamma_{\mathbf{G}}$ and $\gamma_{\mathbf{H}}$ are the gains of the $\mathbf{G}$ block and the $\mathbf{H}$ block, respectively.*

However, as the small gain theorem only utilizes the gain information, harsh stability constrains may be placed on the full system. For example, at certain frequency one block can have a very high gain, which in turn requires the other block to have a rather low gain, potentially resulting to poor design performance.

#### *2.1.1.2 Passivity Theorem*

**Definition 2.** *(Passive systems) A general square system with input $\mathbf{e} \in L_2[0,\infty)$ and output $\mathbf{y} \in L_2[0,\infty)$ mapped through the operator $\mathbf{M} : L_2[0,\infty) \to L_2[0,\infty)$ is 'very strictly passive' if there exist $\delta > 0$ and $\epsilon > 0$ such that*

$$\int_0^T \mathbf{e}^T(t)\mathbf{y}(t)dt \geq \delta \int_0^T \mathbf{e}^T(t)\mathbf{e}(t)dt + \epsilon \int_0^T \mathbf{y}^T(t)\mathbf{y}(t)dt \quad \forall \mathbf{e} \in L_2[0,\infty), \forall T \geq 0. \tag{2.4}$$

Furthermore, if $\delta > 0$ and $\epsilon = 0$, the system is said to be 'input strictly passive'; if $\delta = 0$ and $\epsilon > 0$, the system is said to be 'output strictly passive'; if both $\delta = 0$ and $\epsilon = 0$, the system is a 'passive' system. Similarly, the passivity for LTI system can be examined along the frequency axis. The input passivity parameter is calculated as

$$\delta = \inf_{\forall \omega} \delta(\omega), \tag{2.5}$$

where $\delta(\omega)$ is the local input passivity parameter

$$\delta(\omega) = \underline{\lambda}\{\mathbf{M}^{\mathrm{H}}(j\omega) + \mathbf{M}(j\omega)\}, \tag{2.6}$$

where $\underline{\lambda}\{\cdot\}$ is the minimum eigenvalue. Additionally, a system that is already input strictly passive with finite gain is output strictly passive, i.e. there exists output passivity parameter $\epsilon > 0$, and hence is very strictly passive.

As passive systems do not produce energy, it is intuitive that the feedback loop composed of two passive blocks is stable, as the total energy stored in the system decreases with time. Accordingly, the passivity theorem states the following useful result for stability assessment.

**Theorem 2.** *(Passivity theorem) The negative feedback interconnection in Fig. 2.2 is $L_2$-stable if $\epsilon_{\mathbf{G}} + \delta_{\mathbf{H}} > 0$ and $\epsilon_{\mathbf{H}} + \delta_{\mathbf{G}} > 0$, where $\delta_{\mathbf{G}}$, $\epsilon_{\mathbf{G}}$ are the passivity parameters of the $\mathbf{G}$ block, and $\delta_{\mathbf{H}}$, $\epsilon_{\mathbf{H}}$ are those of the $\mathbf{H}$ block.*

Furthermore, the following corollary states that the passivity theorem can be applied to check the system stability only based on the input passivity parameter $\delta$ of each block.

**Corollary 1.** *The feedback system is $L_2$-stable if both $\delta_{\mathbf{G}} > 0$ and $\delta_{\mathbf{H}} > 0$.*

Similar to the small gain theorem, the passivity theorem which merely utilizes the passivity condition, is not a silver bullet guaranteeing stability as well. Practical active circuit blocks such as voltage regulators can only exhibit local passivity within a certain frequency range.

### 2.1.2 Hybrid Stability Theorem

It is highly desirable to combine the two basic theories together to guarantee stability, leading to the following hybrid stability theorem.

**Theorem 3.** *(Hybrid stability theorem) The negative feedback interconnection in Fig. 2.2 is $L_2$ stable if for all frequencies at least one of the following conditions is met:*

*1) (Passivity condition) $\delta_{\mathbf{G}}(\omega) > 0$ and $\delta_{\mathbf{H}}(\omega) > 0$.*

*2) (Gain condition) $\gamma_{\mathbf{G}}(\omega)\gamma_{\mathbf{H}}(\omega) < 1$.*

Appealingly, HST exploits two complimentary principles, passivity and gain, making it possible to rather efficiently ensure system stability by judiciously selecting one of the two principles as the most cost-effective mechanism at each given frequency. Thus, the entire frequency domain can be broken into passivity region ($\Omega_{pass}$) and gain region ($\Omega_{gain}$), where passivity condition and gain condition are utilized respectively. While HST offers greater design freedoms, it only provides *sufficient* condition for stability.

## 2.2 HST-Based Stability Checking for PDNs

### 2.2.1 PDN System Partitioning

The application of HST to the stability evaluation of PDNs with distributed on-chip voltage regulators [35, 22] is outlined in this subsection. To begin with, we show a model of a PDN system with $n$ on-chip LDOs in Fig. 2.2(a), where the PDN is partitioned into two system blocks: block $\mathbf{G}$ comprising LDOs and block $\mathbf{H}$ containing the bulky passive network. Note that $\mathbf{G}$ interfaces with $\mathbf{H}$ through both the global VDD grids and regulated local power grids. Thus for a PDN embedded with $n$ LDOs, $\mathbf{G}$ and $\mathbf{H}$ both have $2n$ ports, and are described using a $2n \times 2n$ transfer function matrix. We model $\mathbf{G}$ with $Y$-parameters and $\mathbf{H}$ with $Z$-parameters. The identified system-wide negative feedback loops are shown in Fig. 2.2(b). The loops start from $\mathbf{H}$'s current inputs $\mathbf{i_H}$, which induces the voltage outputs (fluctuations) $\mathbf{v_H}$ of $\mathbf{H}$. $\mathbf{v_H}$ feeds $\mathbf{G}$ and produces $\mathbf{G}$'s current outputs $\mathbf{i_G}$. As $\mathbf{i_G}$ and $\mathbf{i_H}$ have the same magnitudes but opposite in sign, an inverter is placed to indicate the negative feedback involved.

### 2.2.2 Hybrid Stability Metrics

To quantitatively measure the passivity and gain of the entire PDN system consisting of the blocks $\mathbf{G}$ and $\mathbf{H}$ as shown in Fig. 2.2, we further define a passivity margin ($PassMarg$) and a gain margin ($GainMarg$) concept at each frequency $\omega$:

$$PassMarg(\omega) = \begin{cases} -\sqrt{\delta_{\mathbf{G}}(\omega)^2 + \delta_{\mathbf{H}}(\omega)^2}, \delta_{\mathbf{G}}(\omega) < 0, \delta_{\mathbf{H}}(\omega) < 0 \\ min\{\delta_{\mathbf{G}}(\omega), \delta_{\mathbf{H}}(\omega)\}, \end{cases} \tag{2.7}$$

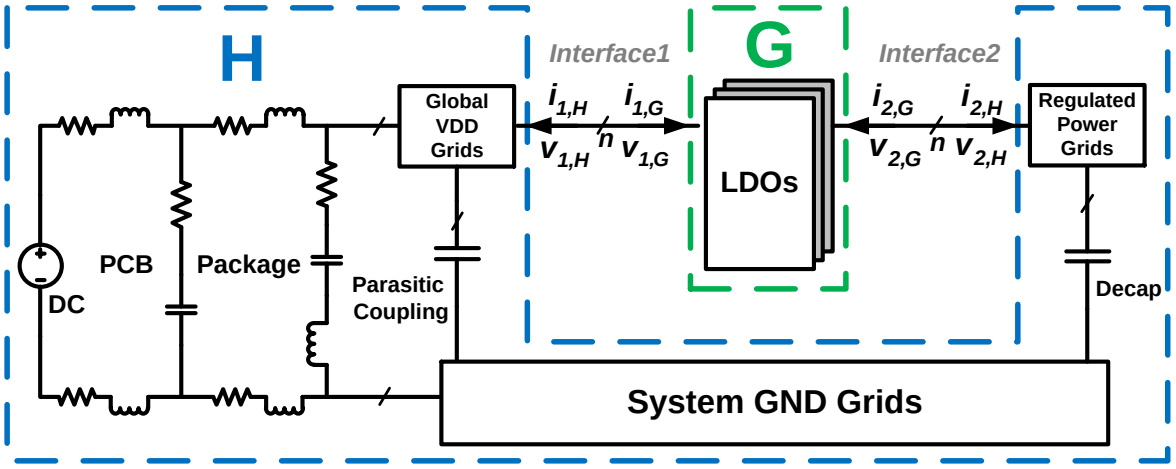


(a)

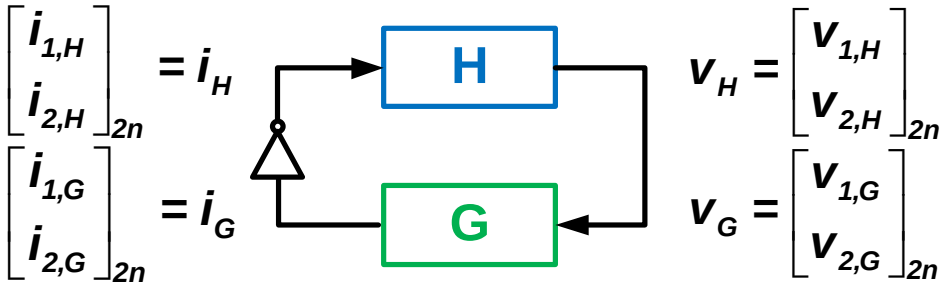


(b)

Figure 2.2: Partition the PDN into a negative feedback loop: (a) system model, and (b) block diagram.

$$GainMarg(\omega) = 1 - \gamma_{\mathbf{G}}(\omega)\gamma_{\mathbf{H}}(\omega) \tag{2.8}$$

Obviously, at each frequency point, the system possesses a positive $PassMarg$ iff both subsystems are passive while a positive $GainMarg$ is observed iff the loop gain of the whole system is less than one. Based on $PassMarg(\omega)$ and $GainMarg(\omega)$, the hybrid stability margin (HSM) at frequency $\omega$ is defined in (2.9), which combines the gain and passivity margins into a single stability measure. Finally, the *global HSM* is defined as the minimum HSM($\omega$) across the entire frequency axis. In light of Theorem 3, the system is stable if the global HSM is positive.

$$HSM(\omega) = \begin{cases} \sqrt{PassMarg(\omega)^2 + GainMarg(\omega)^2}, & PassMarg(\omega), GainMarg(\omega) \geq 0 \\ max\{PassMarg(\omega), GainMarg(\omega)\}, & otherwise \end{cases} \tag{2.9}$$

### 2.2.3 Localized Stability-Ensuring Design Flow for PDNs

In the PDN design process, active voltage regulators, such as LDOs, may be re-designed multiple times with the passive power grids fixed in order to meet certain specifications on stability and performance. As such, by separating the bulky passive part of the PDN from LDOs, the system partitioning in Fig. 2.2(a) allows us to spend only a one-time effort in characterizing the gain and passivity of the RLC passive network and then focus solely on the design of LDOs, leading to an efficient localized regulator design methodology while ensuring the full system stability. A localized stability-ensuring design flow for PDNs is illustrated in Fig. 2.3. Note that only the low-cost regulator characterization is repeated in design iterations (as in the grey box in Fig. 2.3). However, the inherent conservativeness of HST can lead to large over-design, i.e. performance loss, which is the key challenge we aim to address in this work.

It should be noted that the underlying HST framework works on LTI system and we need to describe the $\mathbf{G}$ and $\mathbf{H}$ blocks using linear models. Although it is convenient to characterize the linear analog LDOs via AC simulation, it is also possible to extend the HST-based design flow in Fig. 2.3 to other regulator topologies such as digital LDOs and DC-DC buck converters by

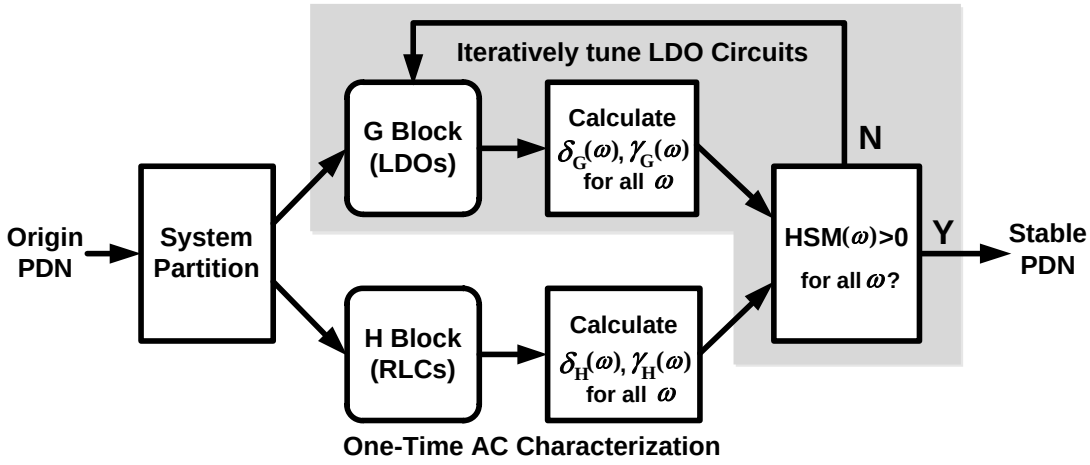


Figure 2.3: Localized stability-ensuring design flow for PDNs.

obtaining an approximated linear model for those voltage regulators. In the past several decades, there has been a large body of work focusing on the modeling methodologies such as [46], [32] and [47].

## 2.3 Proposed Admittance Splitting

### 2.3.1 Motivations

Recall that HST is only a sufficient but not necessary condition for stability checking. It comes with the following important design implications. Meeting the HST stability criterion immediately confirms the stability of the system under discussion. However, failing it with a negative global HSM does not fully exclude the possibility of stability. In the latter case, striving for a positive HSM margin could lead to unnecessary over-design due to the uncertainty imposed by the HST based stability checking.

We make the following key observation regarding over-design reduction. While HST assesses the stability of an interconnection of two subsystems (Theorem 3), how the overall system shall be partitioned (the first step in Fig. 2.3) is completely left to us. This motivates us to leverage system partitioning to significantly to reduce the conservativeness of the hybrid stability theory.

We use a realistic PDN design to shed light on the above idea. The PDN is first partitioned into the $\mathbf{G}_0$ and $\mathbf{H}_0$ blocks as in Fig. 2.2(a), and then $GainMarg$ and $PassMarg$ are examined

in Table 2.1. The minimum $GainMarg$ is $-\infty$ due to the nonexistence of DC paths between the regulated power grids and the GND grids. The minimum $PassMarg$ is -31.9, showing a severe passivity violation. Next, we repartition the same PDN without altering the overall system in two different ways. We move 10% of each self-admittance of the $2n \times 2n$ port model of $\mathbf{G}_0$ from $\mathbf{G}_0$ to $\mathbf{H}_0$ to form a new partition of $\mathbf{G}_1$ and $\mathbf{H}_1$. In the opposition direction, we move 10% of each self-admittance of $\mathbf{H}_0$ from $\mathbf{H}_0$ to $\mathbf{G}_0$ to form another partition of $\mathbf{G}_2$ and $\mathbf{H}_2$. Table 2.1 shows that these two repartitions dramatically improve the system's min $GainMarg$ and min $PassMarg$, respectively, suggesting that partitioning has a large impact on the HSM.

Table 2.1: Gain and passivity of different partitions.

| Partition | $Min\{GainMarg\}$ | $Min\{PassMarg\}$ |
|---|---|---|
| $\mathbf{G}_0$+$\mathbf{H}_0$ | $-\infty$ | $-31.9$ |
| $\mathbf{G}_1$+$\mathbf{H}_1$ | $-11.9$ | – |
| $\mathbf{G}_2$+$\mathbf{H}_2$ | – | $-3.7$ |

### 2.3.2 Frequency-Dependent Bidirectional Admittance Splitting

From Theorem 3 and (2.9), it is evident that maximizing the global hybrid stability margin (HSM) for a given PDN could maximize its chance to pass the HST checking and therefore reduce the theoretical pessimism. This can be in turn achieved by maximizing the local HSM for each frequency as in (2.9). To do so, it is instrumental to improve either $PassMarg$ or $GainMarg$, or both by finding the best system partitioning for each frequency $\omega$, leading to the idea of frequency-dependent admittance splitting. One further development is to allow *bidirectional* splitting of self-admittances (or self-impedances) across the boundary of the $\mathbf{G}$ and $\mathbf{H}$ blocks as below.

#### *2.3.2.1 Small-Gain Enhancement via Admittance Splitting*

As shown in Fig. 2.4, splitting self-admittances from $\mathbf{G}$ to $\mathbf{H}$ helps lower the overall loop gain, and therefore improves $GainMarg$. Intuitively, this is because lowering the self-admittance values of the $Y$-parameter model of the $\mathbf{G}$ block tends to reduce its voltage-to-current gain. At

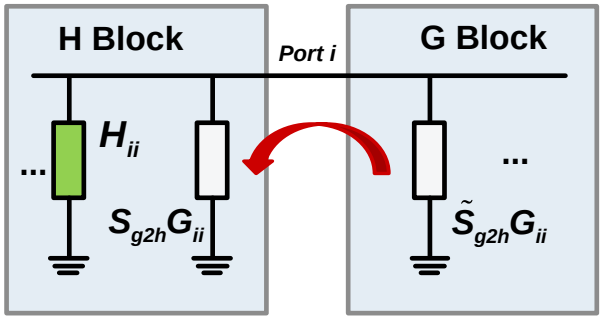


Figure 2.4: Splitting admittance from G block to H block to improve the gain margin.

the same time, adding additional self admittances to the **H** block lowers the impedance from each port to the ground, reducing the gain of the **H** block. Formally, we define a frequency-dependent splitting factor $S_{g2h}(j\omega)$ to split each self-admittance of the **G** block into two parts: the portion to be moved from **G** to **H** corresponds to $S_{g2h}(j\omega)$, and the remaining portion corresponds to $\widetilde{S}_{g2h}(j\omega) = 1 - S_{g2h}(j\omega)$.

#### 2.3.2.2 *Passivity Enhancement via Admittance Splitting*

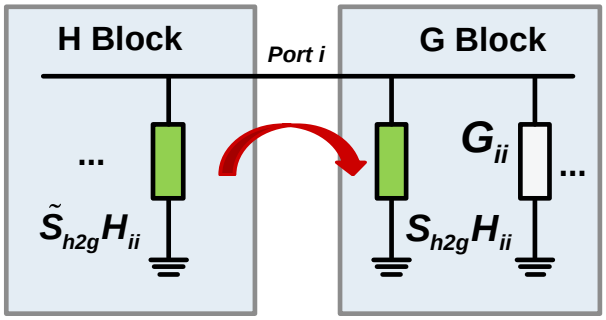


Figure 2.5: Splitting admittance from H block to G block to improve the passivity margin.

Fig. 2.5 shows the possibility of splitting self-admittances from the **H** block to the **G** block using similarly defined splitting factors $S_{h2g}(j\omega)$ and $\widetilde{S}_{h2g}(j\omega)$ to improve $PassMarg$. This is based on the observation that the original **H** network consists of passive RLC components and is always passive. It is possible to remove certain self-admittances from the **H** block while maintain-

ing its passivity. In the meantime, adding additional self-admittances into the $\mathbf{G}$ block enhances its passivity property. To demonstrate this important trend, the detailed analysis is derived as below.

Firstly note that the LDOs in the $\mathbf{G}$ block are isolated to each other, accordingly the original $\mathbf{G}_{2n\times 2n}$ is a block diagonal matrix and its local passivity $\delta_{\mathbf{G}}(\omega)$ can be evaluated via each individual LDO [22]:

$$\delta_{\mathbf{G}}(\omega) = min_{i=1,2,...n}\delta_{\mathbf{Y}_i}(\omega) \tag{2.10}$$

where $\mathbf{Y}_i$ is the $2 \times 2$ $Y$-parameter matrix of the $i$-th LDO. According to (2.6), the passivity of a single LDO at a certain frequency $\omega$ can be calculated as:

$$\delta_{\mathbf{Y}} = r_{11} + r_{22} - \sqrt{(r_{11} + r_{22})^2 - (4r_{11}r_{22} - \Delta)} \tag{2.11}$$

where $r_{ij} = \mathrm{Re}\{Y_{ij}\}$, $m_{ij} = \mathrm{Im}\{Y_{ij}\}$ and $\Delta = (r_{12} + r_{21})^2 + (m_{12} - m_{21})^2$. Evidently, the LDO behaves like a passive circuit at this frequency iff the following conditions are met:

$$\begin{cases} r_{11} + r_{22} > 0 & \text{(2.12a)} \\ 4r_{11}r_{22} > (r_{12} + r_{21})^2 + (m_{12} - m_{21})^2 & \text{(2.12b)} \end{cases}$$

Notice that $\mathbf{Y}$'s first-column elements $Y_{11}$ and $Y_{21}$ are smaller in magnitude than the second-column elements $Y_{12}$ and $Y_{22}$, approximately by a factor of $A_{LL}$ (denoted as local loop gain of LDO). Additionally, the output admittance of the regulator often has a positive real part, i.e. $r_{22} > 0$. Therefore (2.12a) can be easily satisfied over a wide range of frequencies. On the other hand, usually (2.12b) can be met only beyond its unit-gain bandwidth (UGB) where $A_{LL}$ drops below unit. This is corresponding to the phenomenon that the original $\mathbf{G}$ block exhibits local passivity merely at high frequencies.

Now we split self-admittances from $\mathbf{H}$ to $\mathbf{G}$. It is obvious that such admittance splitting helps increase $r_{11}$ and $r_{22}$ and boost the left side of (2.12b) dramatically, thus the passivity of the regulator

is enhanced. Besides, the partial differential equations can be derived as

$$\begin{cases} \dfrac{\partial \delta_{\mathbf{Y}}}{\partial r_{11}} = 1 - \dfrac{r_{11} - r_{22}}{\sqrt{(r_{11} - r_{22})^2 + \Delta}} > 0 & \text{(2.13a)} \\ \dfrac{\partial \delta_{\mathbf{Y}}}{\partial r_{22}} = 1 - \dfrac{r_{22} - r_{11}}{\sqrt{(r_{22} - r_{11})^2 + \Delta}} > 0 & \text{(2.13b)} \end{cases}$$

indicating that $\delta_{\mathbf{Y}}$ is a monotonically increasing function in terms of $r_{11}$ and $r_{22}$. According to (2.10), the passivity of the $\mathbf{G}$ block grows by receiving admittances from the $\mathbf{H}$ block.

The analyses in this subsection illustrate that at a certain frequency, splitting admittance from $\mathbf{G}$ to $\mathbf{H}$ can improve $GainMarg$ while splitting admittance from $\mathbf{H}$ to $\mathbf{G}$ improves $PassMarg$. As will be described later, an optimization-based approach has been developed to find the optimal admittance splitting direction and the amount of splitting for each frequency, which maximizes the chance of the system to pass the hybrid stability checking. Clearly, the resulting optimal HSM provides a more truthful assessment of stability with reduced pessimism.

### 2.3.3 System Invariance with Admittance Splitting

While system repartitioning via the proposed frequency-dependent bidirectional admittance splitting has the potential for overdesign reduction, this benefit is meaningful only if such admittance splitting does not alter the physical PDN system. We rigorously prove that the overall PDN transfer function is indeed not altered by the admittance splitting, which only changes the boundary of the PDN partitioning and how the internal feedback loops are represented. To begin with, in the close-loop feedback system shown in Fig. 2.2(b), we identify the input-output transfer matrix where the input current noises induce the output voltage fluctuations as

$$\mathbf{\Phi}(j\omega) = (\mathbf{I} + \mathbf{H}(j\omega)\mathbf{G}(j\omega))^{-1}\mathbf{H}(j\omega). \tag{2.14}$$

Then the following important property is derived.

**Property 1.** *The frequency-dependent bidirectional admittance splitting technique does not change the input-output transfer matrix $\mathbf{\Phi}(j\omega)$ in the close-loop system in Fig. 2.2(b).*

*Proof.* Firstly, the inverse transfer matrix can be written as

$$\mathbf{\Phi}^{-1}(j\omega) = \mathbf{G}(j\omega) + \mathbf{H}^{-1}(j\omega). \tag{2.15}$$

Considering Fig. 2.2(a), again, before applying the proposed admittance splitting, let us denote the admittance matrices of the passive network and regulator blocks by $\mathbf{H}_0^Y(j\omega)$ and $\mathbf{G}_0^Y(j\omega)$, respectively. To relate to the negative feedback interconnection of Fig. 2.2(a), we invert $\mathbf{H}_0^Y(j\omega)$ to get the impedance representation of the passive network. In other words, we have: $\mathbf{H}^{-1}(j\omega) = \mathbf{H}_0^Y(j\omega)$ and the corresponding $\mathbf{\Phi}^{-1}(j\omega)$ without admittance splitting is

$$\mathbf{\Phi}_0^{-1}(j\omega) = \mathbf{G}_0^Y(j\omega) + \mathbf{H}_0^Y(j\omega). \tag{2.16}$$

Next, we apply the proposed frequency-dependent bidirectional admittance splitting scheme to the PDN system. As mentioned earlier, the two splitting factors $S_{g2h}(j\omega)$ and $S_{h2g}(j\omega)$, which correspond to the admittance splitting in two different directions, operate complementarily along frequency axis attempting to increase the local HSM at each frequency $\omega$. According to Fig. 2.5 and Fig. 2.4, the re-partitioned $\mathbf{G}(j\omega)$ and $\mathbf{H}^{-1}(j\omega)$ are

$$\mathbf{G}(j\omega) = \mathbf{G}_0^Y(j\omega) - S_{g2h}(j\omega)\widehat{\mathbf{G}}_0^Y(j\omega) + S_{h2g}(j\omega)\widehat{\mathbf{H}}_0^Y(j\omega) \tag{2.17}$$

$$\mathbf{H}^{-1}(j\omega) = \mathbf{H}_0^Y(j\omega) + S_{g2h}(j\omega)\widehat{\mathbf{G}}_0^Y(j\omega) - S_{h2g}(j\omega)\widehat{\mathbf{H}}_0^Y(j\omega) \tag{2.18}$$

where $\widehat{\mathbf{G}}_0^Y(j\omega)$ and $\widehat{\mathbf{H}}_0^Y(j\omega)$ are the two diagonal matrices consisting of the diagonal entries of $\mathbf{G}_0^Y(j\omega)$ and $\mathbf{H}_0^Y(j\omega)$, respectively:

$$\widehat{\mathbf{G}}_0^Y(j\omega) = \begin{bmatrix} G_{1,1}(j\omega) & 0 & \dots & 0 \\ 0 & G_{2,2}(j\omega) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & G_{2n,2n}(j\omega) \end{bmatrix}_{2n\times 2n},$$

$$\widehat{\mathbf{H}}_0^Y(j\omega) = \begin{bmatrix} H_{1,1}(j\omega) & 0 & \dots & 0 \\ 0 & H_{2,2}(j\omega) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & H_{2n,2n}(j\omega) \end{bmatrix}_{2n\times 2n}.$$

Making the appropriate substitutions for $\mathbf{G}(j\omega)$ and $\mathbf{H}^{-1}(j\omega)$, it follows that:

$$\begin{aligned} \mathbf{\Phi}^{-1}(j\omega) &= \mathbf{G}(j\omega) + \mathbf{H}^{-1}(j\omega) \\ &= (\mathbf{G}_0^Y(j\omega) - S_{g2h}(j\omega)\widehat{\mathbf{G}}_0^Y(j\omega) + S_{h2g}(j\omega)\widehat{\mathbf{H}}_0^Y(j\omega)) \\ &\quad + (\mathbf{H}_0^Y(j\omega) + S_{g2h}(j\omega)\widehat{\mathbf{G}}_0^Y(j\omega) - S_{h2g}(j\omega)\widehat{\mathbf{H}}_0^Y(j\omega)) \\ &= \mathbf{G}_0^Y(j\omega) + \mathbf{H}_0^Y(j\omega) \end{aligned}$$

By comparing (2.16) with above equation, it is clear that the frequency dependent bidirectional admittance splitting does not change the transfer matrix of the entire system. □

The proposed admittance splitting technique does not change the physical PDN system by virtue of the fact that the total amount of admittance at each port remains the same. It means that we do not add or remove elements into or from the system. In this scheme, only the way modeling the system is changed. As the HST is only a sufficient condition, different admittance splitting schemes can lead to different degrees of pessimism and therefore provides us potential opportunities to reduce the theoretical conservativeness.

## 2.4 Theoretical Basis for Admittance Splitting

We have demonstrated the great opportunity for over-design reduction brought by the proposed frequency-dependent admittance splitting scheme. However, to apply the HST in a rigorous way, each individual partitioned block needs to maintain certain key theoretical properties upon which the HST is anchored.

Most physical systems of interest such as PDNs are causal. For a general system $\mathbf{M}$, it is causal if and only if its impulse response satisfies $m(t) = 0, \forall t < 0$. Note that the signals that are considered in the HST framework are defined on $L_2[0,\infty)$-space $\{\mathbf{v} : \mathbb{R}^+ \to \mathbb{R}^m | \int_0^\infty \mathbf{v}^T(t)\mathbf{v}(t)dt < \infty\}$ which are assumed to be causal. Importantly, this implies that the proposed frequency-dependent admittance splitting must maintain the following key system properties: the resulting system blocks $\mathbf{G}$ and $\mathbf{H}$ must be causal; and the outputs of each block in response to inputs in $L_2[0,\infty)$ must be in $L_2[0,\infty)$.

While the above requirements are a nonproblem in well-behaving physical systems, artificially performing admittance splitting can violate these theoretical conditions and thus jeopardize the foundation of the hybrid stability analysis. However, since the proposed frequency-dependent admittance splitting technique is performed in the frequency domain, it is not convenient for us to examine whether the re-partitioned $\mathbf{G}$ and $\mathbf{H}$ blocks satisfy the time-domain requirements as above. Therefore, we need to interpret the requirements from time domain to frequency domain or $s$ domain to constrain the frequency dependent slitting technique and maintain the theoretical rigor. For this purpose, we introduce the definition of $\mathcal{H}_2$-space and a theoretical result [48, 49].

**Definition 3.** *($\mathcal{H}_2$-space) The Hardy 2-space is defined as*

$$\mathcal{H}_2 = \{F : F(s) \textit{ is analytic in } Re(s) > 0 \textit{ and } \|F\|_2 < \infty\},$$

*in which* $\|F\|_2 = \{\sup_{\alpha>0} \frac{1}{2\pi} \int_{-\infty}^{\infty} F^*(\alpha + j\omega)F(\alpha + j\omega)d\omega\}^{\frac{1}{2}}$.

**Theorem 4.** *Suppose $\Lambda u$ is the Laplace transform of function $u$ in time domain, then:*

*(a) If $u \in L_2[0,\infty)$, then $\Lambda u \in \mathcal{H}_2$.*

*(b) If $\hat{u} \in \mathcal{H}_2$, there exists $u \in L_2[0,\infty)$ satisfying $\Lambda u = \hat{u}$.*

Theorem 4 states that $\mathcal{H}_2$-space is isomorphic to $L_2[0,\infty)$-space. In other words, a system in time domain maintains the input-output mapping $L_2[0,\infty) \to L_2[0,\infty)$, if and only if its Laplace transform maintains the mapping $\mathcal{H}_2 \to \mathcal{H}_2$. This class of systems is known as $\mathcal{H}_\infty$ which is defined as:

**Definition 4.** *($\mathcal{H}_\infty$-space) The Hardy $\infty$-space is defined as*

$$\mathcal{H}_\infty = \{\mathbf{M} : \mathbf{M}(s) \textit{ is analytic in } Re(s) > 0 \textit{ and } \|\mathbf{M}\|_\infty < \infty\},$$

*in which* $\|\mathbf{M}\|_\infty = \sup_{\alpha>0}\{\sup_{\omega} \overline{\sigma}\{\mathbf{M}(\alpha + j\omega)\}\}$.

This implies us to perform the proposed admittance splitting technique with the re-partitioned $\mathbf{G}$ and $\mathbf{H}$ blocks in $\mathcal{H}_\infty$-space.

## 2.5 Theoretically Rigorous Admittance Splitting Scheme

### 2.5.1 Specific Constraints for Practical Admittance Splitting

According to the discussion in Section 2.4, the repartitioned $\mathbf{G}$ and $\mathbf{H}$ blocks must lie in $\mathcal{H}_\infty$ space while performing admittance splitting. Practical systems such as PDNs considered in this work have rational transfer functions. We further construct new $\mathbf{G}(s)$ and $\mathbf{H}(s)$ using rational splitting factors $S_{g2h}(s)$ and $S_{h2g}(s)$:

$$\mathbf{G}(s) = \mathbf{G}_0^Y(s) - S_{g2h}(s)\widehat{\mathbf{G}}_0^Y(s) + S_{h2g}(s)\widehat{\mathbf{H}}_0^Y(s) \tag{2.19}$$

$$\mathbf{H}(s) = (\mathbf{H}_0^Y(s) + S_{g2h}(s)\widehat{\mathbf{G}}_0^Y(s) - S_{h2g}(s)\widehat{\mathbf{H}}_0^Y(s))^{-1} \tag{2.20}$$

Clearly, a rational transfer function is in $\mathcal{H}_\infty$ space iff it has no RHP pole and is proper with the relative degree (defined as the difference between the degree of the numerator and the degree of the denominator) less than or equal to 0. As the blocks in our practical power delivery network before repartitioning are stable proper systems with a relative degree of 0 for each entry in the matrix transfer function, $\widehat{\mathbf{G}}_0^Y(s)$ and $\widehat{\mathbf{H}}_0^Y(s)$ are all in $\mathcal{H}_\infty$ space. Based on that, we summarize the following two important constraints for practical admittance splitting: 1) the splitting factors $S_{g2h}(s)$ and $S_{h2g}(s)$ shall be well-behaving in $\mathcal{H}_\infty$ space to ensure that $\mathbf{G}(s)$ and $\mathbf{H}^{-1}(s)$ are in $\mathcal{H}_\infty$ space as well; 2) the matrix inversion in (2.20) shall not generate unstable RHP poles to ensure that $\mathbf{H}(s)$ is also in $\mathcal{H}_\infty$. In the remainder of this section, a theoretically rigorous admittance splitting scheme which satisfies the above two constraints is presented.

### 2.5.2 Constructing Appropriate Splitting Factors

#### *2.5.2.1 Naive Ideal Admittance Splitting*

Recall that we partition the frequency axis into two disjoint sets of $\Omega_{gain}$ and $\Omega_{pass}$, including the frequency points at which the direction of the admittance splitting is from $\mathbf{G}$ to $\mathbf{H}$ and the opposition direction, respectively. To maximize frequency-dependent admittance splitting, it might appear to be appropriate to determine the best splitting strategy at each frequency completely independent of other frequencies. We call this approach the *ideal* splitting.

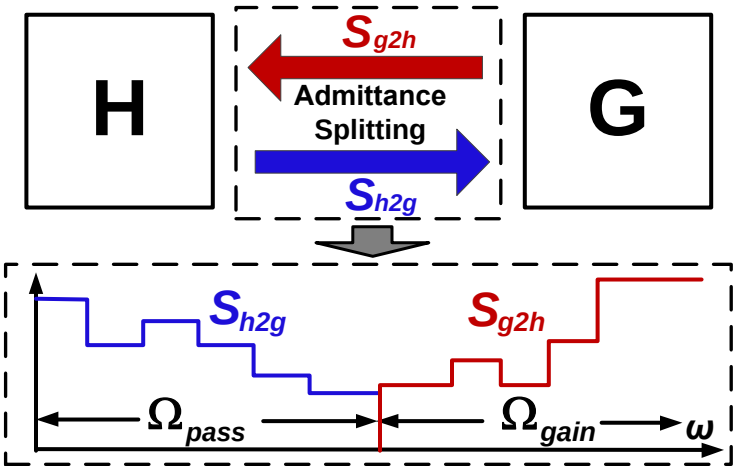


Figure 2.6: Naive ideal admittance splitting.

However, the splitting factors $S_{g2h}$ and $S_{h2g}$ resulted from the ideal splitting can have an arbitrary number of jumps and become discontinuous (Fig. 2.6). Although it seems to be beneficial as it provides the maximum freedom in admittance splitting, the corresponding $S_{h2g}$ and $S_{g2h}$ can have singularities leading to the violation of $\mathcal{H}_\infty$ requirement. Clearly, this choice and many other artificial admittance splitting will destroy the key system properties discussed in Section 2.4.

#### *2.5.2.2 Proposed Smooth Admittance Splitting*

To employ HST with theoretical rigor, we propose a smooth admittance splitting scheme in consistent to the first constraint in Section 2.5.1. Our key idea is to use a series of N-th order

practical filters $\{A_1(j\omega), A_2(j\omega), ..., A_n(j\omega)\}$, which are rational and well-behaving, to construct $S_{h2g}(j\omega)$ and $S_{g2h}(j\omega)$. As such, discontinuities as shown in Fig. 2.6 would be eliminated.

The practical implementation of this idea entails making the first filter as a low-pass filter, the last one a high-pass filter, and all other filters band-pass. While it is possible to use other types of filters, Butterworth filters are a good choice. The N-order Butterworth low-pass filter is given by

$$A_1(s) = \frac{M_1}{\prod_{k=1}^{N}((s - s_k)/\omega_c)}, \tag{2.21}$$

where $M_1$ is the magnitude and $s_k = \omega_c e^{(j(2k+N-1)\pi/2N)}$. Note that $Re\{s_k\} < 0$ and $A_1(s) \in \mathcal{H}_\infty$. Butterworth bandpass and highpass filters are also in $\mathcal{H}_\infty$. Now $S_{h2g}$ and $S_{g2h}$ are constructed as follows:

$$S_{h2g}(j\omega) = \sum_{\Omega_i \in \Omega_{pass}} A_i(j\omega), \tag{2.22}$$

$$S_{g2h}(j\omega) = \sum_{\Omega_i \in \Omega_{gain}} A_i(j\omega), \tag{2.23}$$

where $\Omega_i$ is the pass band of the $i$-th filter. The filters are assigned to two groups to construct $S_{h2g}(j\omega)$ and $S_{g2h}(j\omega)$, respectively. Clearly, $S_{h2g}(s)$ and $S_{g2h}(s)$ are in $\mathcal{H}_\infty$ space as they are superpositions of the Butterworth filters. Fig. 2.7 gives a illustration of the filter-based splitting technique where a single admittance are distributed into the $\mathbf{G}$ and $\mathbf{H}$ blocks according to filter transfer functions. Obviously, the well-behaved splitting factors will not introduce any RHP singularity.

Compared to the ideal splitting, finite roll-offs and phase shifts are introduced by the filters adopted in the proposed approach. It seems that high-order filters are more desirable since the reduced roll-off can make the admittance splitting more flexible with small granularity. However, since the total phase shift for an $N$-th order filter is $N \times 90°$, it inevitably introduces large phase shift and may jeopardize the splitting for expected direction around the cutting-off frequency. For example, for the frequencies with phase shift equals to $-180°$, the splitting direction is completely inversed. On the other hand, experimental studies show that small splitting granularities are un-

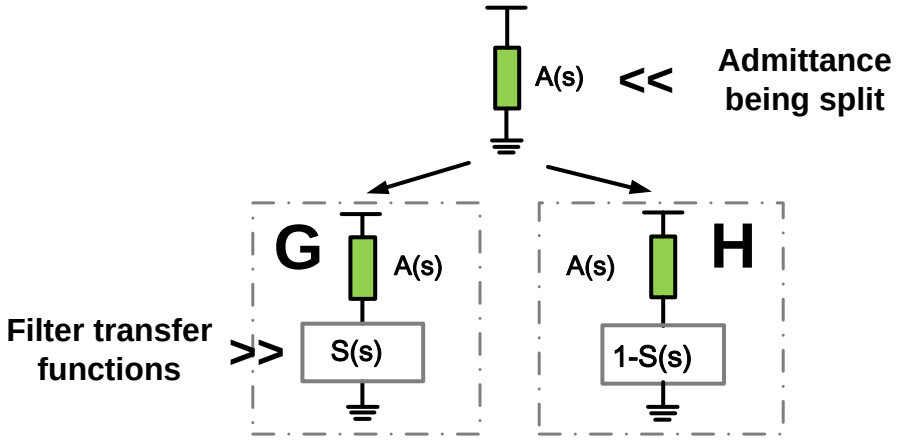


Figure 2.7: A single admittance re-partitioned via smooth admittance splitting.

necessarily needed and therefore low-order filters are used in our experiments to control the phase shifts. In the next section, we present an automated design flow which optimizes the gains and the cutoff frequencies (i.e. placements) of the filters to optimally realize $S_{h2g}$ and $S_{g2h}$ over the frequency axis so as to minimize the conservativeness in stability evaluation.

### 2.5.3 Stability Checking for H Block

With the help of the proposed filter-based admittance splitting technique, the admittance matrices of the re-partitioned $\mathbf{G}$ and $\mathbf{H}$ blocks have no RHP singularities. However, in order to relate to the negative feedback interconnection in Fig. 2.2(a), we need to invert $\mathbf{H}$'s admittance matrix to get the impedance representation. This may create unstable RHP poles in $\mathbf{H}(s)$ and violate the second constraint in Section 2.5.1. It should be noted that the original $\mathbf{H}$ block comprising of only RLC elements must satisfy this constraint as both of its admittance and impedance matrices are stable. This is based on the fact that a bounded input cannot produce an unbounded branch current or node voltage in a passive RLC circuit. Thus the above issue only exists in the repartitioned $\mathbf{H}$ block with artificial splitting. To maintain the theoretical rigor, we proposed to confirm the stability of $\mathbf{H}(s)$ based on the following generalized (MIMO) Nyquist stability criterion [50][51].

**Theorem 5.** *Let $P_{ol}$ denote the number of unstable poles in the open-loop transfer matrix $\mathbf{L}(s)$.*

*Then the closed-loop negative feedback system* $(\mathbf{I}+\mathbf{L}(s))^{-1}$ *is stable, if and only if the Nyquist plot of det*$(\mathbf{I}+\mathbf{L}(s))$*:*

*1) makes* $P_{ol}$ *anti-clockwise encirclements of the origin, and*

*2) does not pass through the origin.*

Different from the single-input-single-output (SISO) case, the net encirclement of the origin in Theorem 5 is evaluated via the determinant of the matrix $(\mathbf{I}+\mathbf{L}(s))$ as $s$ traverse the Nyquist contour shown in Fig. 2.8(a). In our case we have

$$\mathbf{I}+\mathbf{L}(s)=\mathbf{H}^{-1}(s)=\mathbf{H}_0^Y(s)+S_{g2h}(s)\widehat{\mathbf{G}}_0^Y(s)-S_{h2g}(s)\widehat{\mathbf{H}}_0^Y(s). \tag{2.24}$$

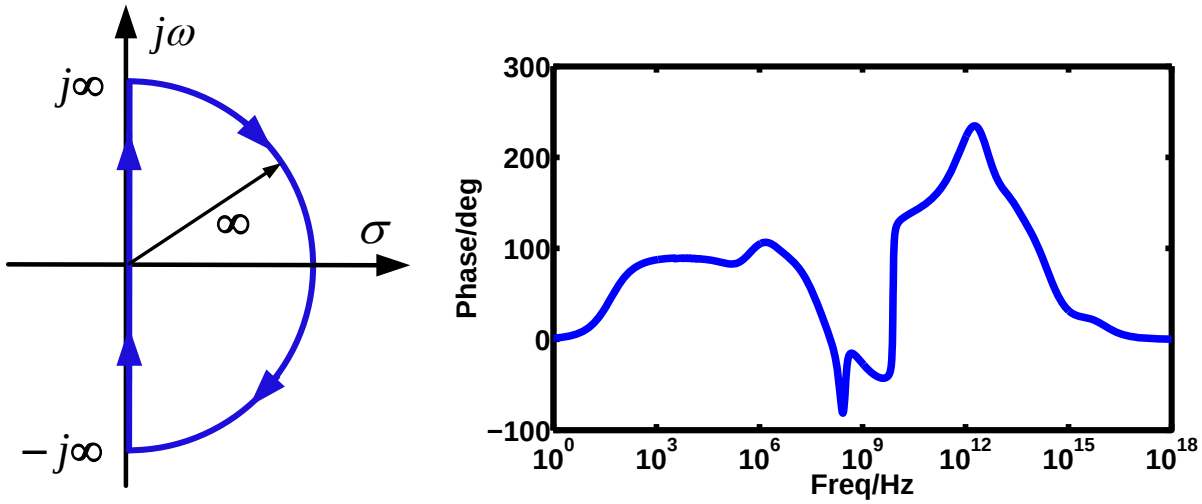


(a) (b)

Figure 2.8: (a) Nyquist contour, (b) phase shift in the bode plot of det$(\mathbf{H}^{-1}(j\omega))$ for a realistic design.

Then the encirclement can be obtained by examining $[\Delta\arg\ \det(\mathbf{H}^{-1}(s))]/2\pi$, where $\Delta\arg$ denotes the total change in argument as $s$ traverses the Nyquist contour. As $\mathbf{H}^{-1}(s)$ is proper, the phase shift of det$(\mathbf{H}^{-1}(s))$ is zero when $s$ traverses the semicircle with an infinite radius. Therefore the total phase shift can be conveniently evaluated in the frequency domain, i.e. by tracking the

phase curve in the Bode plot of $\det(\mathbf{H}^{-1}(j\omega))$. A realistic example in shown in Fig. 2.8(b). As each term in the right side of (2.24) is in $\mathcal{H}_\infty$ space which gives $P_{ol} = 0$, $\mathbf{H}(s)$ has no RHP poles iff the number of the encirclements (or the total phase shift in Bode plot) is zero.

## 2.6 HST-Based PDN Design Flow

This section illustrates a bottom-level HSM optimization flow which is used to maximize the HSM for a more truthful evaluation of the system-wide stability. After that, an efficient top-level automated design flow which incorporates the optimal HSM evaluator is presented for the PDN design optimization.

### 2.6.1 Optimal HSM Evaluator

The HSM optimization flow solves an optimal frequency-dependent admittance splitting problem to maximize the network-wide HSM corresponding to the objective function

$$f = HSM(\mathbf{M}, \boldsymbol{\omega}), \tag{2.25}$$

where $\mathbf{M} = \{M_1, M_2, ..., M_n\}$ defines the gain of each filter and $\boldsymbol{\omega} = \{\omega_1, \omega_2, ..., \omega_{n-1}\}$ specifies the cut-off frequencies. The optimization problem is subject to two constraints:

1. Filter magnitude constraint:

$$-1 \leq M_1, M_2, ..., M_n \leq 1. \tag{2.26}$$

The $i$-th filter belongs to $S_{g2h}$ when $M_i > 0$; otherwise it belongs to $S_{h2g}$. This supports bidirectional splitting.

2. Filter frequency constraint:

$$\omega_L < \omega_1 < \omega_2 < ... < \omega_{n-1} < \omega_H, \tag{2.27}$$

where $\omega_L$ and $\omega_H$ are the lowest and highest frequency between which admittance splitting takes

place.

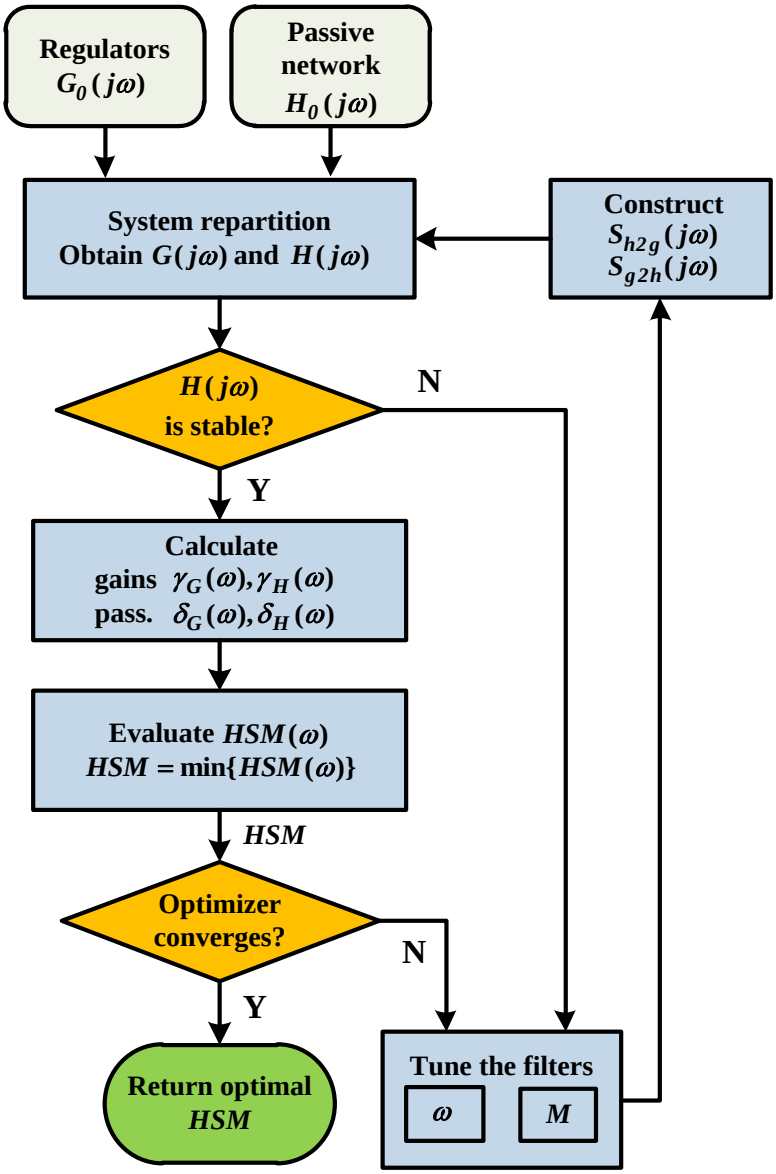


Figure 2.9: Flowchart of the optimal HSM evaluation for a given PDN design.

The optimal HSM evaluation flow is illustrated in Fig. 2.9. The transfer matrices $\mathbf{G}_0(j\omega)$ and $\mathbf{H}_0(j\omega)$ are precomputed for the original regulator block and passive network. In each iteration, the two splitting functions $S_{h2g}(j\omega)$ and $S_{g2h}(j\omega)$ are constructed based on the current magnitudes $\mathbf{M}$ and cutoff frequencies $\boldsymbol{\omega}$ of the filters. Then, the system is re-partitioned into the new $\mathbf{G}$

and $\mathbf{H}$ blocks according to (2.19) and (2.20). It is followed by a stability checking for the $\mathbf{H}$ block based on the MIMO Nyquist stability theory as in Section 2.5.3. Our experimental study shows that the re-partitioned $\mathbf{H}$ block doesn't contain RHP poles in most cases. After that, the gain and passivity information for each block are characterized and the temporary HSM for the entire network is calculated. The filter parameters $\mathbf{M}$ and $\boldsymbol{\omega}$ are adjusted in each iteration till the optimizer converges, outputting the optimal HSM. Note that this flow invokes no circuit simulation but only matrix operations for the system repartition and HSM evaluation. Since the size of the matrices is only $2n \times 2n$ where $n$ is the number of LDOs in PDN, this optimal HSM evaluator is implemented with high efficiency.

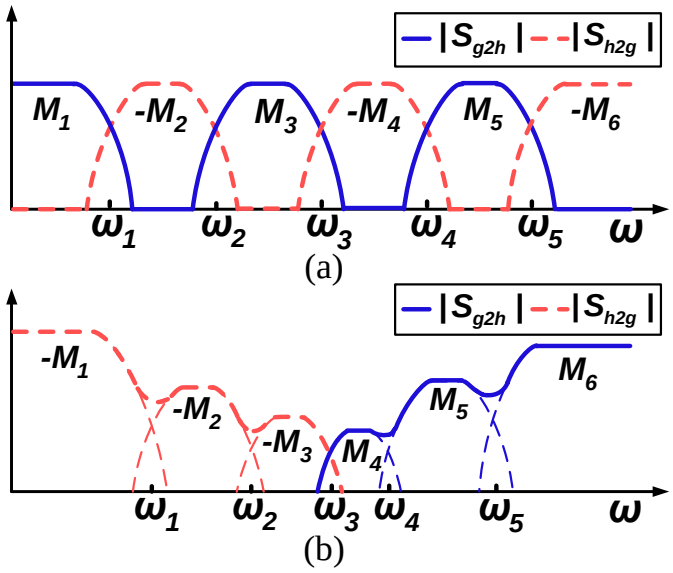


Figure 2.10: Frequency-dependent admittance splitting: (a) before optimization, (b) after optimization.

Fig. 2.10 gives an intuitive picture of the above optimization process. Initially, all filters are evenly placed along the frequency axis (Fig. 2.10(a)). The optimizer adjusts the magnitudes and cut-off frequencies of the filters to realize the optimal bidirectional admittance splitting. A typical optimal splitting is shown in Fig. 2.10(b), where the splitting from the $\mathbf{H}$ to the $\mathbf{G}$ block happens at low frequencies while the opposite splitting direction takes place at high frequencies.

### 2.6.2 Automated PDN Design Flow

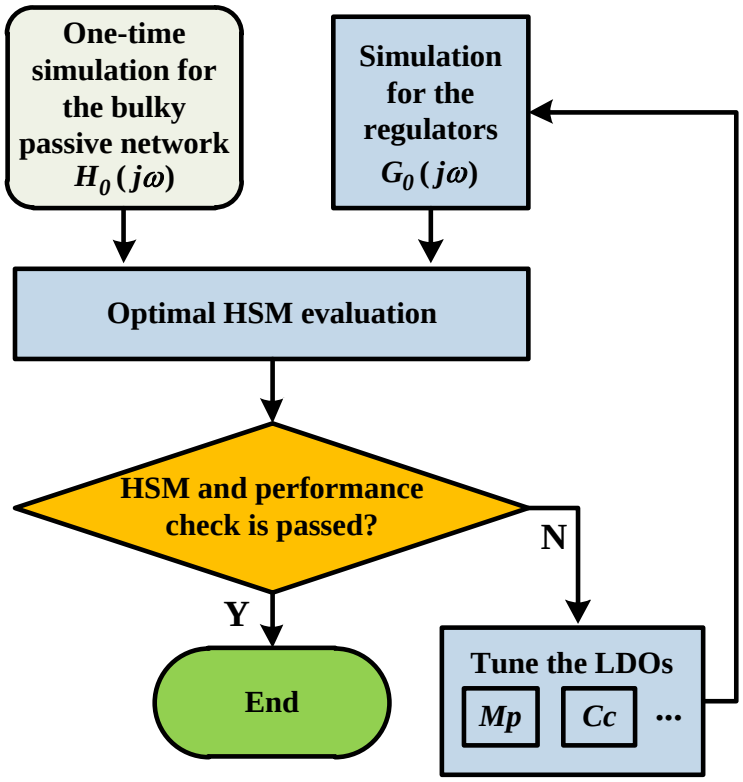


Figure 2.11: Automated flow of the PDN design optimization. The optimal HSM evaluator flow in Fig. 2.9 is adopted.

We present an efficient automated design flow for PDNs with distributed LDOs. As shown in Fig. 2.11, it employs the optimal HSM evaluator in 2.6.1 to evaluate network stability. In each iteration, it adjusts key LDO design parameters (e.g. the compensation capacitor $Cc$ and transistor sizing) to jointly optimize a given set of specifications such as HSM, regulation performance, power efficiency, and area overhead.

Appealingly, the flow is local to the design of LDOs in the sense that after one-time AC (admittance) characterization of the large passive network, it tunes each LDO *locally* while ensuring the stability of the entire PDN. As only the low-cost LDO circuits simulation is repeated in each

iteration, the entire PDN optimization flow is computationally efficient.

## 2.7 Experimental Studies

The automated design flow has been implemented using C/C++ and the APPSPACK 5.0.1 optimization package [52]. The LDO topology from [1] as shown in Fig. 2.12 is adopted to design a number of LDOs based on a commercial 90nm CMOS technology. The adopted package model is similar to [53] as illustrated in Fig. 2.13 and the parameters are shown in Table 2.2. As in [53], the parasitic values are chosen to match the measured off-chip impedance of the Pentium 4 processor and the used system in the experiment is reasonable to model modern IC designs. HSPICE is employed for AC characterization of each circuit block using 200 samples per decade. To provide a reference, we implement a simpler HSM checking approach which splits all self-admittances from the $\mathbf{G}$ block to the $\mathbf{H}$ block to lower the PDN loop gain in the frequency range where gain condition is targeted. A similar approach is employed in [22]. This scheme is a restricted version of our presented work with single direction admittance splitting, offering a meaningful reference of comparison for the proposed optimal bidirectional frequency-dependent admittance splitting scheme. Since the stability must be maintained under all load conditions and the worst case stability for each LDO usually happens at light load condition, we use the light load as the operating point to characterize stability and HSM.

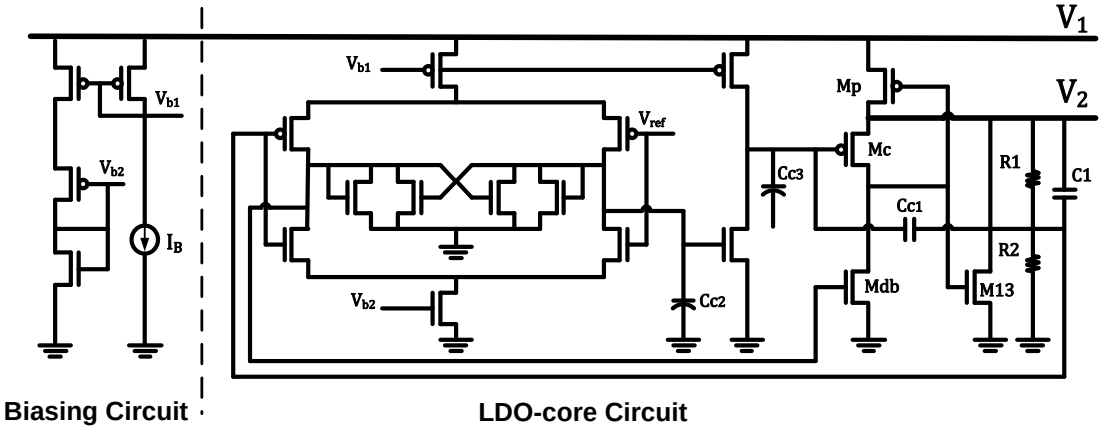


Figure 2.12: A multi-loop LDO [1] adopted in the experiment setups.

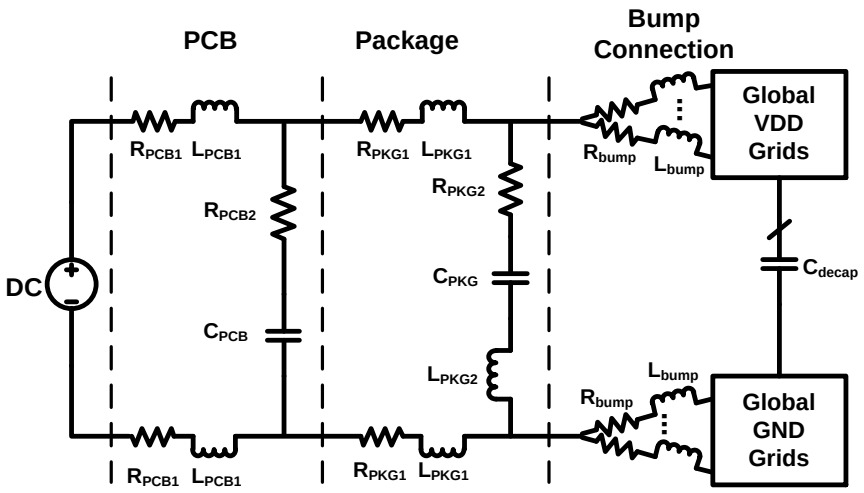


Figure 2.13: Package model.

Table 2.2: Parameters for the package model.

| Resistance | Value | Inductance | Value | Capacitance | Value |
|---|---|---|---|---|---|
| $R_{PCB1}$ | 0.094 mohm | $L_{PCB1}$ | 21 pH | $C_{PCB}$ | 240 uF |
| $R_{PCB2}$ | 0.167 mohm | | | | |
| $R_{PKG1}$ | 1 mohm | $L_{PKG1}$ | 120 pH | $C_{PKG}$ | 26 uF |
| $R_{PKG2}$ | 0.542 mohm | $L_{PKG2}$ | 5.61 pH | | |
| $R_{bump}$ | 0.4 mohm | $L_{bump}$ | 120 pH | | |

### 2.7.1 Pessimism Reduction in Stability Analysis

Recall that the sufficient nature of HST can lead to conservativeness in stability evaluation. To demonstrate, we start with a small PDN of about 200 nodes and 4 LDOs, such that the stability can be feasibly analyzed using the pole analysis as the ground truth. We examine the PDN stability by performing time-domain transient analysis (Fig. 2.14(a)). We assume the load current modeled with current sources are evenly distributed across the regulated grids. It is observed that even under a large injected load current ($IL$) variation stepping between 0.36A and 0.44A with 0.1ns transition time, the voltage responses of the global VDD grids ($GV_{DD}$) and the regulated local

VDD grids ($V_{REG}$) get settled quickly. The stability of the network is more rigorously confirmed by a full-blown pole analysis which shows no RHP poles.

To see how the HSM based stability evaluation works, we first apply the reference approach described at the beginning of this section and show its results in Fig. 2.14(b). The frequency axis can be broken into four sub-regions: region "A" with only a positive $PassMarg$, region "B" with only a positive $GainMarg$, region "C" where both margins are non-positive, and region "D" with both $PassMarg$ and $GainMarg$ being positive. This frequency labeling convention is used throughout the work. The existence of the "C" region in Fig. 2.14(b) might suggest "potential" instability of the system, or at least prevents us from drawing any firm conclusion about stability. The fact that the system is stable clearly shows the conservativeness of the reference method.

The same PDN is re-checked by our proposed optimal HSM evaluator in Fig. 2.14(c), where the optimized $S_{g2h}$ and $S_{h2g}$ operate in a complimentary way and region "C" is eliminated, formally confirming the stability and demonstrating the effectiveness of the proposed technique.

### 2.7.2 Stability-Ensuring PDN/LDO Design

The traditional phase and gain margins can lead to misleading conclusions about stability since they are only applicable to single-loop stability analysis. To see this, we plug 4 LDOs depicted in Fig. 2.12 into a small PDN with around 200 nodes. By adopting the active frequency compensation scheme in the LDO topology, each LDO achieves a large phase margin of over 100°. The transient simulation of Fig. 2.15(a), however, indicates that integrating these seemingly stable LDOs into the PDN actually creates instability, which is verified by a pole analysis that detects a pair of complex RHP poles. The proposed stability evaluator can detect this instability by identifying a potential unstable region (region "C") that is between 180KHz and 27MHz as shown in Fig. 2.15(b).

We now apply our automated design flow to tune this LDO design locally to fix the detected instability by achieving a positive HSM. The frequency-wise checking of the system is shown in Fig. 2.15(c), where positive local HSM is achieved for all frequencies. As shown in Fig. 2.15(d), the transient simulation of the optimized system confirms that the instability is indeed fixed.

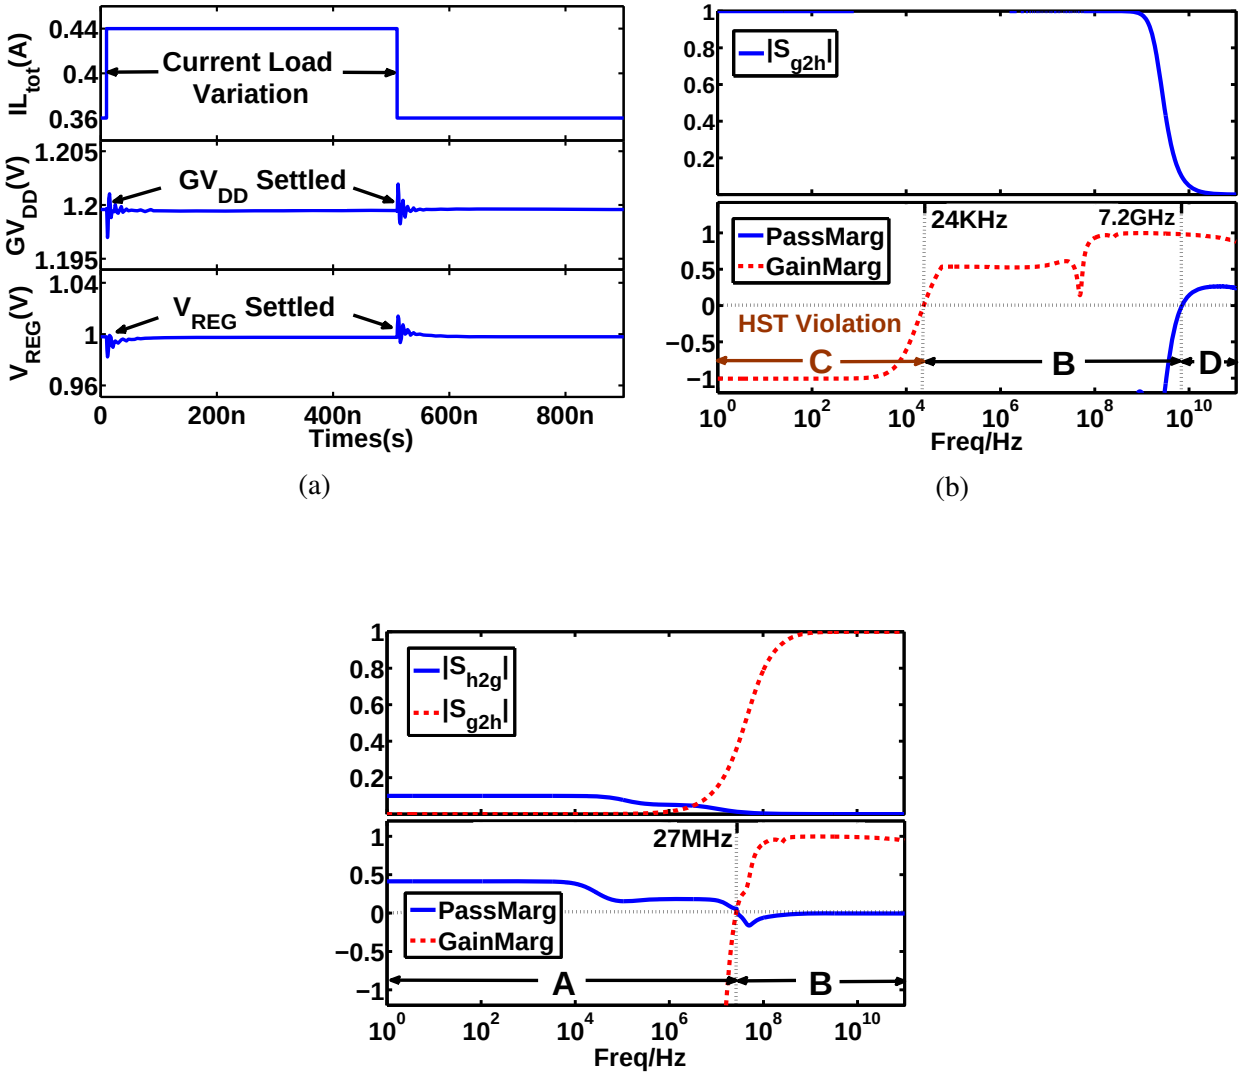


Figure 2.14: HSM checking of a stable PDN: (a) transient simulation, (b) frequency-wise stability checking using the reference method, and (c) frequency-wise stability checking using the proposed evaluation.

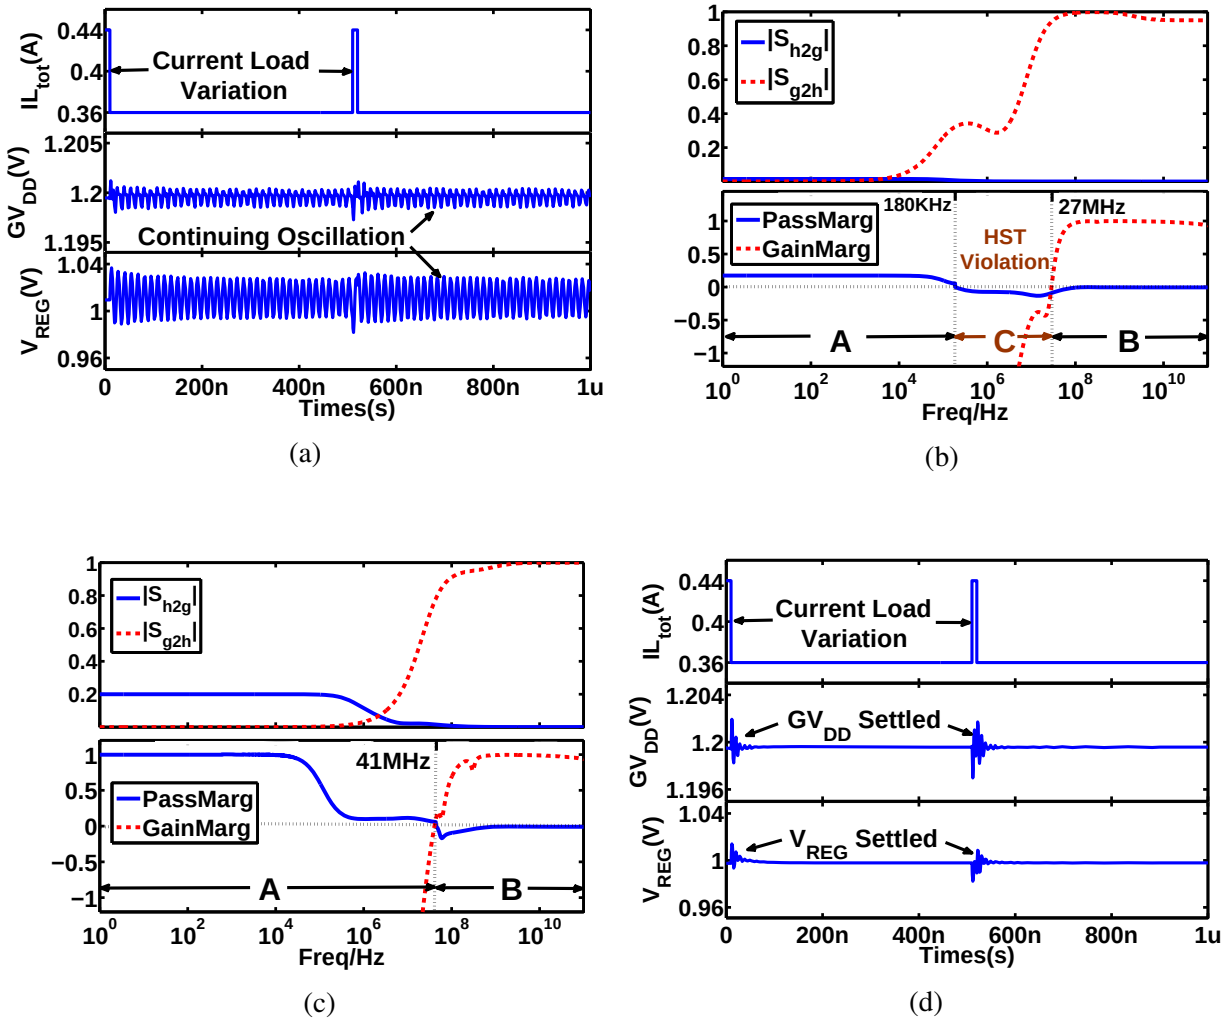


Figure 2.15: Stability-ensuring PDN design: (a) transient analysis showing the instability of a PDN integrating 4 LDOs with a large phase margin, (b)the proposed frequency-wise stability checking confirming the in-stability, (c) frequency-wise stability checking with the optimized stable design, and (d) transient analysis of the optimized LDO design confirming the fixed instability.

### 2.7.3 Joint Performance Optimization

Our automated design flow allows the designer to jointly optimize competing design specifications while ensuring the network's stability. To demonstrate this, we consider several key LDO specifications: DC accuracy of load regulation $ACC$, unit-gain bandwidth $UGB$, quiescent current $I_q$, average output admittance $Y_{avg}$ (representing the average load regulation performance within the frequency range of interest), and $Y_{avg}/I_q$, a measure of design efficiency, in addition to the PDN's global HSM. We consider two PDN examples with the first PDN (PDN1) having about 200 nodes and 4 integrated LDOs, and the second PDN (PDN2) having over 100K nodes and 9 LDOs. Initially, the LDO is designed to achieve a large phase margin under a lumped capacitive load using the standard analog design process.

We apply our proposed automated design flow to tune this LDO design to ensure network stability while optimizing all other performances. We also replace our optimal HSM evaluator in the design flow by the reference HSM checking method to come up a reference design flow for comparison purposes. The performances based on the initial LDO design ($LDO_1$), the reference approach ($LDO_2$), and the proposed approach ($LDO_3$) are reported in Table 2.3 and Table 2.4. The last column of each table shows the performance boosts of the proposed approach over the reference approach.

Table 2.3: Performance comparisons for PDN1.

| | Initial $LDO_1$ | Opt. $LDO_2$ | Opt. $LDO_3$ | Perform. Gain |
|---|---|---|---|---|
| HSM | -2.47[1]/-0.23[2] | 6.69e-3[1] | 5.77e-2[2] | – |
| Reg. ACC. | 99.96% | 99.95% | 99.92% | $\downarrow 0.03\%$ |
| UGB(MHz) | 380 | 501 | 562 | $\uparrow 12.2\%$ |
| $I_q(\mu A)$ | 639 | 320 | 313 | $\downarrow 1.9\%$ |
| $y_{avg}(S)$ | 5.44 | 1.22 | 2.61 | $\uparrow 114\%$ |
| $y_{avg}/I_q(S/\mu A)$ | 8.51e-3 | 3.81e-3 | 8.34e-3 | $\uparrow 119\%$ |

[1]Evaluated by the reference approach.

[2]Evaluated by the optimal HSM evaluator.

Table 2.4: Performance comparisons for PDN2.

| | Initial $LDO_1$ | Opt. $LDO_2$ | Opt. $LDO_3$ | Perform. Gain |
|---|---|---|---|---|
| HSM | -1.80[1]/-1.69[2] | 1.17e-2[1] | 0.58[2] | – |
| Reg. ACC. | 99.96% | 99.96% | 99.94% | $\downarrow 0.03\%$ |
| UGB(MHz) | 380 | 531 | 661 | $\uparrow 24.5\%$ |
| $I_q(\mu A)$ | 639 | 353 | 352 | $\downarrow 0.3\%$ |
| $y_{avg}(S)$ | 5.44 | 2.00 | 2.83 | $\uparrow 41.5\%$ |
| $y_{avg}/I_q(S/\mu A)$ | 8.51e-3 | 5.67e-3 | 8.05e-3 | $\uparrow 42\%$ |

[1]Evaluated by the reference approach.

[2]Evaluated by the optimal HSM evaluator.

According to Table 2.3 and Table 2.4, both design flows can fix the instability of the initial LDO design by providing a positive HSM margin. However, the proposed design flow outperforms the reference flow noticeably, boosting the performance by up to 119% and 42% in PDN1 and PDN2, respectively. It suggests that the over-design with the proposed design approach is significantly reduced.

We use a PDN with 100K nodes and 4 LDOs to extensively demonstrate the benefits brought by the proposed over-design reduction. By weighting the six design specifications differently in the objective function, we generate 25 LDO designs with a positive HSM using the reference design flow and the proposed flow, respectively. After AC characterization of the passive network, it takes about $9.7$ hours to complete tuning each LDO design using 4 threads on a 2.3 GHz AMD Opteron processor. To have a more truthful evaluation of stability, the HSMs of the designs produced by the reference flow are re-evaluated using the optimal HSM evaluator at the conclusion of the optimization process.

To systematically evaluate the design tradeoffs achieved by the two flows, we define two figures of merit (FOMs): $FOM1 = \frac{UGB}{UGB\_max} \cdot \frac{Y_{avg}}{Y_{avg\_max}} \cdot 10^{HSM} \cdot ACC$ and $FOM2 = \frac{UGB}{UGB\_max} \cdot \frac{Y_{avg}}{Y_{avg\_max}} \cdot ACC$, where $UGB\_max$ and $Y_{avg\_max}$ are the maximum $UGB$ and $Y_{avg}$ values among these design points. Clearly, FOM2 involves important regulation performances while FOM1 considers the overall performance including the HSM. Fig. 2.16(a) plots the FOM1 as a function of

$I_q$, showing the proposed work can significantly boost FOM1 by an average 113% with a quiescent current saving by 11%. Then Fig. 2.16(b) visualizes the tradeoff between the regulation performances (FOM2) and HSM, illustrating that the proposed approach improves the FOM2 and HSM on average by 29% and 91%, respectively. Besides, the 3-D plot in Fig. 2.17 demonstrates the tradeoffs between $I_q$, $UGB$ and $Y_{avg}$ achieved by the two flows, demonstrating the improved design tradeoffs within a large design space.

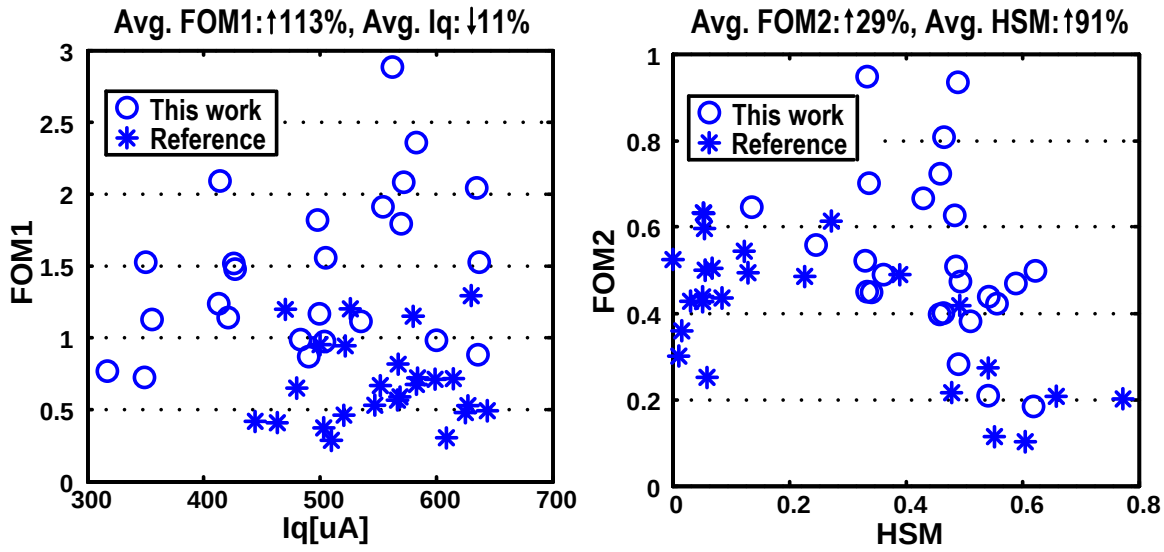


(a) (b)

Figure 2.16: Comparison of optimized design tradeoffs: (a) FOM1 as a function of the quiescent current, and (b) tradeoff between HSM and FOM2.

Additionally, the bar charts in Fig. 2.18 give a full picture of the design tradeoffs within the large design space. As each design achieves a similar high regulation accuracy, we only compare $\frac{HSM}{HSM\_max}$, $\frac{UGB}{UGB\_max}$ and $\frac{Y_{avg}/I_q}{(Y_{avg}/I_q)\_max}$ of these designs, where $HSM\_max$, $UGB\_max$, and $(Y_{avg}/I_q)\_max$ are the maximum $HSM$, $UGB$, and $Y_{avg}/I_q$ values among all 50 optimized designs. It is observed that the proposed work brings forth a 35% improvement on the average overall performance (the average height for each bar). Evidently, the proposed approach produces significantly improved design tradeoffs.

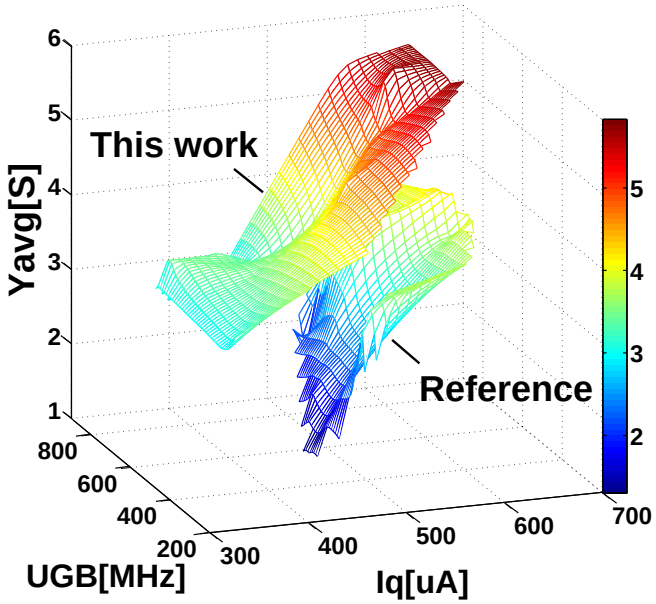


Figure 2.17: Optimized design tradeoffs between three performances.

We also compare the classic phase margin (PM) based design approach with the HST-based approaches. Three groups of optimized PDN designs are generated based on different design methodologies. The first group is optimized using a phase margin based optimization flow. Different from the HST-based flow as in Fig. 2.11, the phase margin of individual LDO is used as the stability metric instead of HSM in the optimization process. After optimization, each individual LDO achieves a large phase margin of over $100°$. The second and the third groups obtain positive HSMs based on the reference and the proposed approach respectively. Different $I_q$ constraints are placed to produce multiple design points. As in Fig. 2.19(a), while the large over-design in the reference approach results in performance degradation, the proposed approach improves FOM2 by up to 53% when compared with the PM-based approach. In the meantime, even though achieving large phase margins in the first group, Fig. 2.19(b) shows that the HST violation may happen among some of the design points which can lead to potential instability.

We end our discussion in this section with the application of the proposed design flow to PDNs with multiple domains of different size. We consider a 100K-node PDN with the first domain integrating 4 LDOs and the second domain integrating 2 LDOs. Through locally tuning the LDOs

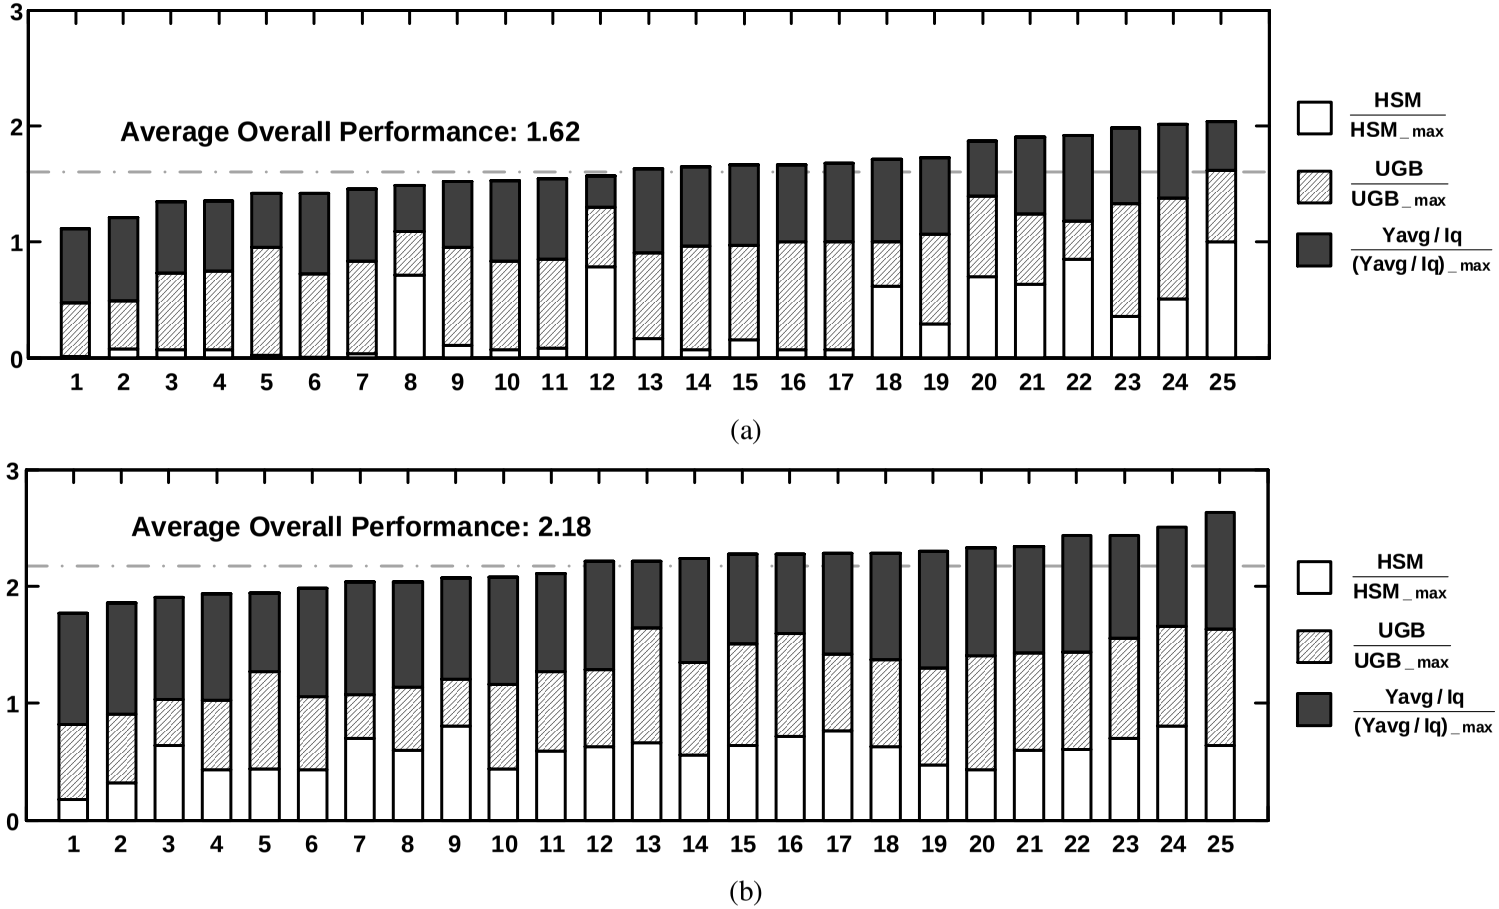


Figure 2.18: Design comparison within a large space: (a) 25 optimal designs with the reference approach, and (b) 25 optimal designs with the proposed approach.

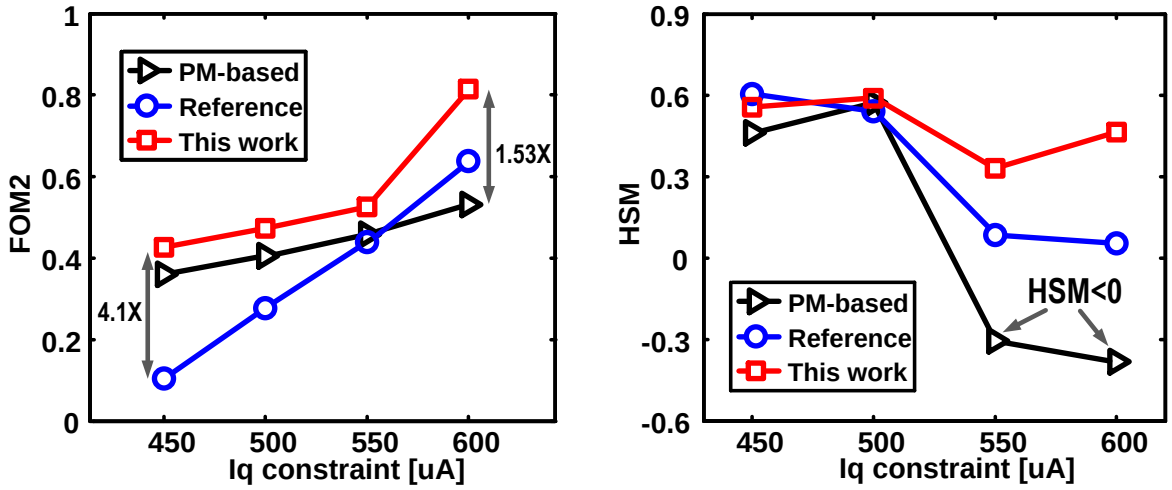


(a) (b)

Figure 2.19: PM-based approach versus HST-based approaches: (a) FOM2 comparison, and (b) HSM comparison.

in each domain, our design flow can jointly optimize the two domains with different performance specifications. As illustrated in Table 2.5, the LDOs in the first domain is optimized by emphasizing high regulation performance, while the LDOs in the second domain is targeting low quiescent power consumption. The achieved positive HSM after optimization indicates guaranteed stability for the entire PDN system.

Table 2.5: Performance optimization for PDN with two domains.

| | LDOs in 1st domain | LDOs in 2nd domain |
|---|---|---|
| HSM | 5.21e-2 | |
| Reg. ACC | 99.96% | 99.96% |
| UGB(MHz) | 562 | 275 |
| $Iq(uA)$ | 467.5 | 397.5 |
| $y_{avg}(S)$ | 3.99 | 3.62 |
| $y_{avg}/I_q(S/uA)$ | 8.53e-3 | 9.11e-3 |

## 2.8 Summary

In this chapter, a bidirectional frequency-dependent system partition technique with a theoretically complete framework has been proposed to reduce the inherent conservativeness of the hybrid stability theorem. After recognizing all the necessary constraints making the partition scheme theoretically rigorous, an efficient optimization-based approach has been presented to assess the stability of large PDNs with distributed regulators. With the reduced pessimism in stability checking, we propose a new automated PDN design flow which can significantly reduce the system overdesign and improve the overall PDN performance while ensuring the stability of the entire network.

# 3. STABILITY-CONSTRAINED DESIGN SPACE EXPLORATION FOR DISTRIBUTED ON-CHIP POWER DELIVERY *

As discussed in Chapter II, a new stability-ensuring PDN design methodology, based on hybrid stability theory (HST), is developed to cope with the complex stability problem of distributed on-chip voltage regulation and enables efficient localized PDN design. As helpful as it is, the new stability margin, combing the system gain and passivity, is unfamiliar to circuit designers and, as a result, the above design approach mostly serves as a block box and cannot offer much intuition to designers. For example, it takes the initial PDN design as the input and after the optimization it outputs a certain PDN satisfying all the specifications on HSM and performances.

In order to offer more design intuitions about the HST-based design of power delivery networks, this chapter systematically explores the relationship between key circuit-level parameters and the hybrid stability metric. We identify several important PDN design parameters which can significantly influence the system-wide stability metric and other design specifications; namely, regulator topology, on-chip passive coupling capacitance and the number of regulators. Being aware of these key parameters and their impact on the overall PDN performance can help IC designers make informed design choices at an early stage. After studying the effect of each design parameter separately, extensive studies are conducted to bring together all the previous cases and explore the circuit design tradeoffs. This chapter is concluded by summarizing all the new design insights obtained from the HST-based design studies. Circuit designers can make use of these meaningful design insights to improve the overall PDN tradeoffs.

## 3.1 PDN Design Study: Tradeoffs between Stability and Performance

### 3.1.1 Important Design Parameters for PDN Tradeoffs

Recall that for the application of HST, the entire PDN system is required to be partitioned into a negative feedback interconnection of two blocks. The voltage regulators are contained in the $\mathbf{G}$ block while the remaining passive elements including power grids and PCB/package parasitics are pushed into the $\mathbf{H}$ block. We first identify several important design parameters within the large PDN design space which can not only impact the design performances but also influence the characteristics of $\mathbf{G}$ and $\mathbf{H}$ blocks thus hybrid stability of the entire system.

#### *3.1.1.1 LDO Design Parameters*

For a specific LDO topology, important transistor-level parameters can be adjusted to ensure stability and optimized for performance. For example, through transistor sizing and tuning compensation capacitors, the unity-gain bandwidth (UGB) of LDO can be effectively controlled. This UGB change will in turn affect $\mathbf{G}$'s admittance transfer matrix $\mathbf{G}(j\omega)$ and further alter $\gamma_{\mathbf{G}}(\omega)$ and $\delta_{\mathbf{G}}(\omega)$. To be more specific, the elements in $\mathbf{G}(j\omega)$ have small magnitudes after UGB and this will be reflected on $\gamma_{\mathbf{G}}(\omega)$ as it is in a sense of admittance. Besides, our experiment results show that the $\mathbf{G}$ block usually exhibits a positive $\delta_{\mathbf{G}}(\omega)$ in a high frequency band above UGB. Therefore, this category of circuit modification will not only affect the regulation performance of each LDO, but also alter the hybrid stability of the whole system.

#### *3.1.1.2 LDO Topology*

Another important degree of freedom in the design comes from the choice of LDO topology. From the hybrid stability point of view, $\mathbf{G}$ blocks with different LDO topologies exhibit different characteristics in $\gamma_{\mathbf{G}}(\omega)$ and $\delta_{\mathbf{G}}(\omega)$ and thus constrain the design process in different ways. On the other hand, a judiciously chosen LDO topology facilitates stability-ensuring PDN design while meeting all other specifications in power and area overhead.

#### *3.1.1.3 On-Chip Decaps*

As well known, inserting on-chip decoupling capacitances is beneficial to suppress the high-frequency switching noise as well as the mid/high-frequency impedance peaking due to off-chip inductive parasitics [54]. From the HST perspective, the suppressed resonance peaking will be reflected on the gain of the $\mathbf{H}$ block since $\mathbf{H}(j\omega)$ is an impedance matrix. This gives another important degree of freedom for the design from the side of passives.

#### *3.1.1.4 Number of LDOs*

Having more LDOs would make them a stronger regulation system as they can share work loads to respond to sudden changes in load current requirements. However, recall that the stability issue comes from the inter-LDO interactions. Varying the number of on-chip LDOs alters the inter-LDO coupling strength and therefore may affect the system stability.

In the rest of this chapter, a comprehensive set of design studies will be conducted to reveal the specific impacts of the above design factors on the overall PDN design tradeoffs. New design insights coming from those experiments will be discussed.

## 3.2 Experimental Results and Analysis

The HST-based PDN optimization flow [38, 39] is adopted to conduct the PDN design study. It is implemented using C/C++ and the APPSPACK 5.0.1 optimization package [52]. A standard 90nm CMOS technology is used to design a number of LDOs. The same PCB/package model in [53] is adopted. Since the stability must be maintained under all load conditions and the worst-case stability for each LDO usually happens at light load condition, we use the light load as the operating point to characterize stability and HSM.

### 3.2.1 Impact of LDO Topology

It is well understood how the LDO topology with certain gain and phase characteristics impact the gain/phase margin of a single LDO circuit. However, the gain/phase margin based approach is incapable to predict or analyze the system-level PDN stability in terms of LDO topology. By virtue

of the HST-based PDN optimization flow, we gain the following new design insight: integrating multiple LDOs with high loop gain and unit-gain bandwidth (UGB) may lead to deteriorated hybrid stability of the entire PDN. Intuitively, this is because higher loop gain and UGB contribute to an increased gain of the $\mathbf{G}$ block (represented using admittance transfer matrix), and hence make the gain condition more difficult to satisfy.

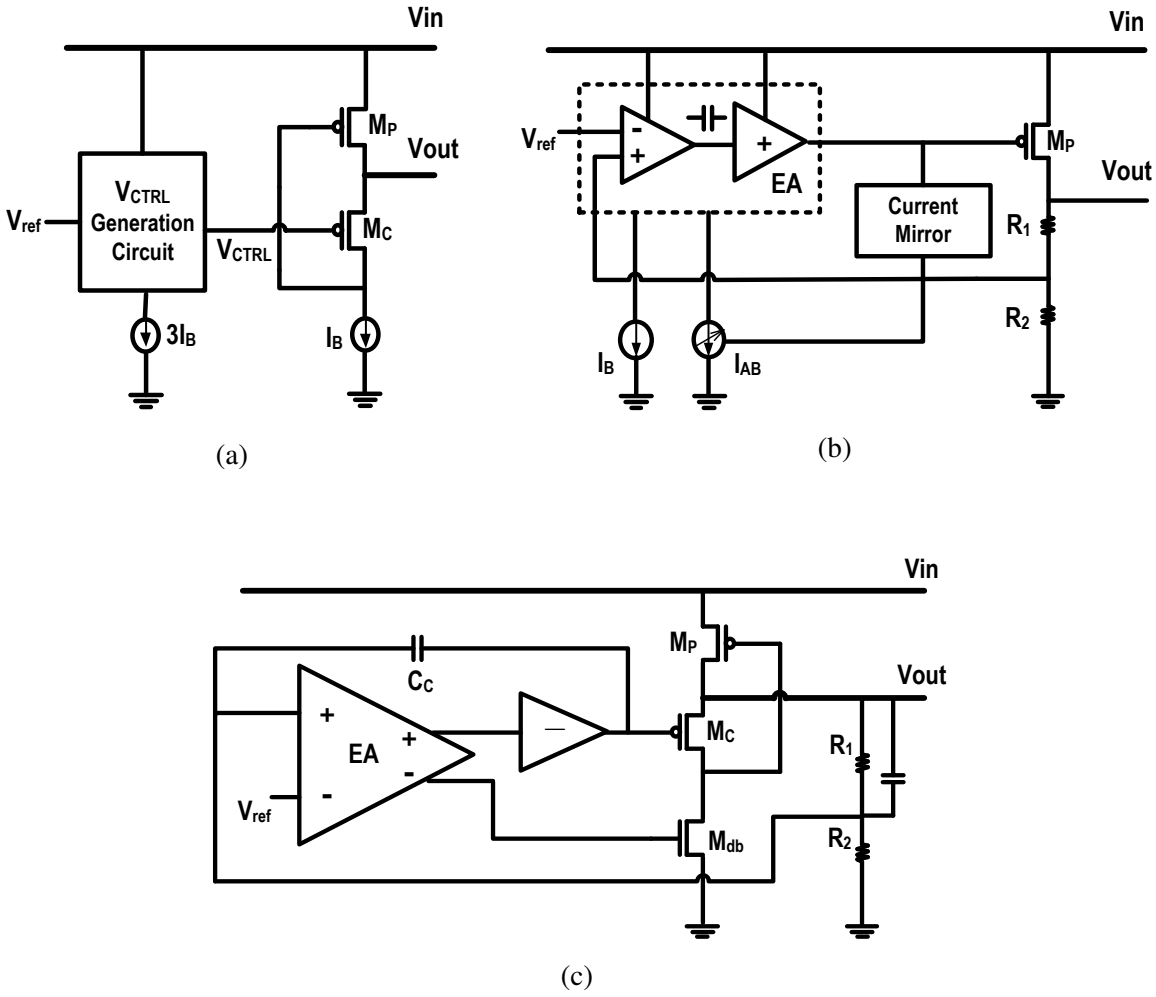


Figure 3.1: LDO topologies adopted in the experiments: (a) FVF LDO [2], (b) AB LDO [3], and (c) ML LDO [1].

To see this, we compare three different LDO topologies as in Fig. 3.1. Based on a flipped voltage follower (FVF) structure, the first LDO [2] (denoted as FVF LDO) achieves fast load

regulation and high bandwidth. However, it is limited by its relatively weak feedback loop with low gain. As contrast, while the steady-state performance of the second adaptive biased LDO [3] (AB LDO) is notably improved with a multi-stage error amplifier, its bandwidth is heavily lowered as more nodes are inserted into the loop. In the third LDO topology [1], an additional feedback path is introduced to the FVF LDO so that both the steady-state performance and the transient response are enhanced. Nonetheless, a large amount of quiescent current is consumed on the active frequency compensation scheme. As multiple feedback loops are employed to boost the performance, we denote the third topology as multi-loop LDO (ML LDO). This LDO topology labeling is used throughout this chapter.

Clearly, the above three LDO topologies demonstrate different tradeoffs between steady-state regulation, transient response and power consumption. To see how they impact the stability of the network, we implement the three LDO circuits with the same output voltage (1V) and maximum load current (50mA). The gain curve of the corresponding G block containing four identical LDOs is depicted in Fig. 3.2(a). It is observed that for most frequencies the G block with ML LDOs has the highest gain. This is caused by the enhanced loop gain and UGB of the ML LDO topology.

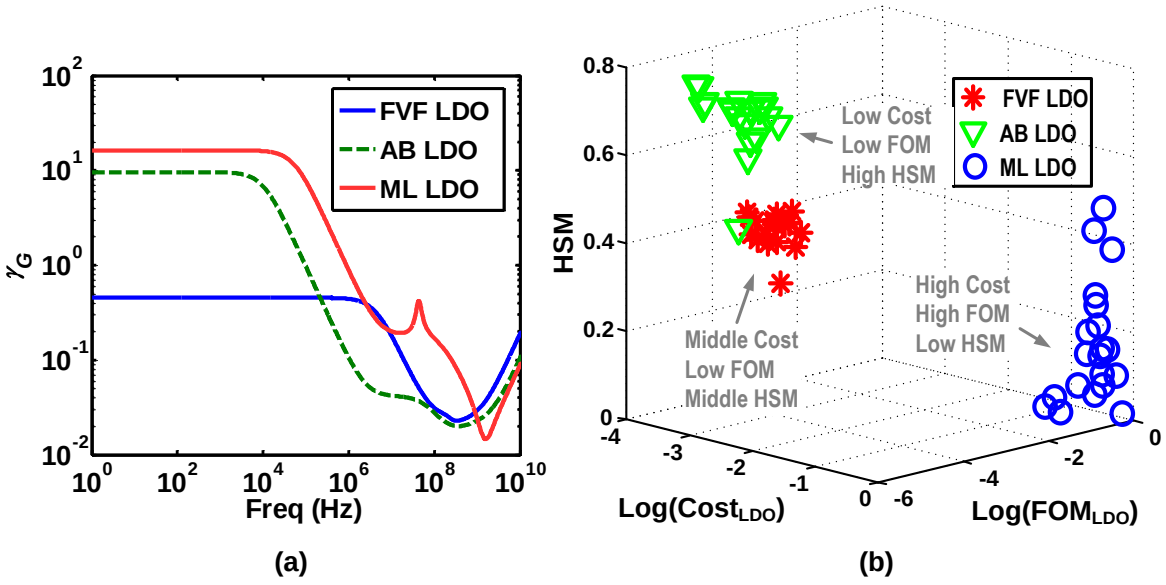


Figure 3.2: (a) $\gamma_{\mathbf{G}}$ for different LDO topologies. (b) Design space exploration for different LDO topologies. $Cost_{LDO}$ and $FOM_{LDO}$ are normalized to the maximum values among all 60 designs.

We use a PDN setup with over 100K nodes and 4 LDOs to extensively explore the impact of LDO topology on the PDN's stability. Although the passive networks are the same for all PDNs, the on-chip LDOs are designed using different LDO topologies. After the one-time AC characterization for the passive network, the HST-based optimization flow is applied to locally tune each LDO circuit and produce 20 optimal stability-ensured designs with emphasis on different LDO specifications such as the unity-gain bandwidth and the output admittance. Using 4 threads on a 2.2 GHz AMD Opteron Processor, the optimization flow takes about 2.5/7.0/13.3 hours to complete each PDN optimization with FVF/AB/ML LDOs. To evaluate the regulation performance of each LDO, we define a figure of merit (FOM) for individual LDO as $FOM_{LDO} = UGB \cdot Y_{avg} \cdot Acc$, where $UGB$ is the unity-gain bandwidth, $Y_{avg}$ is the average output admittance (within the frequencies of interest) and $Acc$ is the DC accuracy of load regulation. Besides, in order to measure the power consumption/area overhead of each LDO, a cost function is defined as $Cost_{LDO} = I_q \cdot Area$, where $I_q$ is the quiescent current and $Area$ is the LDO area.

Fig. 3.2(b) demonstrates the tradeoffs between $FOM_{LDO}$, $Cost_{LDO}$ and $HSM$ among all the 60 designs. While each LDO topology shows its own merits in cost/performance within a large design space, the PDNs with AB and FVF LDOs tend to be more stable. Additionally, Table 3.1 reports the average performance over the 20 design points for each LDO topology. Compared to ML LDO topology, the PDNs with AB and FVF LDOs achieve 2.72X and 3.94X enhancements on HSM, respectively, demonstrating that choosing the LDO topology with lower loop gain/bandwidth can lead to larger stability margin for the PDN network.

### 3.2.2 Insertion of On-Chip Decoupling Capacitance

As well known, the use of on-chip decoupling capacitors (decaps) is one of the most common noise suppression techniques. In this subsection, the HST-based design approach is leveraged to explore the new role of the on-chip decaps in the aspect of system-level stability for PDNs with distributed regulators. Indicated by our design study, the insertion of additional on-chip decaps may be beneficial to relax the hybrid stability constraint of large PDN system and lead to improved PDN performance due to the enlarged design space.

Table 3.1: Average PDN performance with different LDO topologies.

| | FVF LDO | AB LDO | ML LDO |
|---|---|---|---|
| $HSM$ | 0.49 | 0.71 | 0.18 |
| $UGB(MHz)$ | 98 | 8.2 | 641 |
| $Y_{avg}(S)$ | 0.25 | 3.49 | 4.75 |
| $Reg.ACC.$ | 98.83% | 99.97% | 99.96% |
| $I_q(\mu A)$ | 90.3 | 53.6 | 535.9 |
| $Area(\mu m^2)$ | 1948 | 1521 | 1686 |
| $FOM(MHz \cdot S)$ | 20.93 | 28.48 | 3.05e3 |
| $Cost(\mu A \cdot \mu m^2)$ | 1.64e5 | 0.82e5 | 9.09e5 |

From the HST point of view, this can be understood by noting that inserting global decaps (GD) which are allocated between the global VDD grids and GND grids can significantly suppress the mid/high-frequency impedance peaking seen from the global VDD grids. On the other hand, additional local decaps (LD) between the regulated local VDD grids and the GND grids help reduce the impedance observed from the local VDD grids, by providing lower impedance paths to ground. Since the $\mathbf{H}$ block interfaces with the $\mathbf{G}$ block through both the global VDD grids (connected with the input pins of regulators) and local VDD grids (connected with the output pins of regulators), inserting more GD/LD can lower the magnitudes of the entries in $\mathbf{H}$ associated with the global/local VDD grids, respectively. As a consequence, GD and LD can work together to decrease $\gamma_{\mathbf{H}}$ along frequency axis makes the gain condition easier to satisfy.

To see the above trend, we compare the gain curves for different amount of decaps in Fig. 3.3. In the meantime, we study the impact of decap distribution by assigning decaps based on two different patterns for a given amount of decaps. In the first even pattern, all decaps are uniformly distributed. While in the second uneven pattern, the decaps are randomly distributed but with 60% decaps concentrated on a quarter of the chip area. From Fig. 3.3, it is observed that increasing the amount of decaps can obviously reduce $\gamma_{\mathbf{H}}$, while how the decaps are disctributed only has secondary impact around the frequencies where impedance peaking happens. In the rest of the section we assume that on-chip decaps are uniformly distributed.

To fully explore the benefits of inserting decaps, we insert different amounts of on-chip decaps

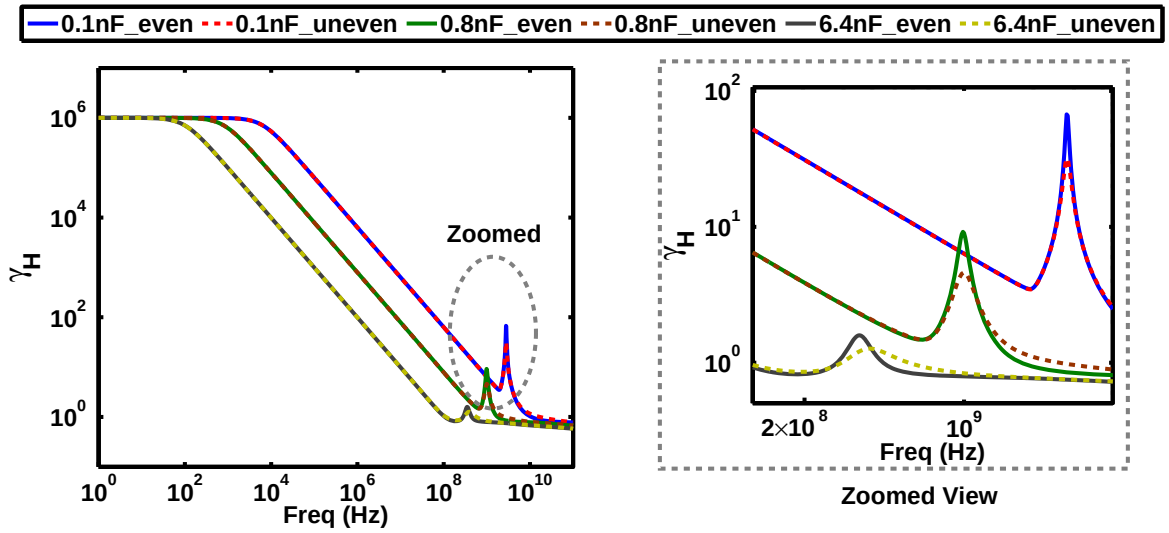


Figure 3.3: The influence of local/global decaps on the gain of H block.

containing GD and LD in a 1:1 ratio into a PDN with 100K nodes and 4LDOs and then apply the HST-based optimization flow to jointly optimize the $HSM$ and other performance. To see how the decaps insertion interacts with the choice of LDO topology, we also integrate LDOs based on different topologies to explore the joint effects of the above two key parameters on the design tradeoffs. In Fig. 3.4, the $HSM$ and $FOM_{LDO}$ are plotted as the functions of decaps. Based on the simulation results, we analyze the impacts of on-chip decaps in two cases.

### *3.2.2.1 PDN with High Loop-Gain/UGB LDOs*

As discussed in Section 3.2.1, the high $FOM_{LDO}$ can be achieved by choosing LDO topology with boosted loop gain and UGB (e.g. the ML LDO) at the price of high $Cost_{LDO}$. From the stability point of view, integrating such LDOs into a PDN will make the gain condition harder to satisfy due to large $\gamma_{\mathrm{G}}$ values. The gain condition will be even harder to satisfy around frequencies where resonance peaking occurs in the $\mathbf{H}$ block. In this case, inserting GD and LD can notably remedy the stability constraint via suppressing the impedance peaking. This trend is verified in Fig. 3.4 that for ML LDO topology, both $HSM$ and $FOM_{LDO}$ are significantly boosted with increased GD/LD. Note that the optimized PDN designs under low GD/LD setup can only have small or even negative $HSM$ due to harsh stability constraint.

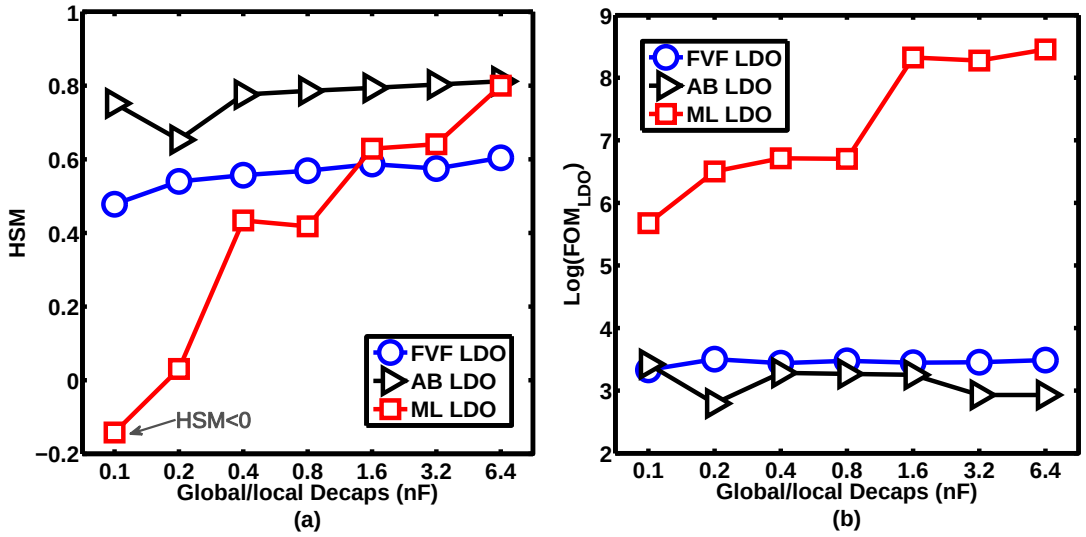


Figure 3.4: Design tradeoffs between decaps and (a) HSM, (b) $FOM_{LDO}$.

### *3.2.2.2 PDN with Low Loop-Gain/UGB LDOs*

In this case, the gain of the G block is already low around the frequencies where resonance peaking happens. As a result, inserting additional decaps cannot play a critical role as in the previous case to relieve the stability constraint. As shown in Fig. 3.4(a) and (b), though the $HSM$ for AB/FVF LDO is slightly improved with increased on-chip decaps, no obvious variations on $FOM_{LDO}$ can be observed.

Besides, we evaluate the impact of decaps on the maximum switching noise $\Delta V_{max}$ for each optimized PDN design. $\Delta V_{max}$ is captured in the local VDD grids under three different load current transitions: 1) between 1mA and 5 mA with 1ns rising/falling (Tr/Tf) time; 2) between 160mA and 200mA with 1ns Tr/Tf; 3) between 10mA and 200mA with 100ns Tr/Tf. We assume the total load current is uniformly distributed on chip. Then Fig. 3.5 depicts variation of $\Delta V_{max}$ with different decaps.

As in Fig. 3.5(a) and (b), more decaps can significantly improve the switching noise for all the LDO cases when the load transition is fast (1ns). However, the improvement for the PDNs with FVF and AB LDOs mainly comes from the enhanced noise suppression of decaps, while for the PDN with ML LDO the decap has a twofold benefit: in addition to the high frequency noise

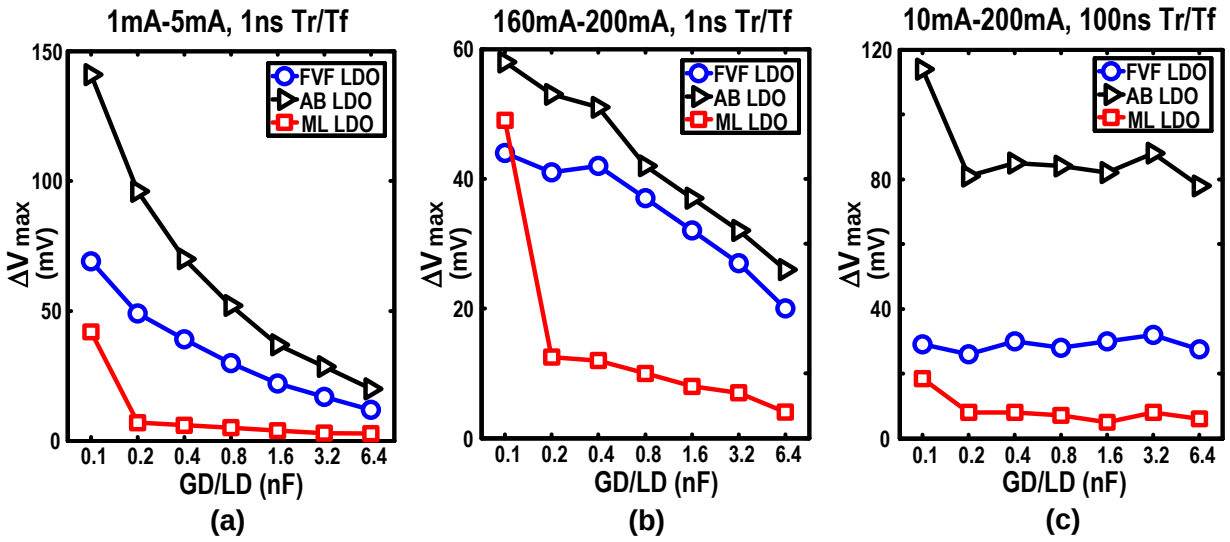


Figure 3.5: Impact of on-chip decaps on $\Delta V_{max}$ measured under three different load transitions.

suppression, it also boosts the ML LDO design due to enlarged design freedom. Also note for the third slower load transition (100ns), the resulted lower frequency noise is mainly regulated by the LDOs. As shown in Fig. 3.5(c), $\Delta V_{max}$ under this load is mainly related to the LDO performance and inserting more decaps cannot play a critical role for better switching noise, especially for FVF and ML LDOs with enough bandwidth.

### 3.2.3 Impact of LDO Number

Recall that the stability problem is caused by the inter-LDO interaction. From the point of view of hybrid stability, increasing the number of on-chip regulators may jeopardize the stability margin of the entire network due to the enhanced coupling strength between LDOs. To shed some light on this, we plug different numbers of pre-designed LDOs into a PDN and plot the gain of $\mathbf{H}$ over frequencies. As in Fig. 3.6, $\gamma_{\mathbf{H}}$ goes up when more LDOs are integrated, making the gain condition harder to satisfy. On the other hand, our investigation shows that the entire network can meet the passivity condition in a narrower frequency band when it contains more active regulators.

Additionally, our experiment results show that with a certain number of LDOs, the LDO distribution doesn't appear to be a strong factor influencing the final achievable $HSM$. Besides, although the LDO distribution could alter the current sharing among LDOs and it is preferred to

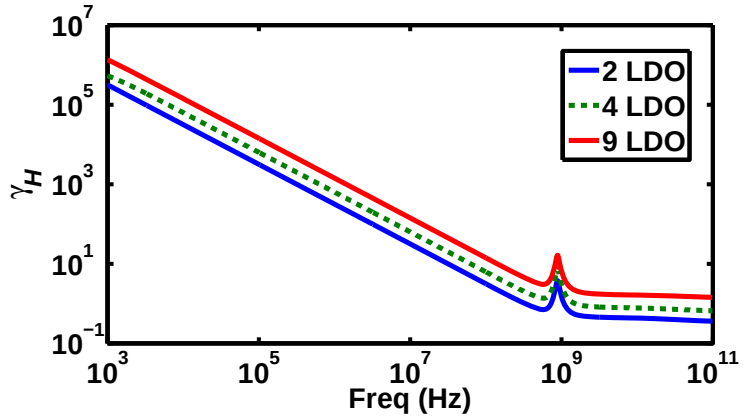


Figure 3.6: $\gamma_{\mathbf{H}}$ curves with different numbers of LDOs.

optimize the LDOs at the corresponding load point, our results show that different current sharing cases may only have limited influence on the whole-system optimization and tradeoffs. Thus we assume the LDOs are evenly distributed in this study.

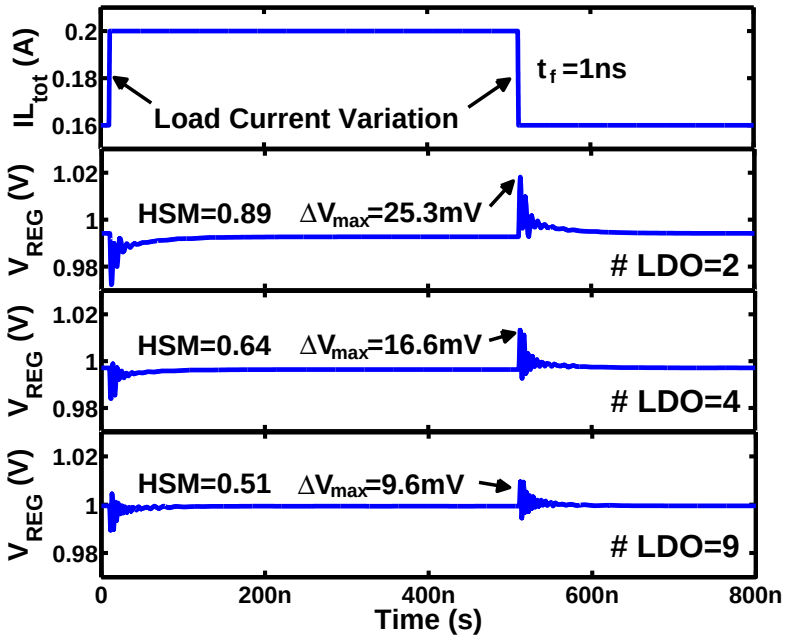


Figure 3.7: The influence of LDO number on HSM and regulation performance.

Through simulations, we examine the variation of $\Delta V_{max}$ and $HSM$ in the PDN with different

number of LDOs. Fig. 3.7 illustrates that when more LDOs are embedded in this PDN, the local voltage regulation is enhanced due to the stronger load sharing effect among different LDOs. However, the HSM of the entire network is degraded. This suggests the opportunity to optimize the number of LDOs for improved design tradeoffs between stability and performance.

### 3.2.4 Joint Effects of All Design Parameters

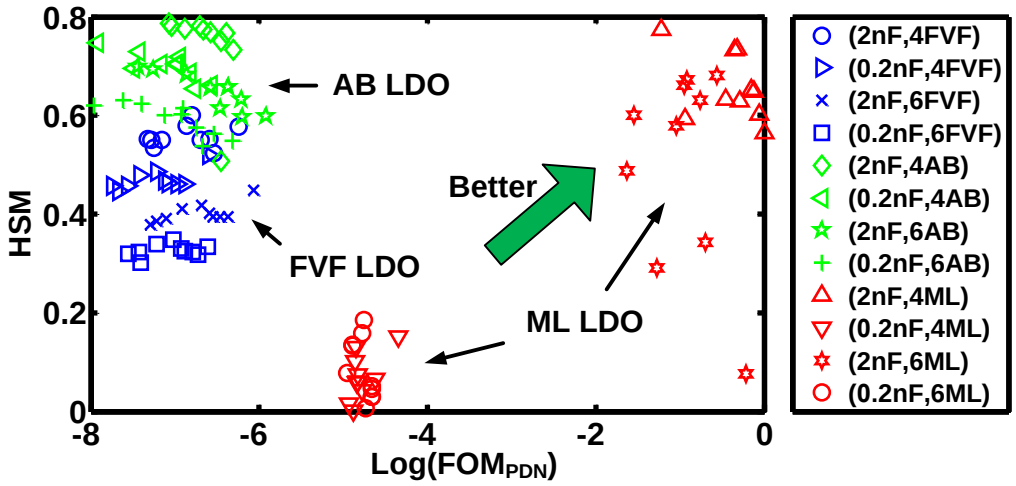


Figure 3.8: Tradeoffs between HSM and $FOM_{PDN}$ among 12 design strategies. $FOM_{PDN}$ is normalized to the maximum value among all designs.

To give a full picture of the joint effects of all design parameters, the following experiment is conducted. We obtain 12 different PDN parameter sets by combining the choice of LDO topology, the amount of decaps and the number of LDOs. To fully explore the design space, the HST-based PDN design flow is applied to generate 10 optimized PDN designs for each strategy via locally tuning the LDO design parameters. To systematically evaluate each PDN design, we define a figure of merit for PDN as $FOM_{PDN} = \frac{UGB \cdot Y_{avg} \cdot Acc}{\Delta V_{max}}$, where $UGB$, $Y_{avg}$ and $Acc$ are the unity-gain bandwidth, average output admittance and DC regulation accuracy of each LDO and $\Delta V_{max}$ is the maximum switching noise in the local VDD grids, measured under load transition from 160mA-200mA with 1ns Tr/Tf. We also define an overall cost function: $Cost_{PDN} = n^2 \cdot I_q \cdot Area \cdot C_{Decap}$,

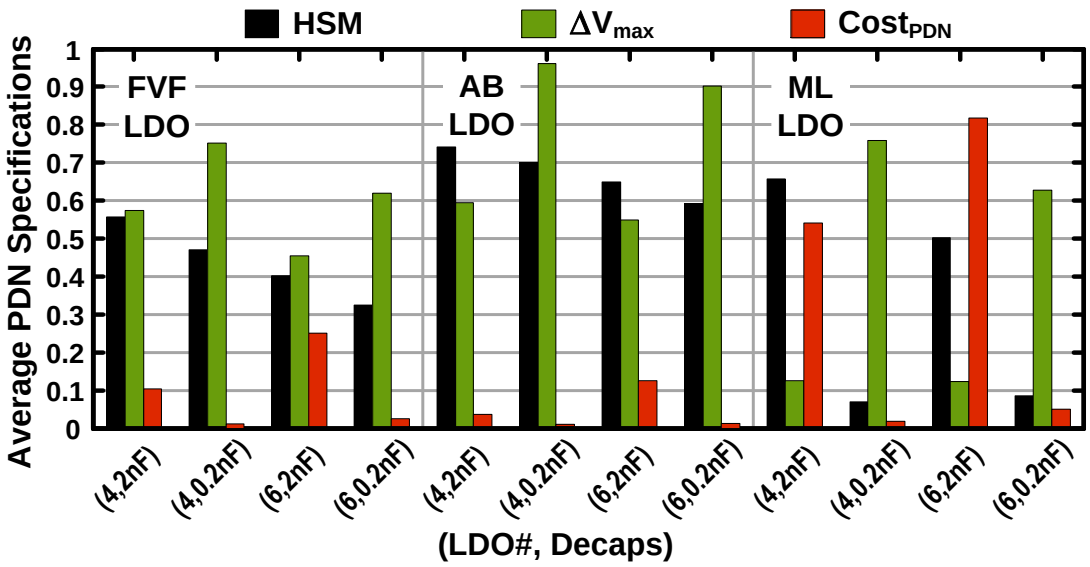


Figure 3.9: Comparison of averaged PDN specifications among 12 design strategies. $\Delta V_{max}$ and $Cost_{PDN}$ are normalized to the maximum values among all designs.

where $n$ is the number of LDOs, $I_q$ and $Area$ are the quiescent current and area of each LDO, and $C_{Decap}$ is the amount of global/local decaps. Then Fig. 3.8 visualizes the tradeoffs between $HSM$ and $FOM_{PDN}$ among all design points. And Fig. 3.9 summarizes the tradeoffs between averaged $HSM$, $\Delta V_{max}$ and $Cost_{PDN}$ for all designs. Based on these results we identify two important design cases and the detailed design suggestions are summarized.

#### *3.2.4.1 PDN Design Targeting High Regulation Performance*

In this case, regulators with high loop-gain/UGB (e.g. ML LDOs) should be adopted to enhance the load regulation of the network. Since the stability condition under this setting is harder to satisfy as discussed in Section 3.2.2, additional decaps are needed to relieve the stability constraint for improved LDO performance. Of course extra decaps will also help suppress the high-frequency switching noises. In this way, the boosted PDN regulation performance is obtained at the price of extra decap overhead as well as the high quiescent current required by ML LDOs. This design trend can be seen in Fig. 3.9 as the lowest $\Delta V_{max}$ is achieved in PDNs with 4/6 ML LDOs and 2nF decaps.

On the other hand, increasing the number of LDOs does not necessarily always enhance the

regulation performance of PDN. This could be explained by noting that more LDOs may exacerbate the stability condition as mentioned in Section 3.2.3 and cause performance degradation. This is verified by noting that with 2nF decaps, similar $\Delta V_{max}$ is achieved for PDNs with 4 and 6 ML LDOs.

#### *3.2.4.2 PDN Designs Targeting Low Cost*

In this circumstance, the regulators with low loop-gain/UGB such as AB and FVF LDOs should be used as they consume much smaller $I_q$. The PDNs with such regulators usually achieve sufficient $HSM$ within a large design space, indicating that less design efforts are needed in the aspect of stability. This can be observed from Fig. 3.8 and Fig. 3.9. On the other hand,inserting more regulators and decaps into the network helps suppress $\Delta V_{max}$ due to the strengthened load sharing between regulators and better suppression on switching noises, demonstrating a tradeoff between cost and regulation performance.

## 3.3 Summary

This chapter identifies critical design parameters impacting the stability of PDN network. Systematic design study is conducted to explore the relationship between those parameters and the PDN design tradeoffs. Based on the HST-based PDN design approach, several new design insights in terms of the network-wide PDN stability can be summarized to facilitate improved PDN design tradeoffs.

# 4. MACHINE LEARNING ENABLED POWER MANAGEMENT FOR HETEROGENEOUS VOLTAGE REGULATION SYSTEM

The work presented in this chapter is based on the vision that the ultimate power integrity and efficiency may be best achieved via a heterogeneous chain of voltage processing starting from on-board switching voltage regulators (VRs), to on-chip switching VRs, and finally to networks of distributed on-chip linear VRs. As such, we first propose a heterogeneous voltage regulation (HVR) architecture encompassing regulators with complimentary characteristics in response time, size, and efficiency. Through holistic exploration of heterogeneous voltage regulators and their systematic adaption considering complex interdependencies between such regulators, we then develop comprehensive workload-aware control policies to adapt HVR with respect to workload change at multiple temporal scales to significantly improve system power efficiency while providing a guarantee for power integrity. The proposed power management techniques are further supported by hardware-accelerated machine learning prediction of non-uniform spatial workload distributions for more accurate HVR adaptation at fine time granularity.

## 4.1 Motivation of Heterogeneous Voltage Regulation

### 4.1.1 Overview of Voltage Regulators

Voltage regulators (VRs) are key components of a power delivery system and the characteristics of VRs have critical impacts on power efficiency and regulation response of the entire system. Switching at a rate of hundreds of KHz to tens of MHz, off-chip buck voltage regulators (VR) can achieve excellent efficiency at the expense of bulky and costly off-chip LC components [7, 8]. Furthermore, off-chip VRs have slow response times, and hence cannot support fine-grained dynamic voltage scaling (DVS). There has been a great deal of progress on fully-integrated buck VRs thanks to on-die/in-package inductors and new magnetic materials [13, 14, 15]. Operating at a frequency of tens or hundreds of MHz, fully integrated buck VRs come with fast response times and promises for efficient local power delivery and fine-grain DVFS. However, Integrating

high-Q power inductors to support large current density with low loss is still a significant challenge [13, 14, 15]. Compared to their off-chip counterparts, on-chip buck VRs incur more conduction and switching losses, leading to lower efficiency, especially at light loads. On-chip linear voltage regulators (e.g. LDOs) are area efficient and can achieve sub-ns response times [16]. Their efficiency drops with increasing dropout voltage, making them inefficient for wide-range voltage conversion. Clearly, those VRs have complimentary characteristics in response time, area and power efficiency and none of them can address the IC power delivery challenge alone.

**Conversion vs. regulation.** While *conversion* and *regulation* are used almost interchangeably, we shall note a fine distinction between them w.r.t. the best ways for realizing conversion and regulation. Switching VRs are well suited for wide-range voltage conversion for which linear VRs suffer from large loss. On the other hand, area-efficient integrated linear VRs provide fast regulation. Table 4.1 summarizes the characteristics of different VRs.

Table 4.1: Comparison of different VRs.

| | Settling time | Area | Efficiency | Function |
|---|---|---|---|---|
| Off-chip buck | 10's of us [7] | Large | High | Conversion |
| On-chip buck | 10's of ns [15] | Medium | Medium | Conversion |
| On-chip LDO | sub-ns [16] | Small | Low w/ large $V_{in} - V_{out}$ | Regulation |

### 4.1.2 Heterogeneous PDN Architecture

Three power delivery architectures are illustrated in Fig. 4.1. The single-stage PDN is managed by only off-chip buck converters, achieving a high efficiency over a wide workload range. However, the board/package parasitics degrade the power quality delivered from the off-chip bucks to the on-chip power domains. Furthermore, the slow response time of off-chip buck converters limits the application of fine-grained DVS. Thanks to the progress of on-die/in-package inductors

and new magnetic materials the buck converters can be integrated on chip. The two-stage PDN (Fig. 4.1(b)) consists of both off-chip and on-chip buck converters, improves the quality of power delivery by lowering the impedance from the power supply to the load circuits, and well supports fine-grained per-core DVS since the integrated VRs can settle much faster. These benefits make this architecture widely used in modern SoCs such as the Intel's Haswell processors [55]. However, the response time of on-chip buck converters can still limits the PDN performance in the case of highly unpredictable load currents which may occur, for example, in server-class processors [21].

We argue that the ultimate quality and efficiency in supply voltage regulation may be only achieved by fully exploiting the heterogeneity in PDN architecture with heterogeneous regulators with complimentary characteristics in response time, power efficiency, and cost. As shown in Fig. 4.1(c), we propose a *heterogeneous voltage regulation* (HVR) architecture with three voltage processing stages: multiple off-chip buck VRs supplying power to multiple clusters of on-chip buck VRs with each cluster powering a network of distributed on-chip LDO driving a power/voltage domain. Fig. 4.2(a) depicts a more detailed view of the three-stage HVR. Clearly, the first stage enjoys the high efficiencies of off-chip buck VRs over wide ranges of workloads. Their slow response is compensated by the second stage of on-chip buck VRs. Bypassing board/package parasitic impedances, on-chip buck VRs can settle much faster, enabling fine grained per core DVS otherwise impossible. Having two stages of buck converters gives the added benefit of lowering the step-down ratio for each stage, improving the efficiency of both off-chip and on-chip buck converters, and reducing sizes of the off-chip passives and power transistors [43]. Leaving most of the *voltage conversion* functionality to the first two stages, the on-chip LDO networks act as the last (main) stage of *voltage regulation*. Due to the small footprint of LDOs, a large number of compact LDOs with ultra-fast response time can be placed on-chip in a distributed manner within a power domain (Fig. 4.2(b)), forming an interconnected active regulation network. In vicinity of on-chip hot spots, on-chip LDO network can respond very quickly to local voltage droops, achieving good regulation performance.

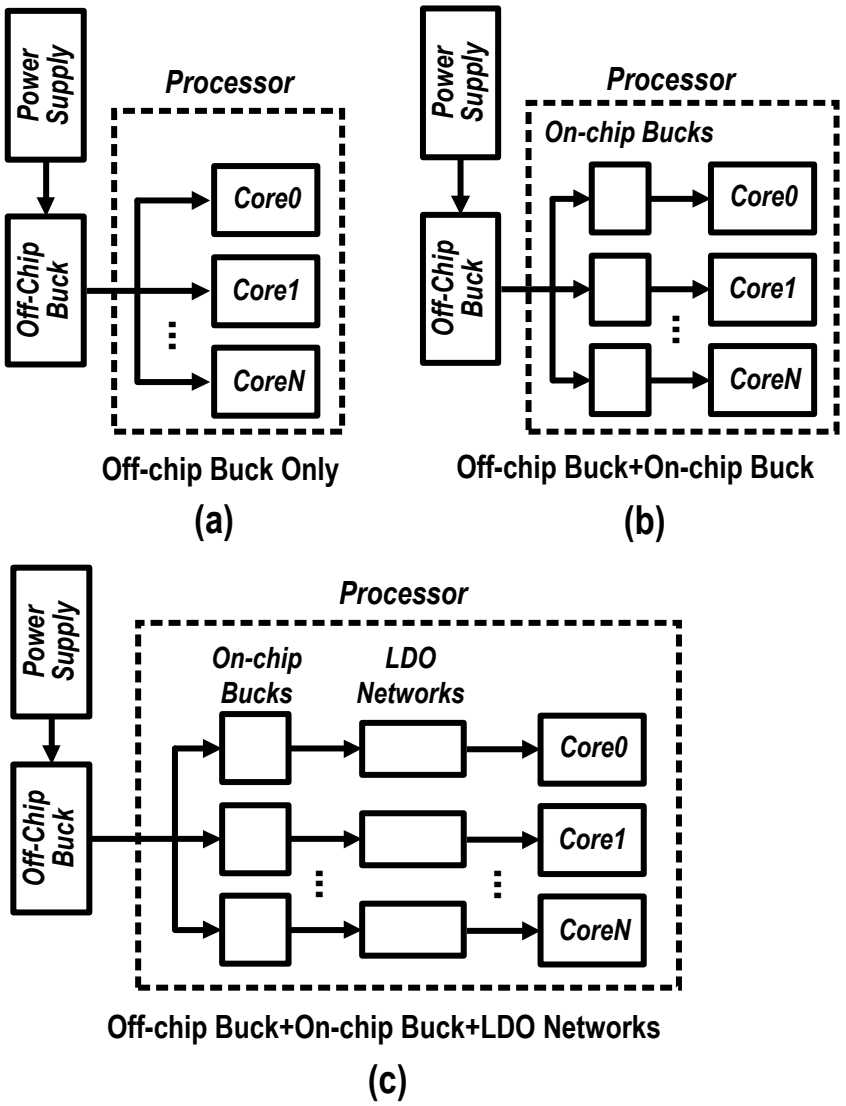


Figure 4.1: PDN architectures: (a) single-stage PDN using off-chip buck converters, (b) two-stage PDN using both on- and off-chip buck converters, and (c) proposed three-stage heterogeneous voltage regulation (HVR).

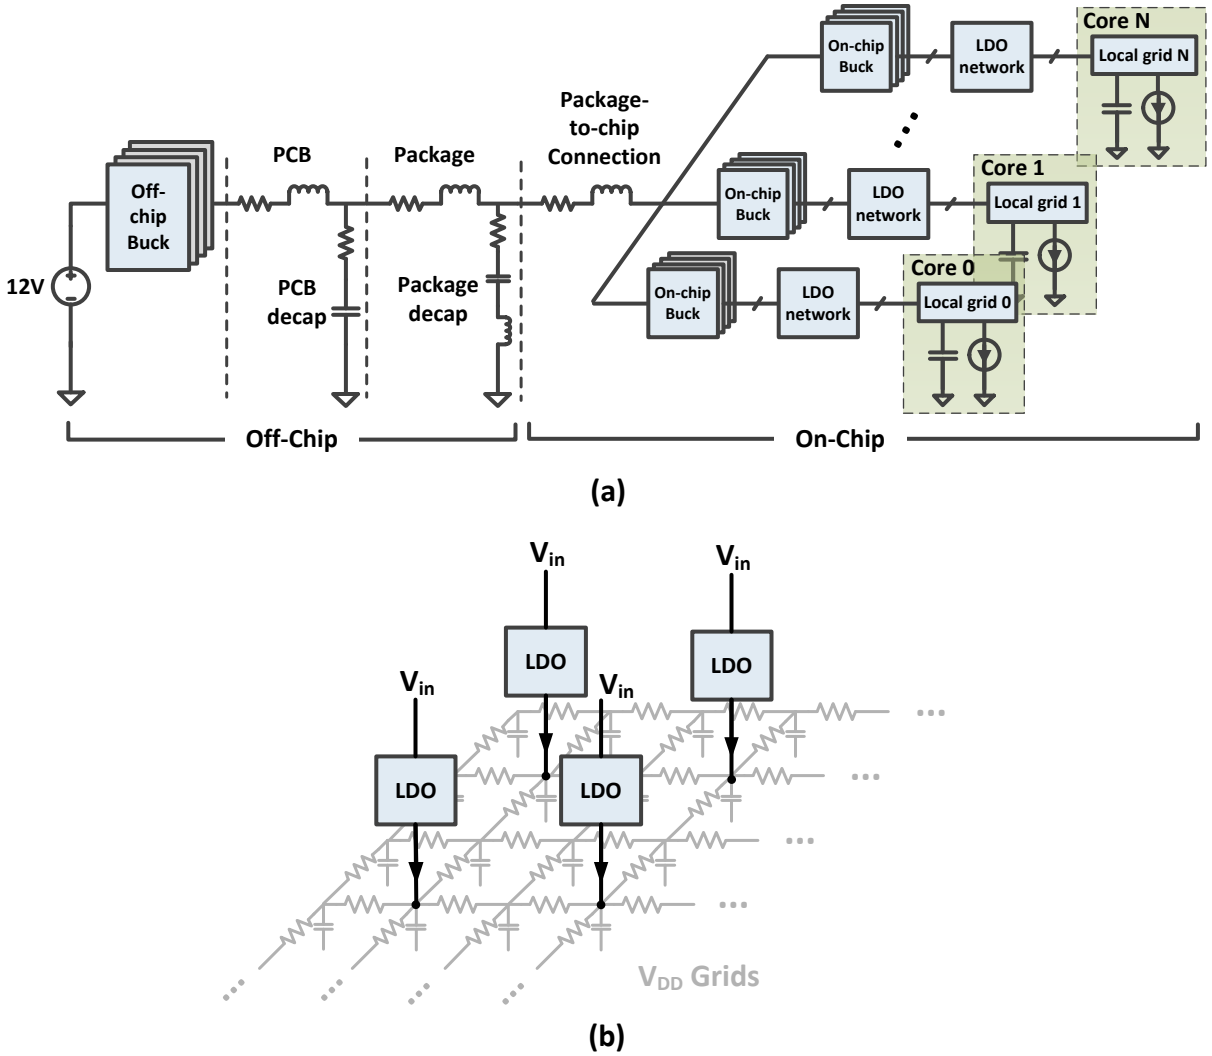


Figure 4.2: (a) Modeling of 3-stage heterogeneous voltage regulation system, and (b) distributed LDO network.

### 4.1.3 Tuning Opportunities in HVR

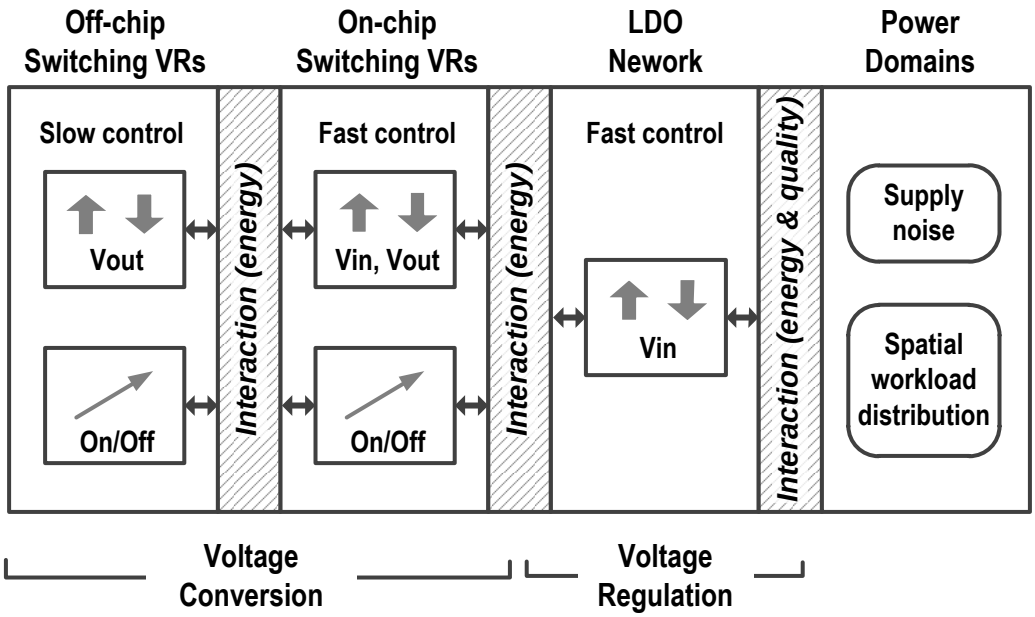


Figure 4.3: Overview of tunability in HVR system.

Heterogeneity brings in a great deal of tunability at multiple HVR stages for workload-aware adaption. The power efficiency of a single VR stage is usually a function of its input/output voltages and current load. For a cluster of VRs, its power efficiency can be optimized according to runtime workload by either tuning its input/output voltages or modulating the number of online VRs, which changes the load per regulator. There are important interdependencies among different voltage processing stages which must be carefully considered in order to optimize the overall energy efficiency and regulation performance. For example, the output of the preceding VR stage is also the input of the subsequent VR stage. Fig. 4.3 summarizes the rich tunability and complicated energy and performance interdependencies in HVR system.

We define several important control variables in Table 4.2, and will use them throughout this chapter. Considering a HVR system consisting of $N$ power domains as in Fig. 4.1(c), the control decision variables are: the number of online converters in each on-chip buck VR cluster $N^{(i)}_{online,on}$,

Table 4.2: Control variables in HVR.

| Variable | Description |
|---|---|
| $N^{(i)}_{online,on}$ | Number of online VRs in the $i$-th on-chip buck cluster |
| $V^{(i)}_{out,on}$ | Output voltage of the $i$-th on-chip buck cluster |
| $N_{online,off}$ | Number of online VRs in the off-chip buck cluster |
| $V_{out,off}$ | Output voltage of the off-chip buck cluster |

and the cluster's output voltage $V^{(i)}_{out,on}$, $i = 1, ..., N$, which is the input voltage to the LDO network driven by the cluster; the number of online converters in the off-chip buck cluster $N_{online,off}$, and its output voltage $V_{out,off}$, which sets the input voltage to all on-chip buck VR clusters in the considered tree.

## 4.2 Modeling of HVR System

Clearly, the HVR voltage processing chain has a tree structure consisting of multiple voltage processing stages starting from a cluster of off-chip buck VRs and ending at the on-die loads in each local power domain. We look into the detailed energy and regulation characteristics at each individual stage first then consider the interdependencies across different stages in the HVR system.

### 4.2.1 Characteristics per Stage

#### *4.2.1.1 On/Off-Chip Buck Clusters*

Fig. 4.4 shows a typical multi-phase buck converter which is widely adopted in modern SoCs. For each single phase of the buck VR, the pulse-width-modulation (PWM) comparator sets the duty cycle of its output voltage waveform which then drives power switches to produce the modulated final output voltage. The multiple parallel time-interleaved phases cancel out the high-frequency output noise and reduce the transient response time at the cost of increased overhead of inductors and control circuits [56].

The major power losses of a buck converter include two parts: the switching loss which is largely independent of load current, and the resistive loss which is a function of load current

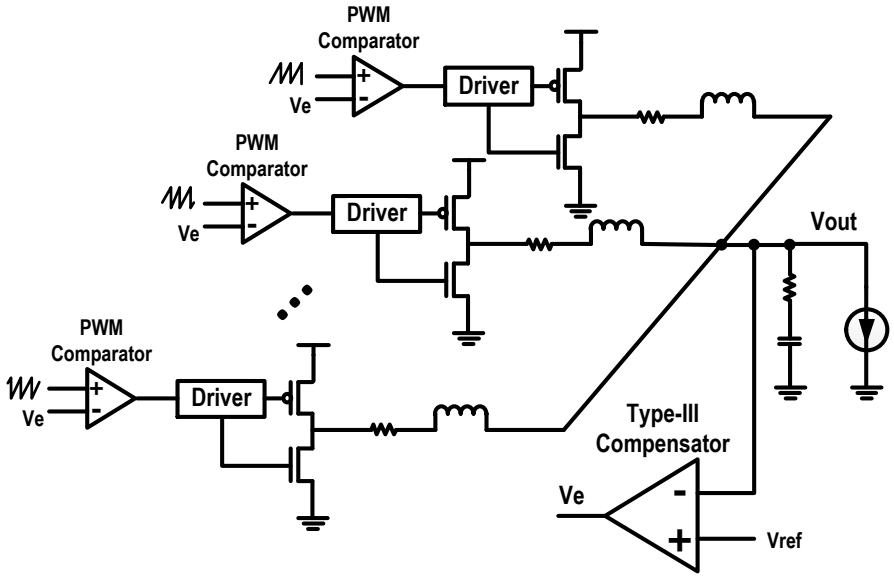


Figure 4.4: Schematic of a multi-phase PWM buck converter.

[57]. The switching loss dominates the power loss at light loads while the resistive loss grows quadratically with increasing load current. Besides, both parts of power loss are functions of the input/output voltages of the buck converter. In a cluster of buck VRs, its overall power efficiency can be further impacted by the number of online VRs $N_{online}$ which varies the total switching loss under the same overall load current. As a result, the general form for the power efficiency of a buck cluster can be written as

$$\eta_{buck} = f(V_{in}, V_{out}, N_{online}, I_L) \tag{4.1}$$

For a given input/output voltages, a single VR achieves the peak efficiency at an optimal load point $I_{opt}$ where the ratio of total loss over the load power is minimized. Besides, Fig. 4.5(a) demonstrates that the power efficiency curves of a buck cluster can be dramatically changed with different number of online VRs $N_{online}$. The peak power efficiency for each curve can only be achieved at a certain optimal current load point, which is roughly $N_{online}I_{opt}$. Therefore, it is intuitive to bring online only a certain number of buck VRs in the cluster such that the load current

per VR stays around $I_{opt}$. The number of required VRs $N_{online}$ can be quantized as

$$N_{online} = min\{N_{max}, \left\lceil \frac{I_L}{I_{opt}} \right\rceil\}. \tag{4.2}$$

where $N_{max}$ is the maximum number of VRs in a cluster and $I_L$ is the total load current.

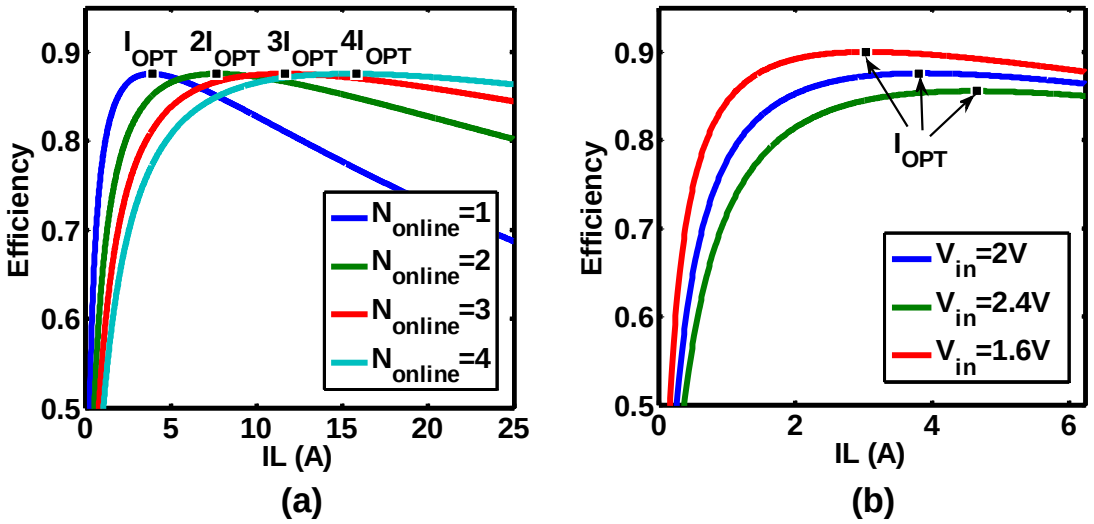


Figure 4.5: (a) Impact of online buck VRs on power efficiency, and (b) impact of input voltage on $I_{opt}$ for a single buck VR.

Note that the chain structure of the HVR makes things much more complicated, because $I_{opt}$ is a function of VR's input/output voltages which can be influenced by the preceding and subsequent stages. Fig. 4.5(b) illustrates the shift of $I_{opt}$ for a single buck converter with varied input voltage. Such effect must be considered in (4.2). As will be discussed later, the adaptive control policy proposed in this work requires short processing latency to enable fine grained temporal control resolutions. Therefore, the complex characteristics of buck converters are stored in two look-up-tables (LUTs) for the ease of online use. For instance, $LUT^{\eta}$ stores the power efficiency characteristics which are indexed by the input/output voltages and load current for each buck VR. As a function of the input/output voltage levels, $LUT^{Iopt}$ stores the optimal load current under

which the peak efficiency is achieved for a single buck VR.

Although the buck converters are more suitable for voltage conversion as discussed earlier, the on-chip buck VRs, which is the final stage in the conventional 2-stage PDN, have to be carefully designed with the consideration of supply noise. The power integrity will be largely determined by the transient response of the on-chip buck VRs. In general, increasing the switching frequency of the buck VRs will help reduce both the transient response time and output voltage ripple but at the price of increased switching power loss. As a result, it is common to integrate on-chip buck converters operating at hundreds of MHz in the 2-stage PDN [55].

#### *4.2.1.2 On-Chip LDO Networks*

The proposed 3-stage HVR system explores the fast voltage load regulation of an additional stage of distributed on-chip LDOs as discussed earlier. In addition, LDOs can be designed with a good power supply ripple rejection (PSRR) to suppress noise from the input voltage (i.e. line regulation) [1]. As a result, the on-chip buck converters in the 3-stage HVR can be optimized to achieve better power efficiency, e.g. by operating them at a lower switching frequency.

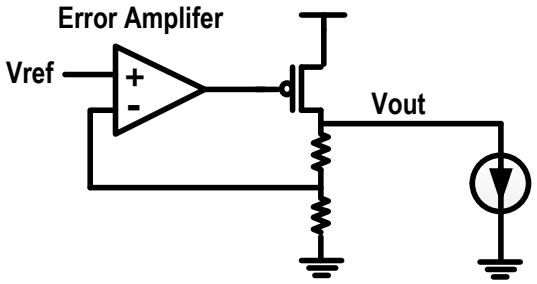


Figure 4.6: Schematic of LDO.

To supply a specific output voltage, a linear LDO converts an input voltage using an error amplifier and feedback loop as depicted in Fig. 4.6. The power efficiency of an LDO is strongly limited by the input-to-output differential voltage $\Delta V = V_{in} - V_{out}$ for a given targeted output

voltage $V_{out}$:

$$\eta_{ldo} = \frac{V_{out}}{V_{out} + \Delta V}. \tag{4.3}$$

At a certain load point, the dropout voltage $V_{drop}$ of an LDO is defined to be the minimum $\Delta V$ at which the LDO ceases to regulate its output voltage, i.e. entering the dropout region from the regulation region. 4.7(a) illustrates $V_{out}$ as a function of $V_{in}$. Therefore, it is desirable to set $V_{in}$ just $V_{drop}$ above $V_{out}$ to keep the LDO at the boundary between the dropout and regulation regions to maximize efficiency. However, setting $V_{in}$ too low may jeopardize the regulation of LDOs and violate power integrity.

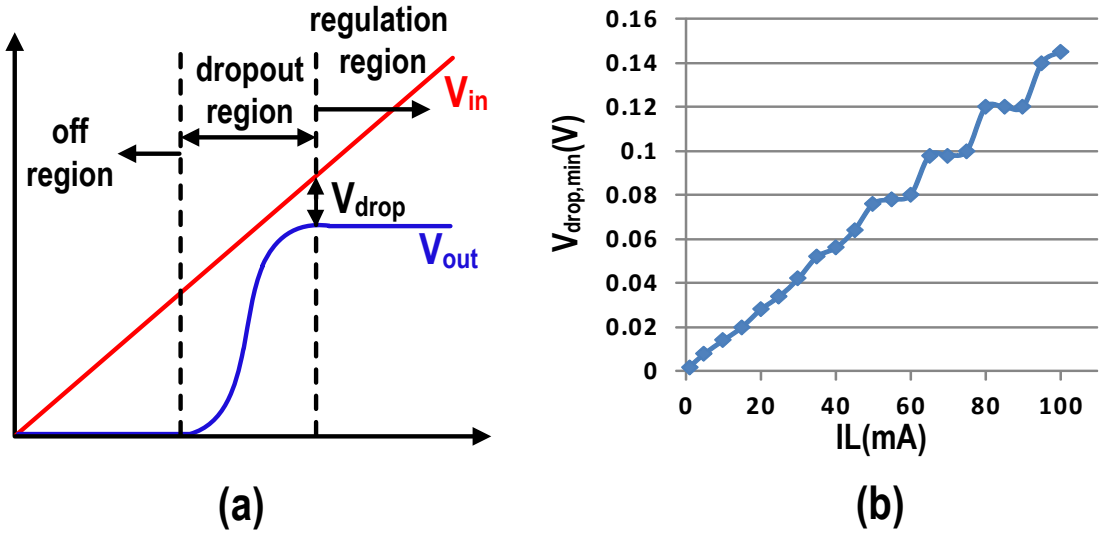


Figure 4.7: Relationship between LDO's dropout voltage and load current.

$V_{drop}$ is a function of the load current $I_L$, which is shown in Fig. 4.7(b) for a realistic LDO design [1]. It can be seen that $V_{drop}$ is approximately linear in $I_L$, hence $V_{drop} \approx \frac{I_L}{I_{L,max}} V_{drop,max}$, where $V_{drop,max}$ is the dropout voltage at the maximum current load $I_{L,max}$. Given a target output voltage $V_{dd}$, e.g. one set by DVS, the optimal LDO's input voltage (output voltage of the on-chip

buck VRs), which leads to the highest of LDO power efficiency, is:

$$V_{in,opt} \approx V_{dd} + \frac{I_L}{I_{L,max}} V_{drop,max}. \tag{4.4}$$

### 4.2.2 Interdependencies between Voltage Processing Stages

According to the above discussion, the power efficiency of a single VR stage largely depends on its input/output voltages and current load. Thus, there are important interdependencies among voltage processing stages which must be carefully considered in order to optimize the overall energy efficiency and regulation performance. Such interdepencies can be observed in (4.5), which describes the overall power efficiency as the product of efficiencies at all stages. Since the input voltage of the off-chip buck VRs is assumed to be constant, it is not considered in the corresponding power efficiency $\eta_{buck,off}$. Under a certain workload $\vec{I}_{L,on} = \{I_{L,on}^{(1)}, I_{L,on}^{(2)}, ...I_{L,on}^{(N)}\}$ and DVS setting $\vec{V}_{dd} = \{V_{dd}^{(1)}, V_{dd}^{(2)}, ...V_{dd}^{(N)}\}$, where $N$ is the number of power domains, the control variables listed in Table 4.2 can simultaneously influence the power efficiencies at multiple stages due to the interdepencies in the voltage regulation chain. For example, the output voltage of the off-chip buck VRs $V_{out,off}$ influences the efficiencies of both off- and on-chip buck VRs. Set by the output of corresponding on-chip buck cluster $V_{out,on}^{(i)}$, the input voltage to an LDO network significantly impacts the power efficiencies of the (preceding) on-chip buck cluster, and the final power quality for the loads observed on the power grids. As a result, such interdependencies have to be considered in the online adaption for maximal power efficiency and noise tradeoffs.

$$\begin{aligned} \eta_{HVR} =& \eta_{buck,off}(V_{out,off}, N_{online,off}, I_{L,off}) \\ & \eta_{buck,on}(V_{out,off}, \vec{V}_{out,on}, \vec{N}_{online,on}, \vec{I}_{L,on}) \eta_{ldo}(\vec{V}_{out,on}, \vec{V}_{dd}) \end{aligned} \tag{4.5}$$

## 4.3 HVR Control Policies

We present our proposed control policies for 3-stage HVR while these policies can be straightforwardly applied to adapt 2-stage HVR consisting of only off- and on-chip switching VRs. Unlike most related work executing power management in the OS or software [41, 42], the proposed poli-

cies can be efficiently implemented in firmware based on simple arithmetics and pre-computed look-up-tables (LUTs) supported by hardware accelerated machine learning prediction of workload.

The settling times of off- and on-chip switching VRs of the the first two stages can differ by several orders of magnitude. Hence, they are adapted using two different control cycle times, denoted by $T_{off}$ and $T_{on}$, respectively. Each $T_{off}$ is split into multiple cycles of $T_{on}$. Accordingly, off- and on-chip switching VRs are adapted by two control procedures, which are shown in Fig. 4.8 for a HVR system with $N$ power domains, one for each core. We estimate core-level workloads $\vec{I}_{L,Toff} = \{I^{(1)}_{L,Toff}, I^{(2)}_{L,Toff}, ..I^{(N)}_{L,Toff}\}$ and $\vec{I}_{L,Ton} = \{I^{(1)}_{L,Ton}, I^{(2)}_{L,Ton}, ..I^{(N)}_{L,Ton}\}$ respectively at the time granularities of $T_{off}$ and $T_{on}$ using power sensors [58] at the output of each on-chip switching VR (buck converter) cluster. At both time scales, we use the workload estimates obtained from the previous control cycle to to generate control actions for the current cycle.

In each off-chip VR control cycle $T_{off}$, the off-chip VR control procedure $VR\_OFF\_OPT$ is invoked to optimize the off-chip VR otuput voltage $V_{out,off}$ and the number of online off-chip VRs $N_{online,off}$ based on $\vec{I}_{L,Toff}$. Each $T_{off}$ is divided into a multiple of much finer grained on-chip VR control cycles $T_{on}$ as shown in Fig. 4.9. The on-chip control procedure $VR\_ON\_OPT$ is invoked in each $T_{on}$ cycle to adjust the output voltage $V^{(i)}_{out,on}$ and the number of online VRs $N^{(i)}_{online,on}$ for each on-chip VR cluster, $i = 1, 2, ...N$, based on the finer grained workload estimation $I^{(i)}_{L,Ton}$. As detailed in Section 4.4, $VR\_ON\_OPT$ relies on a machine learning module utilizing a small number ($N_S$) of voltage sensors to more precisely adjust $\vec{V}_{out,on} = \{V^{(1)}_{out,on}, V^{(2)}_{out,on}, ..., V^{(N)}_{out,on}\}$, based on the spatial distribution of the workload in each power domain. The voltage sensor readings $\vec{V}_{sensor} = \{V^{(1)}_{sensor}, V^{(2)}_{sensor}, ..., V^{(N_S)}_{sensor}\}$ are included as input to $VR\_ON\_OPT$.

Fig. 4.9 shows the timing of the control sequences. There are three steps involved in each $T_{on}$ cycle. The first decision making step executes $VR\_ON\_OPT$ procedure to compute $\vec{V}_{out,on}$ and $\vec{N}_{online,on}$ and the on-chip VRs are adjusted accordingly in the second decision execution step.

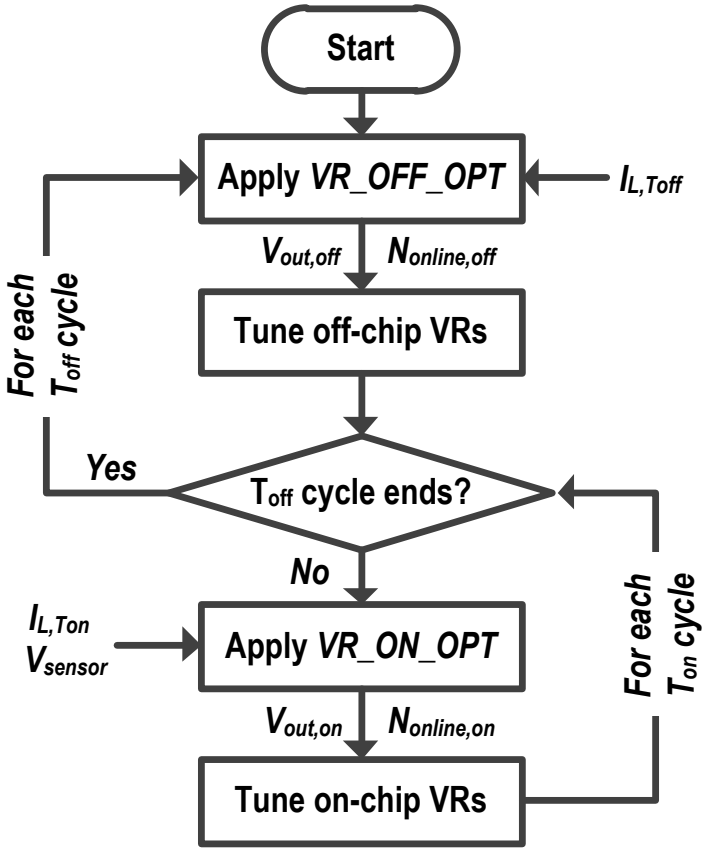


Figure 4.8: Control of off- and on-chip switching VRs at two time scales.

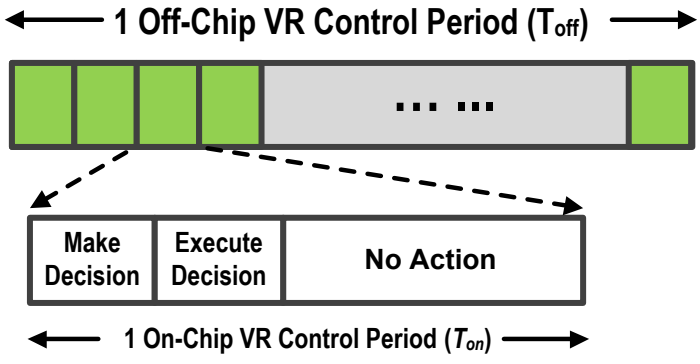


Figure 4.9: Two Control sequences.

### 4.3.1 Off-Chip Switching VR Control

The output voltage $V_{out,off}$ of the off-chip switching VRs is the input voltage to all on-chip switching VR clusters. $V_{out,off}$ impacts the power efficiencies of both on-chip and off-chip buck VRs as well as the resistive power loss due to PCB/package parasitics. As in Algorithm 1, the off-chip control procedure $VR\_OFF\_OPT$ uses the following iterative search to find the optimal $V_{out,off}$ among a set of discretized values of $V_{out,off}$ while considering the above interactions. At each iterative search step with a targeted $V_{out,off}$ value, we first estimate the input voltage to each LDO network $V_{out,on}^{(i)}$ in line 5 as a linear function of workload to maximize the power efficiency of the LDO's as in Section 4.2.1.2 for 3-stage HVR. Otherwise, for 2-stage PDN $V_{out,on}^{(i)}$ is directly set by system's power management (e.g. DVFS) unit as shown in line 7. Then, the optimal load point for each on-chip buck VR $I_{opt}^{(i)}$ is determined in line 9 via a LUT with the known input/output voltages. $N_{online,on}^{(i)}$ is further determined in line 10. The power efficiency for each on-chip buck cluster is conveniently estimated through the use of another LUT in line 11. The total through-package current, which is the sum of the input currents of all on-chip buck clusters is computed in line 13 and used as the load current of the off-chip buck cluster. The power efficiency of the on-chip components of HVR is computed in line 14 considering both integrated buck VRs and LDOs. Our experimental study shows that the resistive loss caused by PCB/package parasitics may not be negligible, which is considered in line 15. Following a similar procedure, $N_{online,off}$ and the off-chip power efficiency are determined in lines 16-18. The overall system power efficiency at the current value of $V_{out,off}$ is the product of the efficiencies of all stages as in line 19. Finally, the combination of the value of $V_{out,off}$ and the corresponding $N_{online,off}$ that maximizes the system efficiency is chosen as the optimal control of the off-chip buck VRs for this $T_{off}$ cycle.

### 4.3.2 On-Chip Switching VR Control

Once the slowly changing variables $V_{out,off}$ and $N_{online,off}$ are determined for each $T_{off}$ cycle, $V_{out,on}^{(i)}$ and $N_{online,on}^{(i)}$ per domain are updated for each finer temporal cycle $T_{on}$ by calling $VR\_ON\_OPT$ shown in Algorithm 2. We follow a flow similar to $VR\_OFF\_OPT$ to deter-

**Algorithm 1** Off-chip control algorithm $VR\_OFF\_OPT$.

**Inputs:**
Workload current estimations $\vec{I}_L$ for each $T_{off}$ cycle.

Maximize $\eta(V_{out,off})$, subject to
  $V_{min\_on} \leq V_{out,off} \leq V_{max\_on}$
  **for** each power domain $i$ **do**
    **if** $3\_stage\_HVR == True$ **then**
      $V_{out,on}^{(i)} = V_{dd}^{(i)} + \alpha I_L^{(i)} / I_{max}^{(i)}$
    **else**
      $V_{out,on}^{(i)} = V_{dd}^{(i)}$
    **end if**
    $I_{opt,on}^{(i)} = LUT_{on}^{Iopt}(V_{out,off}, V_{out,on}^{(i)})$
    $N_{online,on}^{(i)} = \left\lceil I_L^{(i)} / I_{opt,on}^{(i)} \right\rceil$
    $\eta_{on\_sw}^{(i)} = LUT_{on}^{\eta}(V_{out,off}, V_{out,on}^{(i)}, I_L^{(i)} / N_{online,on}^{(i)})$
  **end for**
  $I_{pkg} = \sum_i (I_L^{(i)} V_{out,on}^{(i)} / V_{out,off} / \eta_{on\_sw}^{(i)})$
  $\eta_{on\_chip} = (\sum_i V_{dd}^{(i)} I_L^{(i)}) / I_{pkg} / V_{out,off}$
  $\eta_{pkg} = V_{out,off} / (V_{out,off} + I_{pkg} R_{pkg})$
  $I_{opt,off} = LUT_{off}^{Iopt}(V_{ext}, V_{out,off})$
  $N_{online,off} = \lceil I_{pkg} / I_{opt,off} \rceil$
  $\eta_{off\_chip} = LUT_{off}^{\eta}(V_{ext}, V_{out,off}, I_{pkg} / N_{online,off})$
  $\eta = \eta_{on\_chip} \eta_{pkg} \eta_{off\_chip}$
Return $\{V_{out,off}, N_{online,off}\}$ with maximized $\eta$

**Algorithm 2** On-chip control algorithm $VR\_ON\_OPT$.

**Inputs:**
Workload current estimations $\vec{I}_L$ for each $T_{on}$ cycle.
Voltage sensor readings $\vec{V}_{sensor}$ for each $T_{on}$ cycle.

**for** each power domain $i$ **do**
  **if** $3\_stage\_HVR == True$ **then**
    $V_{out,on}^{(i)} = V_{dd}^{(i)} + \alpha I_L^{(i)} / I_{max}^{(i)}$
  **else**
    $V_{out,on}^{(i)} = V_{dd}^{(i)}$
  **end if**
  $I_{opt,on}^{(i)} = LUT_{on}^{Iopt}(V_{out,off}, V_{out,on}^{(i)})$
  $N_{online,on}^{(i)} = \left\lceil I_L^{(i)} / I_{opt,on}^{(i)} \right\rceil$
  **if** $MachineLearningOption == True$ **then**
    $V_{out,on}^{(i)} = MachineLearning(\vec{S}_{PDN}, \vec{V}_{sensor}^{(i)})$
  **end if**
**end for**
Return $\{\vec{V}_{out,on}, \vec{N}_{online,on}\}$

mine $N^{(i)}_{online,on}$ in lines 2-8. However, if the machine learning is enabled, the final $V^{(i)}_{out,on}$ is fine tuned by the machine learning module with the consideration of fine-grained spatial workload distribution, described next.

### 4.4 Machine Learning Enabled Adaption

One key objective of voltage regulation is to deliver power to on-die devices with ensured power integrity, e.g. without dropping the worst-case voltage from the on-chip power grids below a preset level. Power supply noise hotspots are created due to the non-uniform spatial distribution of workload onchip. To make things even worse, the locations of hotspots can shift during run-time. Such effects can significantly impact the on-die supply noise. Thus, the output voltage of each on-chip switching VR cluster, which is the final point of 2-stage voltage regulation, and also the input voltage to the distributed LDO network in the case of 3-stage HVR, shall be adapted with the considerations of fine-grained spatial workload distribution. However, predicting such spatial workload distribution for the purpose of PDN adaptation is a challenging problem.

Recently, machine learning (ML) has been received a significant amount of interest for power system design. For instance, noise-sensor-based machine learning techniques [59, 60] have been developed to detect voltage emergencies within functional blocks. Different from these works, we leverage machine learning to directly learn the optimal control policy based on the fine-grained spatial workload distribution predicted from a small number of distributed voltage-noise sensors. This enables a very desirable end-to-end ML solution that can lead to additional energy and power integrity benefits.

#### 4.4.1 Machine Learning Problem Formulation

We first formulate the machine learning problem. For a power domain, denote the output voltage of the corresponding on-chip switching VR (buck) cluster $V_{out,on}$. By exploiting the correlation between voltage droops at different nodes in the power grids (including sensor locations) and the distribution of workload, a machine learning model can directly learn the optimal control variable $V^{opt}_{out,on}$ using the voltage sensor readings as input features. Here $V^{opt}_{out,on}$ is defined as the minimum

$V_{out,on}$ value such that the worst-case supply voltage across the entire power grids doesn't fall below a preset safety voltage level. By leveraging the fine-grained spatial information of workload distribution, $V_{out,on}$ can be set in a more accurate way, achieving improved power efficiency and quality. A ML model is used to learn the following mapping:

$$\vec{S}_{PDN}, \vec{V}_{sensor} \rightarrow V^{opt}_{out,on}, \tag{4.6}$$

where $\vec{V}_{sensor}$ is the worst-case voltage values sensed by the voltage sensors during an on-chip VR control cycle $T_{on}$. $\vec{S}_{PDN}$ includes the PDN configurations such as control variables under which the voltage sensor values are measured. The training samples can be collected by circuit simulation by sweeping $V_{out,on}$ within a certain range to obtain the target $V^{opt}_{out,on}$ under the same workload. Fig. 4.10 illustrates the ML module and voltage sensors in a power domain.

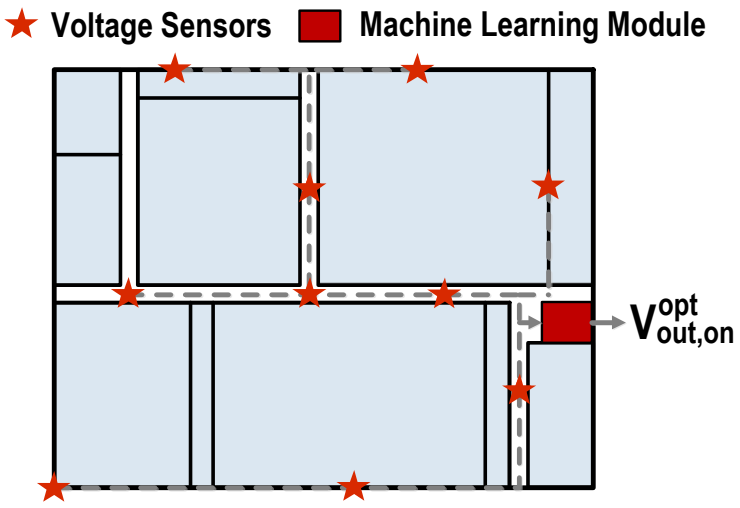


Figure 4.10: Demonstration of machine learning module and voltage sensors.

### 4.4.2 Preliminary of SRKM

We integrate our machine learning module (accelerator) on chip to enable fast real-time workload aware adaption. Such machine learning module must come with sufficient accuracy, low

area/power overhead, and should incur low processing latency to enable HVR adaptation at fine temporal granularity. The recently developed sparse relevance kernel machine (SRKM) [61, 62] is shown to offer great learning accuracy for a variety of applications. Unlike the widely adopted support vector machine (SVM) and relevance kernel machine (RVM), SRKM can lead to significant *sparsity* in both (training) sample and feature space, resulting in lower processing latency and hardware overhead. Therefore, we adopt SRKM as the machine learning algorithm.

A single training sample of the SRKM can be defined as a pair $\{\boldsymbol{x_i}, y_i\}$, where the vector $\boldsymbol{x_i}$ is the input containing $F$ features and $y_i$ is the output. Given $N$ samples, the objective of the SRKM training is to derive the model mapping $\boldsymbol{x} \rightarrow y$. The SRKM training process only selects a certain number of relevance samples ($N_S$) and features ($F_S$), producing the model consists of a sample weight vector $\boldsymbol{w} = \{\boldsymbol{w}(1), \boldsymbol{w}(2), ..., \boldsymbol{w}(N_S)\}$, a feature weight vector $\boldsymbol{v} = \{\boldsymbol{v}(1), \boldsymbol{v}(2), ..., \boldsymbol{v}(F_S)\}$, and a $N_S \times F_S$ shrinked training matrix $\boldsymbol{X}$.

Based on the trained model, the SRKM use the following equation to predict the expectation of the output $y$ for an input vector $\boldsymbol{x}$:

$$y(\boldsymbol{x}, \boldsymbol{w}) = \sum_{i=1}^{N_S} \boldsymbol{w}(i) \cdot K(\boldsymbol{x}, \boldsymbol{X_i}), \tag{4.7}$$

where $\boldsymbol{X_i}$ is the $i$-th relevance vector in $\boldsymbol{X}$ and $K(\boldsymbol{x}, \boldsymbol{X_i})$ is the kernel function computing the similarity between the input vector $\boldsymbol{x}$ and the training vector $\boldsymbol{X_i}$:

$$K(\boldsymbol{x}, \boldsymbol{X_i}) = \sum_{k=1}^{F_S} v(k) \cdot \phi(\boldsymbol{x}(k), \boldsymbol{X_i}(k)), \tag{4.8}$$

where $\boldsymbol{x}(k)$ is the $k$-th feature of the input vector $\boldsymbol{x}$, $\boldsymbol{X_i}(k)$ stands for the $k$-th feature of the $i$-th sample in the shrinked training set $\boldsymbol{X}$, and $\phi(\boldsymbol{x}(k), \boldsymbol{X_i}(k))$ measures the similarity of $\boldsymbol{x}(k)$ and $\boldsymbol{X_i}(k)$:

$$\phi(\boldsymbol{x}(k), \boldsymbol{X_i}(k)) = e^{-\gamma(\boldsymbol{x}(k) - \boldsymbol{X_i}(k))^2}. \tag{4.9}$$

In our work, we train an SRKM model offline based on 2,000 training samples collected from

circuit simulation. It achieves a Normalized Mean Square Error (NMSE) of 4.3e-3, demonstrating excellent prediction accuracy. As mentioned earlier, the trained SRKM model is mapped to an hardware accelerator for efficient runtime application.

### 4.4.3 Pipelined Parallel VLSI Architecture for SRKM Accelerator

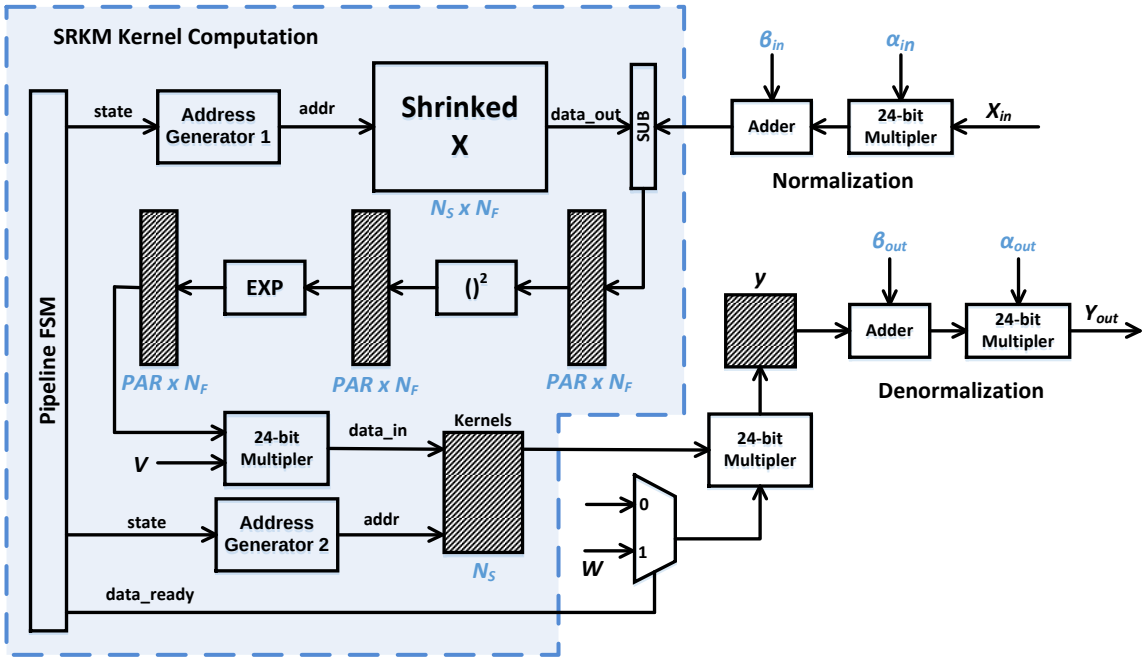


Figure 4.11: VLSI architecture of SRKM accelerator.

We further present the design of VLSI SRKM accelerator circuit and address some design decision makings such as control of parallelism and pipelining. The overall architecture of our SRKM prediction circuit is shown in Figure 4.11. The computation of the kernel $K(\boldsymbol{x}, \boldsymbol{x_i})$ is the main function of the SRKM prediction and takes the majority of the design. Besides, for practical applications, input features $\boldsymbol{x_{in}}$ with different amplitudes are usually normalized before getting into to the SRKM kernel module. The input normalization function is written as:

$$\boldsymbol{x} = \boldsymbol{x_{in}} \cdot diag(\boldsymbol{\alpha_{in}}) + \boldsymbol{\beta_{in}}, \tag{4.10}$$

where $\boldsymbol{x}$ is the normalized sample vector input, $\boldsymbol{\alpha_{in}}$ and $\boldsymbol{\beta_{in}}$ are the input scaling vector and offset vector, respectively. Accordingly, the normalized prediction output $y$ needs to be denormalize to the actual final prediction result $y_{out}$:

$$y_{out} = (y + \beta_{out}) \cdot \alpha_{out}, \tag{4.11}$$

where $\alpha_{out}$ and $\beta_{out}$ are the output scaling scalar and offset scalar, respectively.

As shown in Figure 4.12, the computation of the kernel $K(\boldsymbol{x}, \boldsymbol{x_i})$, which is the main function of the SRKM prediction, can be further divided into four steps: feature difference computing (*SUB*), feature difference squaring (*SQR*), exponential mapping (*EXP*) and weighted sum (*W_SUM*) for kernel finalization. Among these steps, all calculations operate on vectors and there is no data dependency among those vectors. In this case, at each step, we can process multiple vectors at a time to speed up the prediction. However, a high level of parallelism among vectors requires a large number of hardware resources, especially multipliers, which could be costly in terms of circuit area. We will compare multiple SRKM designs with different configurations of the parallelism level in terms of area, power, and latency in the following subsection.

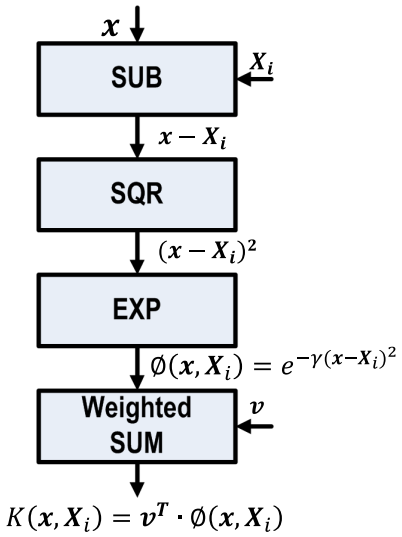


Figure 4.12: Decomposition of the SKRM kernel computation. Equations shown here are for element-wise operation.

Besides, the four steps of kernel computation are processed in pipeline to further speed up the SRKM prediction as well as improve the cost efficiency of the hardware implementation. The flow diagram of the kernel computation pipeline is depicted in Figure 4.13. The number of sample vectors processed at each clock cycle is decided by the configured parallelism level denoted as $PAR$. With the proposed pipeline scheme, the total number of clock cycles $N_c$ that needed to finish the kernel computation is

$$N_c = \frac{N_S}{PAR} + 3, \tag{4.12}$$

where $N_S$ is number of vectors in the shrinked data set. In this design, different pipeline stages have the same parallelism level, thus the same $PAR$.

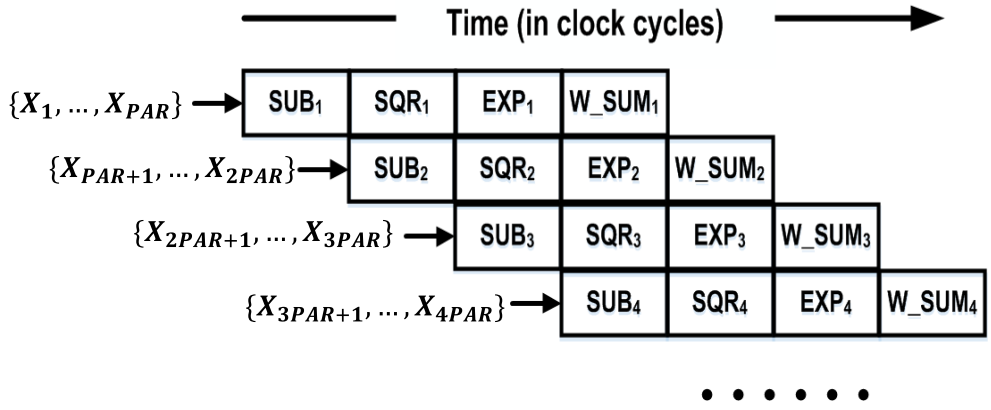


Figure 4.13: Flow diagram of the 4-stage SKRM kernel computation pipeline. The subscript number index the group of sample vectors that are currently under this stage.

When implementing the pipelined SRKM prediction, each of the four stages are realized by a corresponding functional module, which consists of only combinational logic thus can be finished within one clock cycle. Besides, the output of each pipeline stage is registered, which are indicated by the shaded blocks in Figure 4.11. A finite state machine (FSM) is used to control the state transaction and the data transfer. It also controls the two address generators to generate the address for reading the sample vector values in *SUB* stage and storing the kernel values in the *W_SUM* stage based on the current clock cycle count.

Then, in order to get $\phi(\boldsymbol{x}(k), \boldsymbol{x_i}(k))$ in (4.9), the square calculation (*SQR*) is carried out using multipliers and the exponential calculation (*EXP*) follows. We adopt a 1024-entry LUT to efficiently implement the exponential function $e^{-\gamma \cdot x}$ in equation (4.9) in our VLSI SRKM prediction circuit. At last, the kernel function value for the input vector is obtained by computing the weighted sum (*W_SUM*) of differences between input vectors and relevance vectors over all features. The resulting kernel function values (i.e. $data\_in$ in Figure 4.11) are stored in the kernel register array whose depth is the number of input vectors. When the kernel function value corresponding to a certain input vector is available, it will be written to the targeted address in the register array.

After all kernel function values have been updated in kernel registers, a $data\_ready$ control signal will assert. This control signal enable computation of the normalized prediction output $y$ in which the kernel register output vector is multiplied with the pre-trained sample weight vectors $\boldsymbol{w}$. The entire flow is finalized by denormalizing $y$ to the predicted target $y_{out}$.

### 4.4.4 Hardware Tradeoffs

Several SRKM accelerator using the proposed architecture are designed with Verilog HDL. The hardware designs are synthesized with a 45nm CMOS standard cell library using Synopsys Design Compiler under 167-MHz clock frequency. Cadence Innovus is used to perform floorplanning, placement, and routing. We report the area, static analysis power and prediction latency of hardware SRKM accelerator as below.

Table 4.3: Hardware result for SRKM accelerator with different parallelism.

| Parallelism | Area ($mm^2$) | Power (mW) | Latency (ns) | Energy (nJ) | Speedup (X) |
|---|---|---|---|---|---|
| 1 | 0.545 | 26.03 | 393 | 10.23 | 21.1 |
| 2 | 0.844 | 43.97 | 219 | 9.63 | 37.9 |
| 4 | 1.418 | 74.53 | 132 | 9.84 | 62.9 |
| 8 | 2.571 | 118.04 | 90 | 10.62 | 91.9 |

To explore the tradeoffs between area, power and prediction latency impacted by different

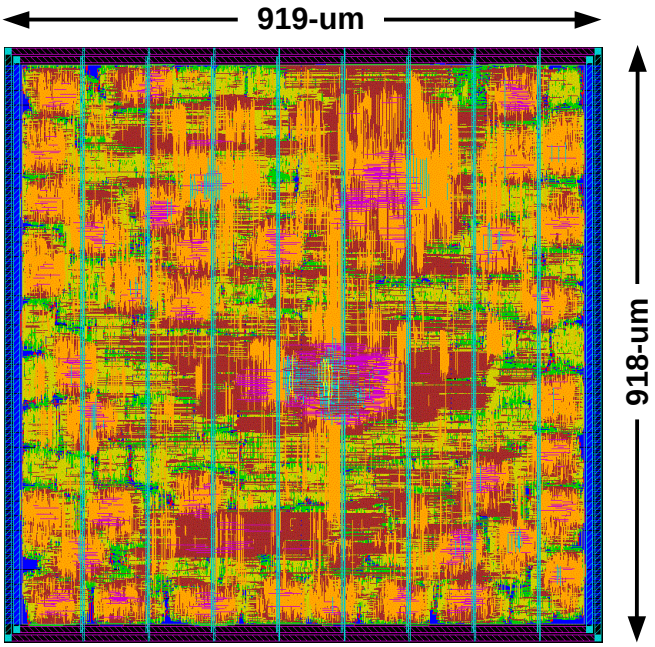


Figure 4.14: Layout of SRKM predictor with parallel parameter PAR=2.

levels of parallelism, we implement four VLSI designs with parallel parameter $PAR = 1, 2, 4, 8$. We set the number of relevance vectors ($F_S$) to be 100 and the number of relevance features ($N_S$) to be 11 considering the number of voltage sensors in each power domain. The main hardware results are summarized in Table 4.3, demonstrating good tradeoff between processing latency, power and area overhead. For example, the estimated power and area of an SRKM design with parallelism degree four are 74.53$mW$ and 1.418$mm^2$, respectively. By using the hardware SRKM accelerator circuit, we can achieve up to 91.9X speedup compared to the software SRKM running on an Intel Xeon E5-2697A Processor and the processing latency is only 90$ns$. The achieved low latency guarantees the effectiveness of applying the control algorithm in Section 4.3 at a fine temporal control granularity. It is also interesting to note that the consumed energy per prediction (defined as the product of the power and latency) are kept in the similar level for different parallelism. The layout of the SRKM hardware with 2-way parallelism is shown in Fig. 4.14.

## 4.5 Experimental Evaluations

### 4.5.1 Experimental Setup

#### *4.5.1.1 Multi-Core Processor Model and Power Analysis*

We use the full-system multi-core simulator GEM5 [63] to generate run-time statistics with the granularity of 100 ns and then feed them into the power analysis tool McPAT [64] to produce realistic workload current traces. The 45nm four-core processor model of Table 4.4 is evaluated using the PARSEC benchmark suite [65]. The total core area estimated by McPAT is 211.4$mm^2$ and the peak workload current per core is 25A. As illustrated in Fig. 4.15, each CPU core is divided into five blocks according to their functionality: IFU (instruction fetch unit), RU (renaming unit), LSU (load/store unit), MMU (memory management unit) and EXEU (execution unit). Those five blocks are further divided into eleven sub-blocks. The current workload of each block, derived from McPAT, is evenly distributed within the block to load the power delivery network (PDN).

Table 4.4: Processor configuration.

| # Cores | 4 | Vdd | 1V |
|---|---|---|---|
| Frequency | 1.8GHz@45$nm$ | $I_{max}$ | 25A |
| Branch Predictor | 2K entries | Core area | 40.4$mm^2$ |
| ALU/MUL/FPU | 6/2/6 | I/D-TLB | 48/64 entires |
| Load/Store buffer | 32 | ROB size | 192 |
| L1 I/D-Cache | 32KB, 2-way, 2-cycle latency | L2 Cache | Shared 2MB, 20-cycle latency |

#### *4.5.1.2 Power Delivery Network*

To enable the comparison across different PDN architectures, we consider the widely used 2-stage PDN in Fig. 4.1(b) with on/off-chip buck VRs as the reference. The main structure of the reference system is similar to the 3-stage HVR except that the centered on-chip buck converters are used as the last voltage processing stage instead of the distributed LDO network. We adopt

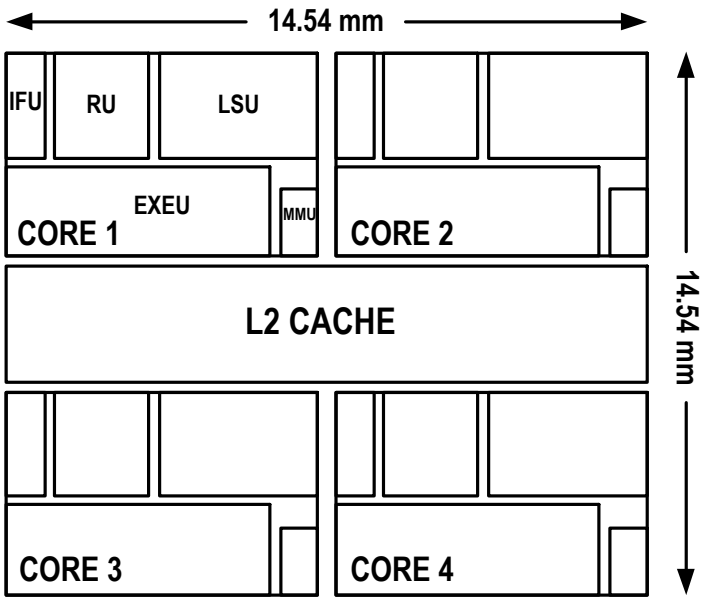


Figure 4.15: Floor plan of a 4-core processor.

a PCB/package model similar to [53] for both PDNs. The on-chip power grids of the PDN are modeled using an RC network with more than 3,000 nodes.

In the regulation chain of each PDN, a cluster of 5 off-chip buck VRs is used to drive 5 on-chip buck clusters with each cluster containing 4 identical on-chip VRs. In 3-stage HVR, each on-chip VR cluster further drives a network of 250 on-chip LDOs for each core (power domain). The topology from [1] is adopted for on-chip LDOs with maximum 100mA load capability. The off-chip and on-chip buck converters are designed using PowerSoC [56], which finds the key design parameters such as switching frequency, filter inductance, and size of MOS switches under a static nominal load condition. Considering the on-chip buck converters are the final regulation stage in the 2-stage reference PDN, they are designed with more emphasis on regulation performance at the cost of more energy loss. As a result, the on-chip buck VRs of the 2-stage PDN operate at 291MHz, while those of the 3-stage HVR operate at 107MHz. For fair comparisons, the total area budget of on-chip VRs (including LDOs) is set to $15mm^2$ for all PDNs.

#### 4.5.1.3 *Control Scheme Setup*

The on- and off-chip VR control periods $T_{on}$ and $T_{off}$ are set to 1us and 100 us, respectively to suit the response times of the considered on- and off-chip switching VRs.

As shown in Algorithm 2, the machine learning enabled control scheme takes the voltage sensor readings as input to predict the optimal output voltage of on-chip switching VRs. However, obtaining the voltage sensor readings for each PARSEC benchmark during runtime through the simulation of our complex PDN model is prohibitively computationally expensive. To speed up the evaluation process, we once again leverage machine learning but for the purpose of fast estimation of voltage sensor readings. We train another SRKM model offline which performs the following mapping:

$$\vec{S}_{PDN}, \vec{I}_{block}(n), \vec{I}_{block}(n-1) \rightarrow \vec{V}_{sensor}(n), \tag{4.13}$$

where $\vec{I}_{block}(n)$ and $\vec{I}_{block}(n-1)$ are the block-level workloads at the current and past 100ns time steps, representing the fine-grained workload transition, $\vec{V}_{sensor}(n)$ is the worst-case voltage sensor readings caused by the corresponding transitions. Based on the traces of $\vec{V}_{sensor}(n)$, the worst case voltage sensor readings for each control cycle $T_{on}$ can be computed as the input to the machine learning module. Similar to the online SRKM module in Section 4.4, the PDN state variables $\vec{S}_{PDN}$ are included as part of the input features for this off-line SRKM model to estimate $\vec{V}_{sensor}(n)$. Trained with 4,000 samples, this off-line SRKM model is very accurate and achieves an average NMSE of 1.52e-4.

### 4.5.2 Online Machine Learning Overhead

The area and power overhead of the proposed machine learning (ML) enabled HVR adaptation comes from the voltage sensors and SRKM accelerators. The voltage sensors can be implemented based on low-power high-speed analog-to-digital converters (ADCs) [66, 67]. The ADC design in [67] is considered to estimate the sensor cost. In our study, 10 voltage sensors and a compact SRKM accelerator (parallelism parameter equaling 4) are placed in each core. As summarized in Table 4.5, the proposed ML approach only incurs an overhead of 2.5% on area and 0.2% on power

but comes with great benefits.

Table 4.5: Additional area and power overhead (%). Area is normalized to the original on-die area. Power is normalized to half of the peak power.

| | Amount | Area overhead | Avg. power overhead |
|---|---|---|---|
| Voltage sensors | 40 | 0.442% | 0.166% |
| SRKM accelerators | 4 | 2.090% | 0.063% |
| Total | - | 2.532% | 0.229% |

### 4.5.3 Evaluation of Ideal Control Schemes

Firstly, intuitive study is conducted to shed some light on the possible energy saving brought by the adaptive 3-stage HVR system through joint optimization of the control variables. To do so, we evaluate four different ideal control schemes where important factors such as multiple control granularities and on-die workload distribution are not considered at this moment. They are covered in the presented overall control scheme and will be evaluated in the following subsections.

As shown in Fig. 4.16(a), four ideal schemes include a static reference with all control variables fixed and several adaptive ones considering different number of control variables are evaluated. For example, the last scheme $N_{online,on/off} + V_{out,on/off}$ means that it enables the tuning of the number of online VRs, output voltage of both on- and off-chip switching VRs. Since the purpose of this study is to demonstrate the potential of adaptive HVR system for energy reduction, we simplify the load condition and assume that each power domain consumes the same amount of power. And at each workload point, the input voltage of each LDO network $V_{out,on}$ is set according (4.3) based on LDO's characteristics while all the other variables are optimized via exhaustively comparing the overall power efficiency for all combinations of control variables. Then Fig. 4.16(a) plots the power efficiency of each scheme as a function of normalized workload. It illustrates that the first static scheme achieves the lowest overall efficiency among all load conditions. $N_{online,on}$ and $N_{online,off}$ are optimized in the second scheme with evident energy improvement under light load

condition. The third scheme with additional consideration of $V_{out,on}$ demonstrates the effectiveness of tuning LDOs' input voltage for further energy saving. The last scheme, with all control variables considered, further increases the energy efficiency by better trading-off the power losses among all components.

To see how the optimal control variables vary with the change of workload, Fig. 4.16(b) further plots these variables as functions of workload for the last control scheme. It is observed that in general, the increasing workload requires more buck converters at both off-chip and on-chip stages to go online. Besides, the output of off-chip converter $V_{out,off}$ tends to be lowered at light load to reduce the switching loss of integrated DC-DC converters while increase at heavy load to reduce the package-through current and the corresponding $I^2R$ power loss. However, it is interesting to note that every time when more on-chip buck converters are activated, $V_{out,off}$ will suddenly steps back to compensate such increased switching power loss, demonstrating that the overall power efficiency may only be achieved by considering the interaction of all VR stages.

### 4.5.4 Power Integrity and Adaptive Overall Control

#### *4.5.4.1 Power Integrity*

Next we evaluate the overall control scheme presented in Section 4.3 with the considerations of spatial workload distribution and multiple control granularities. We first examine the power quality of several adaptive PDNs through detailed circuit-level simulation. Verilog-A models with PWM control are used to model the on-chip buck converters based on design parameters obtained from PowerSoC. Ideal voltage source is used for the off-chip VRs since they have little impact on power supply noise. The complexlity of our PDN model with large number of VRs causes significant simulation challenge. It takes around 112 hours to simulate a 100us segment of benchmark workload with 4 threads on a Intel Xeon E5-2697A processor @2.60GHz. We select a 100us workload segment from each PARSEC benchmark, forming a workload simulation set. This set contains a representative worst-case workload segment from the *fluidanimate* benchmark and random segments from other benchmarks, serving as typical workload conditions. As described in

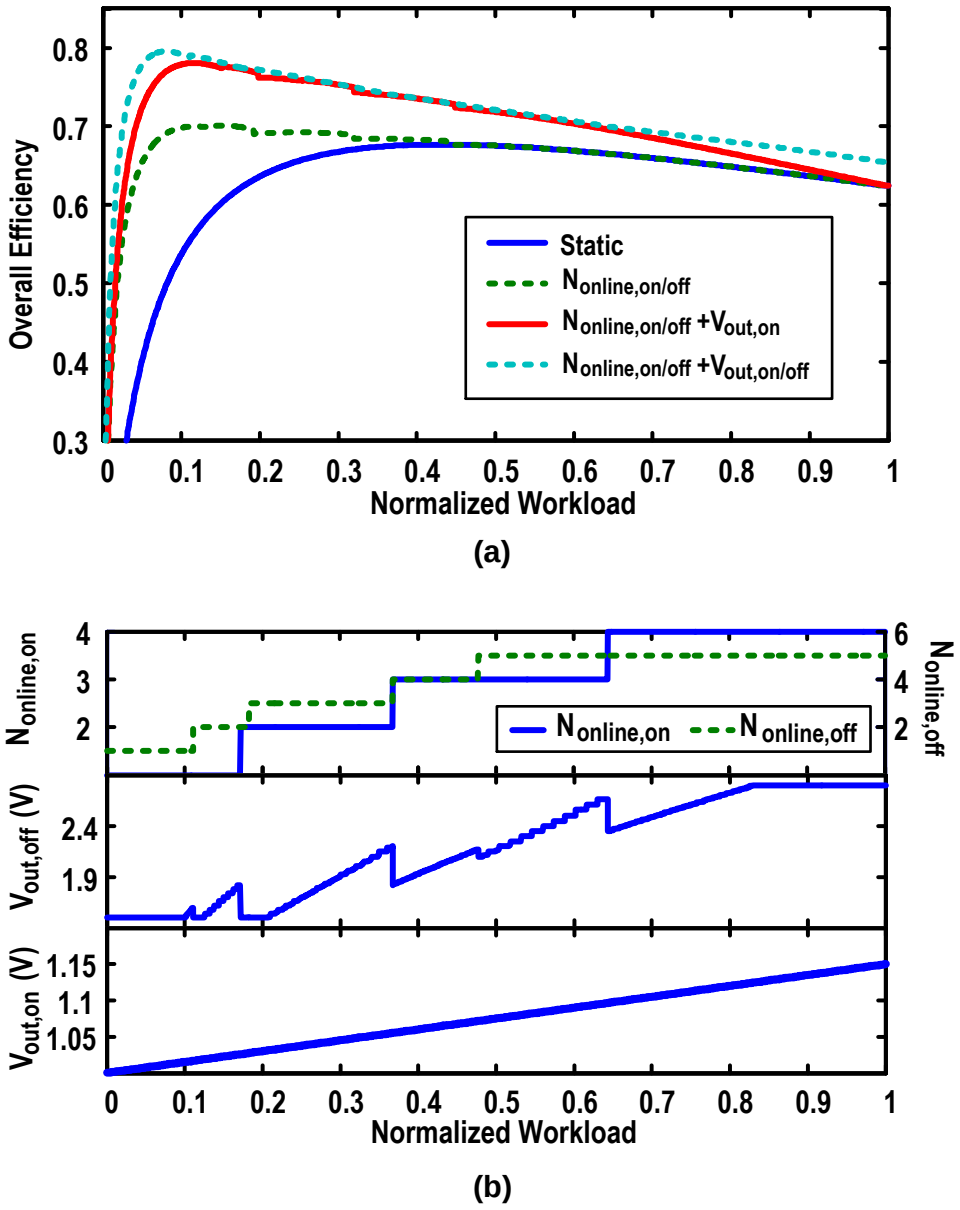


Figure 4.16: (a) Power efficiencies of four different ideal control schemes, and (b) optimal control variables versus workload current.

Section 4.3, our control algorithm supports both 2-stage and 3-stage PDNs and also provides two options with and without machine learning module. This creates four adaptive PDNs and they are simulated based on the aforementioned workload simulation set.

To avoid using large supply voltage guard bands, we allow rare occurrences of voltage emergencies (VE), i.e. supply voltage drops below a preset level. We assume that the processor under study is equipped with a common fail-safe mechanism such as the rolling-back recovery [68] or adaptive frequency tuning [69] in the events of rare VE. In this study, voltage emergency (VE) is considered to occur when the maximum voltage droop on the on-chip power grids exceeds 10% of the nominal Vdd. The count of VEs is used as a metric for power quality.

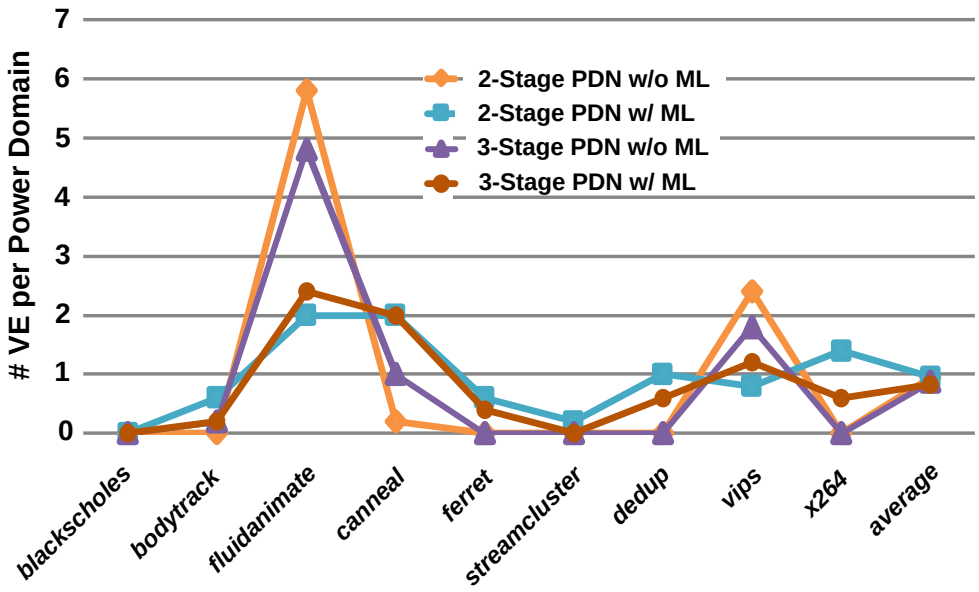


Figure 4.17: Number of VEs per benchmark segment.

Fig. 4.17 plots the average count of VEs per power domain under the workload segment of different PARSEC benchmarks. On average VE only occurs about once in each power domain for all PDNs. In other words, all PDNs have the same power integrity level. Under this equal power quality condition, we will compare these PDNs in terms of energy efficiency in Section 4.5.5.

#### 4.5.4.2 Case Study for Adaptive Control

Next, we use two simulation examples of the adaptive 3- and 2-stage PDN systems with and without machine learning module to shed some light on how the proposed control polices adapt to the workload change and the benefits brought by machine learning. Fig. 4.18 shows the transient waveforms based on *fluidanimate*. Such workload segment represents a worst-case scenario since the total load current suddenly increases to the maximum 25A peak current from light load condition. The fast and large load variations as such tend to cause a considerable amount of power supply noise, imposing a significant regulation challenge. The resulting worst voltage $V_{PG}$ in the entire on-chip power grids is plotted. The dash line indicates the supply voltage level under which VE is considered to happen. It can been seen that for both PDN architecture, the system armored with machine learning (ML) can more accurately set the output voltage of the on-chip buck converters $V_{out}$. That is, the $V_{out}$ that is further fine tuned by the proposed ML module becomes lower under lighter load conditions, reducing the energy at the corresponding stage. On the other hand, $V_{out}$ can be quickly increased in response to the arrival of heavier workloads. The number of online on-chip buck converters $N_{online}$ is also well adapted to the workload variation for energy saving.

Fig. 4.19 shows a more typical workload example from the *streamcluster* benchmark. The corresponding power trace exhibits periodic behavior resulted from a for loop in the program. Although no VE happens in all PDNs, it is evident that for both PDN architectures, the machine learning solution further improves energy efficiency due to lower values of $V_{out}$.

### 4.5.5 Overall Energy Evaluation

#### 4.5.5.1 Energy Comparison

The overall energy efficiencies of different PDN architectures with various control schemes are compared. We name all considered PDNs in the top of Fig. 4.20. There are four 3-stage PDNs denoted by 3-S1 to 3-S4 with different control policies. Take the 3-S4 PDN with configuration $3stage\text{-}N_{online,on/off}\text{-}V_{out,on/off}(ML)$ for example. The configuration means that the system utilizes a 3-stage HVR PDN architecture, enables the tuning of the number of online VRs, output

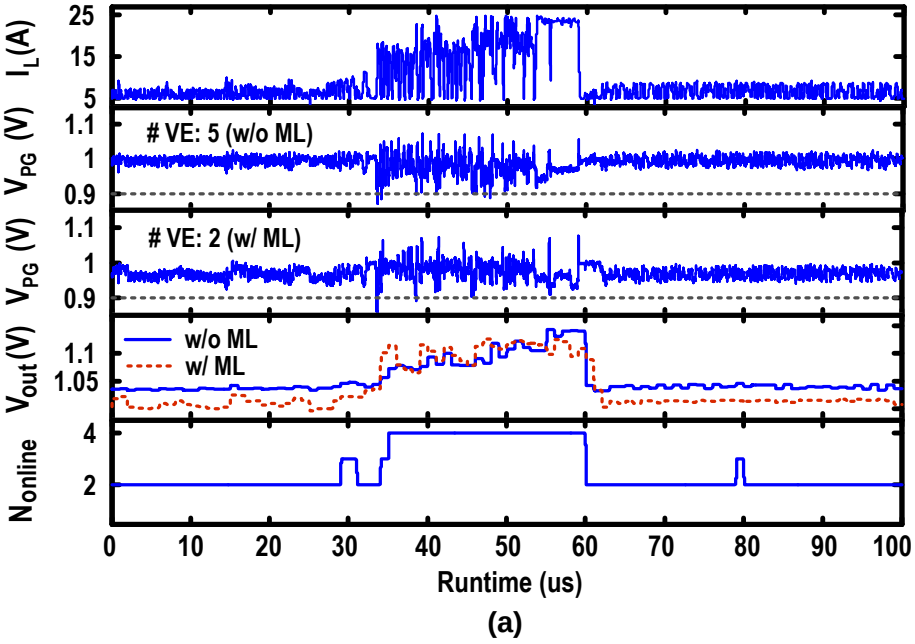


(a)

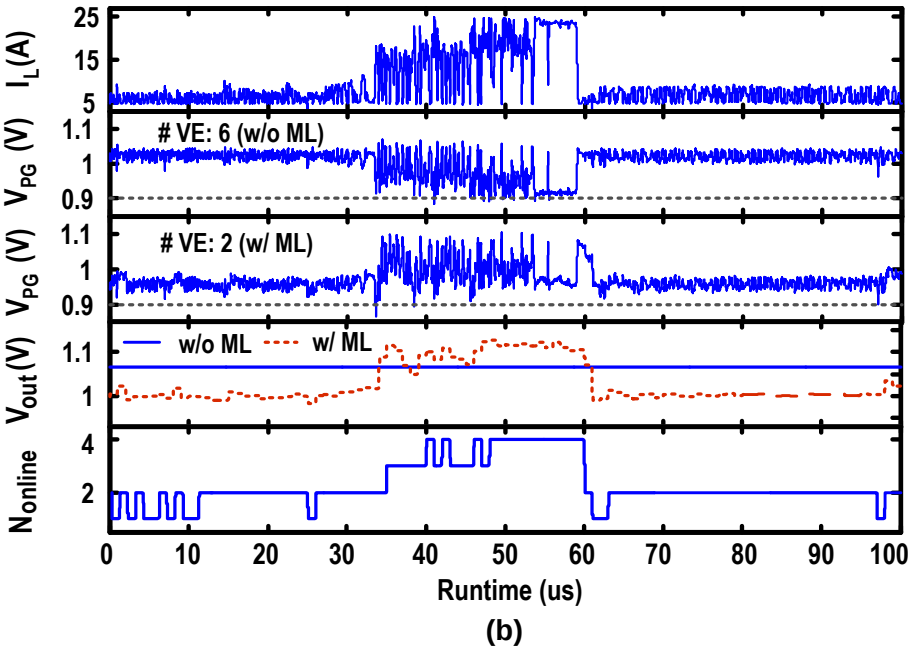


(b)

Figure 4.18: Transient waveforms of fluidanimate benchmark for (a) 3-stage HVR, and (b) 2-stage PDN.

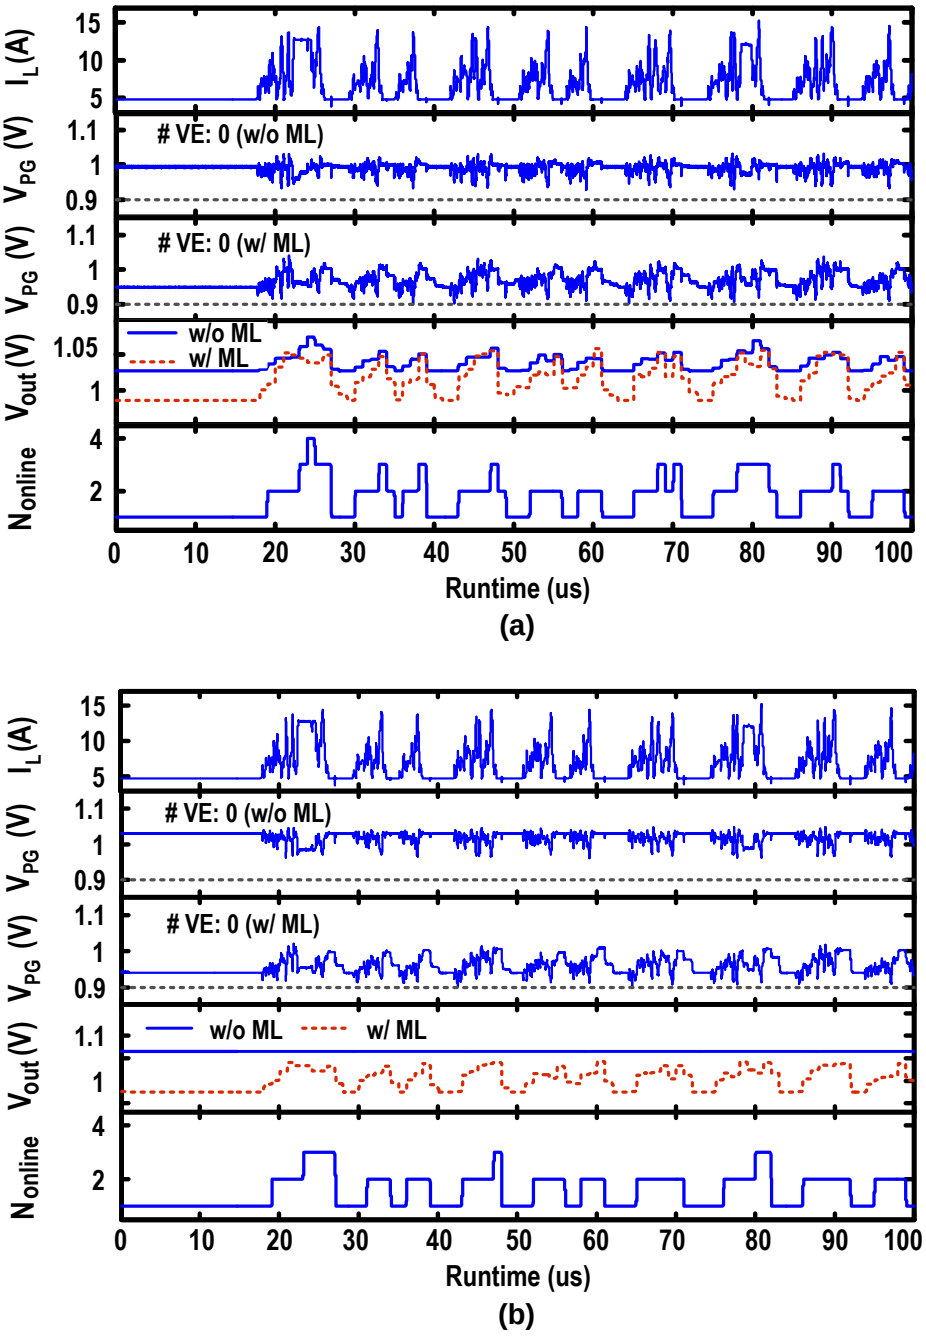


Figure 4.19: Transient waveforms of streamcluster benchmark for (a) 3-stage HVR, and (b) 2-stage PDN.

voltage of both on- and off-chip switching VRs, and it integrates the machine learning module. The first 3-S1 system indicates a static 3-stage PDN with no runtime adaptiaton. Similarly, we have four different PDNs with a 2-stage architecture. We highlight several observations from Fig. 4.20.

- Without any adaptive control, the static 3-S1 outperforms 2-S1 with an energy reduction of 4.0% on average, demonstrating the potential of leveraging HVR for improved energy and performance tradeoffs.

- 2-S2 adopts a simple control scheme similar to [43] by tuning the number of online on/off-chip buck VRs in the 2-stage PDN. It is observed that it reduces the energy by 8.5% over the static 2-S1. However, Adding $V_{out,off}$ into 2-S3 can bring in an additional 2.1% energy saving on average, since such scheme captures more interdependency among the regulation chain. By comparing 2-S4 with 2-S3, the proposed machine learning module offers upto 4.1% reduction of system energy by utilizing the spatial workload distribution information.

- The highest energy efficiency is achieved by the proposed machine learning enabled adaptive 3-S4 system. The 3-S4 system reduces the total system energy dissipation by upto 17.9% and 12.2% on average compared to the static 3-S1. Compared with the conventional static 2-S1, our 3-S4 with runtime control reduces system energy by upto 23.9% and 15.7% on average.

Fig. 4.21 further decomposes the energy consumption for 2-S1, 2-S4, 3-S1 and 3-S4 systems. It is observed that in general the processor in the 3-stage HVR consumes less energy compared to that of the 2-stage PDN. That is because the distributed LDO network enhances the supply noise suppression and thus enables lower supply voltage while maintaining the same power integrity, demonstrating the benefit of HVR in voltage regulation. By setting the output voltage of on-chip buck VRs in a more accurate way, the use of machine learning module significantly improves the LDO's power efficiency in the 3-stage HVR system while reducing the processor's energy consumption in the 2-stage system. With full consideration of the energy interdependency in the

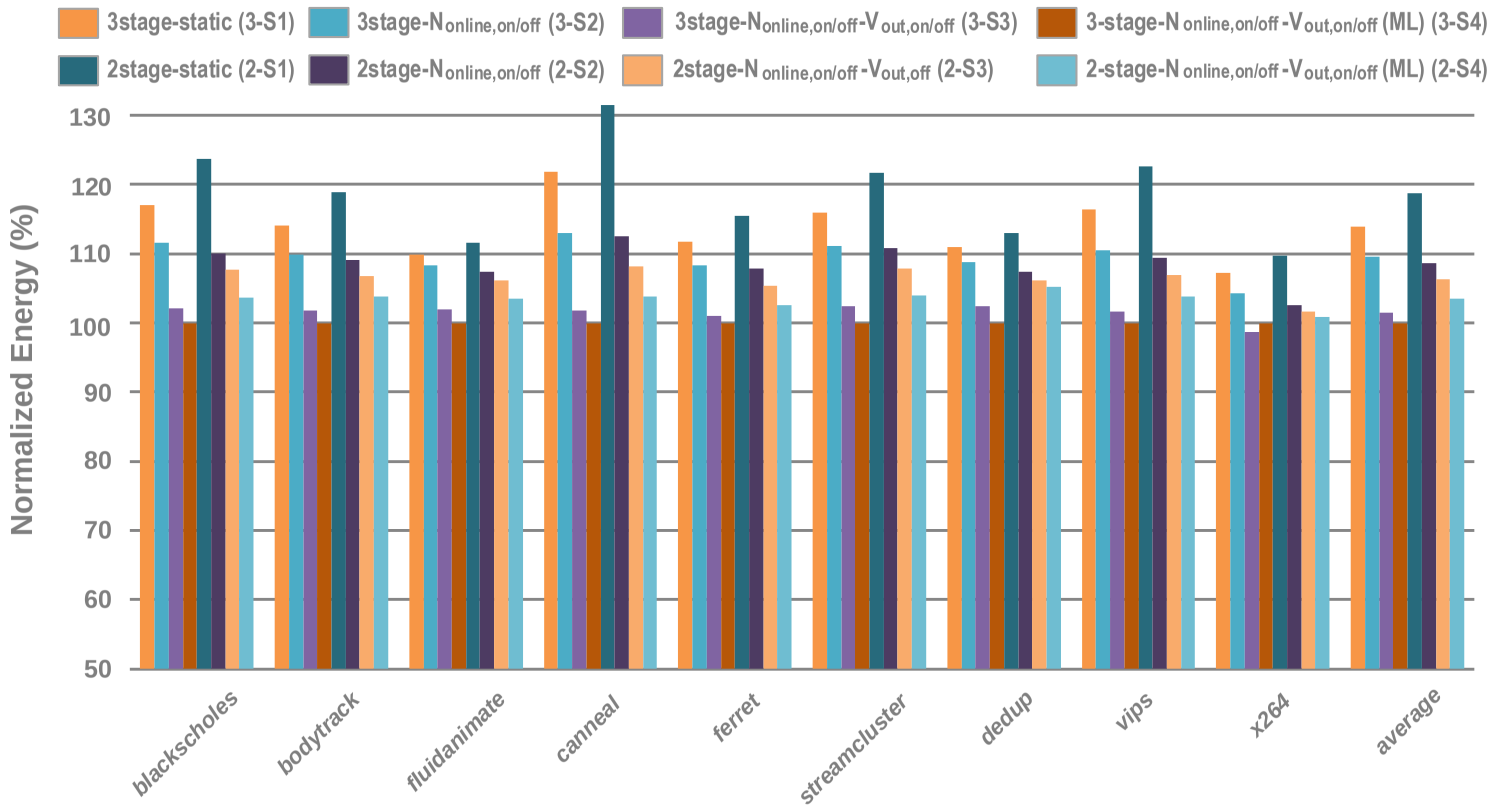


Figure 4.20: Overall energy estimation for different PDN designs.

regulation chain, the proposed control policy achieves a near-optimal overall power efficiency by carefully trading-off power loss at different stages.

#### *4.5.5.2 Impact of Control Granularity*

As discussed earlier, great benefits of adaptive control may be achieved at the finest possible temporal granularity by tracking the workload more closely. To demonstrate it, Fig. 4.22(a) shows the corresponding power loss increments for the 3-S4 system by applying coarser on-chip control granularities. Enlarging $T_{on}$ from 1us to 10us and 100us, the total power loss increases by up to 5% and 10%, respectively, demonstrating the benefits of fine-grained adaptive control. However, it is observed in Fig. 4.22(b) that even with a coarser $T_{on}$, significant power reduction can still be achieved over the static 3-S1 system, demonstrating the effectiveness of the proposed adaptive HVR over a wide range of control granularity.

## 4.6 Summary

Targeting multi-stage heterogeneous voltage regulation (HVR) systems, this chapter develops comprehensive workload-aware control policies acting at multiple temporal granularities based on complimentary characteristics of on-chip and off-chip VRs. The considered control variables are jointly optimized to improve the overall power efficiency according to important interdependencies existing in the regulation chain. Our control policies are further supported with an integrated machine-learning module to cope with fine-grained spatial distributions of workload, achieving further improved power quality and efficiency. We show that the proposed adaptive HVR and control policies reduce system energy by up to 17.9% and 23.9% over a static 3-stage HVR and conventional 2-stage PDN, respectively.

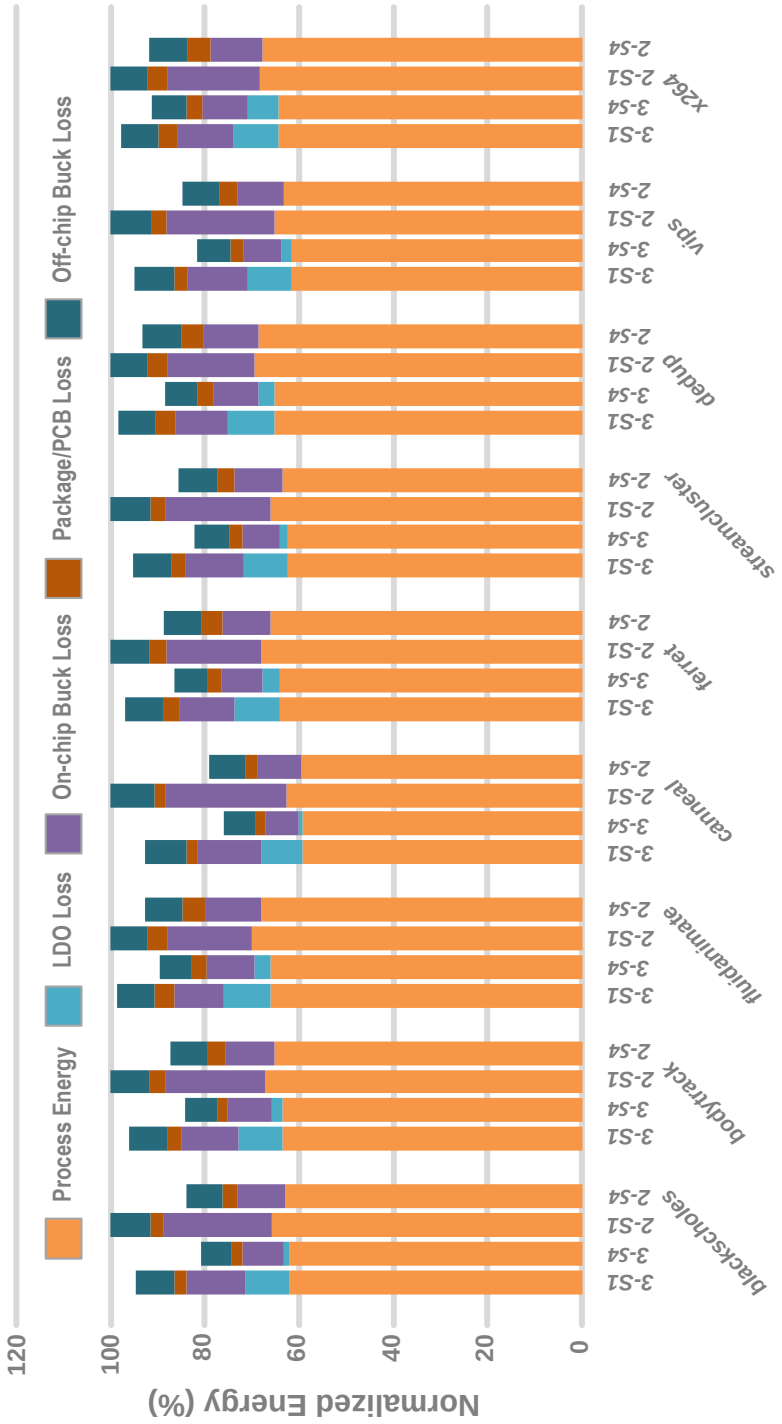


Figure 4.21: Detailed energy breakdown for four different PDNs.

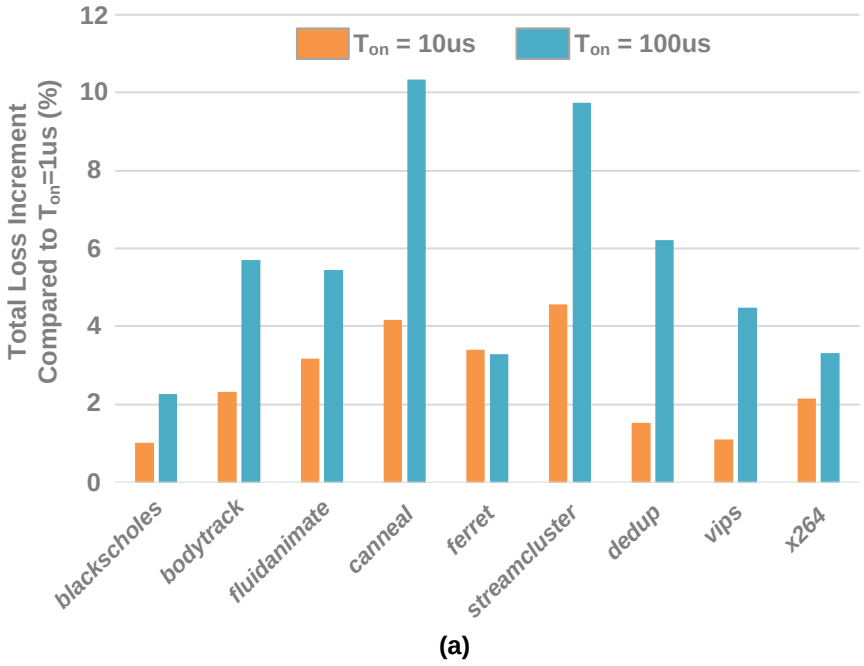


(a)

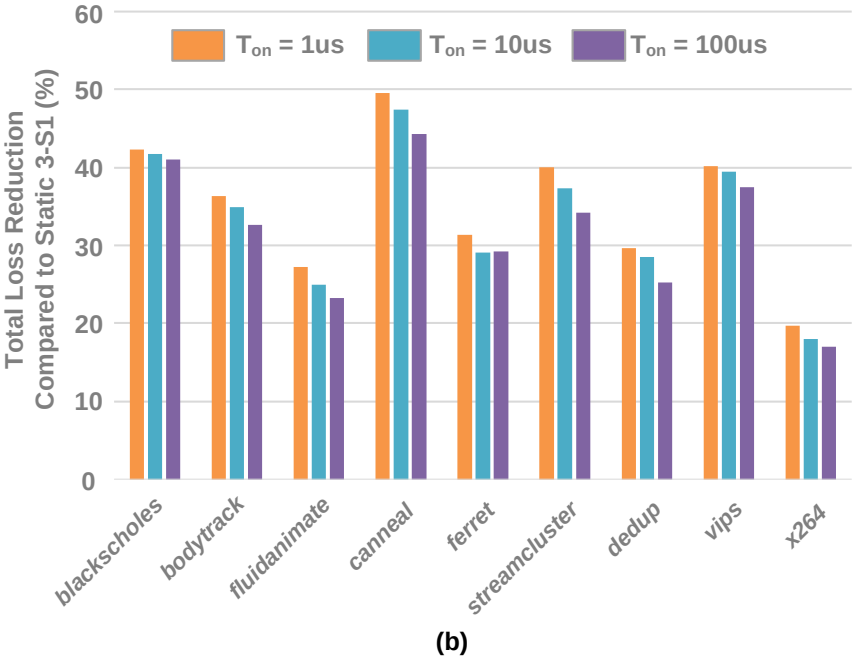


(b)

Figure 4.22: Impact of different control granularities on the power loss of 3-S4 PDN: (a) total loss increment compared to $T_{on}$=1us, and (b) total loss reduction over static 3-S1 PDN.

# 5. CONCLUSION AND FUTURE WORK

## 5.1 Conclusion

This dissertation first addresses the stability design challenge for large PDNs with distributed voltage regulation, which is caused by complicated interactions between multiple voltage regulators and the passive network of surrounding RLC parasitics. Although the recently developed hybrid stability theorem (HST) is promising to deal with the stability issue of distributed regulation by efficiently capturing all the effects of interactions, the intrinsic conservativeness can lead to large amount of pessimism in stability assessment and therefore cause significant overdesign. To remove such overdesign as much as possible and enable the highly desirable localized regulator design with guaranteed network-wide stability, this dissertation first identifies the great conservativeness reduction opportunities brought by appropriate system partitioning. Then we propose a frequency-dependent bidirectional admittance splitting technique to re-partition the system to largely reduce the pessimisms in stability analysis. By systematically exploring the theoretical foundation of the HST framework, we recognize all the critical constraints under which the partitioning technique can be performed rigorously to remove conservativeness while maintaining key theoretical properties of the repartitioned subsystems. Based on that, an efficient and localized stability-ensuring automatic design flow is developed for large power delivery systems with distributed on-chip regulation. Comprehensive design studies show that the proposed approach achieves significant performance gain over both the traditional phase margin based design approach and a reference hybrid stability approach with fixed system partitioning. In use of the proposed approach, we further discover new design insights targeting improved system tradeoffs between stability and performances by comprehensively exploring the vast design space of PDNs with distributed regulation.

Besides stability, modern PDN must adapt itself according to the workload variation and optimize the tradeoffs between power efficiency and quality during runtime. In this dissertation, we

argue that the ultimate quality and efficiency in supply voltage regulation may be only achieved by fully exploiting the heterogeneity in PDN architecture with heterogeneous regulators with complimentary characteristics in response time, power efficiency, and cost. Based on a proposed multi-stage heterogeneous voltage regulation (HVR) architecture, this dissertation aims to answer the following key question for the first time. Given a desired power supply voltage set by a higher-level power management policy, how shall the voltage regulators in the HVR system be adapted autonomously with respect to workload change at multiple temporal scales to significantly improve system power efficiency while providing a guarantee for power integrity? Our solution is a systematic workload-aware HVR control scheme which can jointly optimize power efficiencies of all voltage processing stages to maximize the overall system power efficiency. The system power loss is minimized via considering interdependencies across the entire voltage processing chain and adapting HVR system at multiple temporal scales given the significantly different VR response times. Besides, we further propose an integrated machine learning solution to cope with fine-grained spatial distribution of workload, achieving further improved power quality and efficiency by tuning the control variable in a more accurate way. Such machine learning solution is enabled by a small number of voltage-noise sensors and an efficient machine learning hardware accelerator with low silicon overhead, power consumption and prediction latency. This provides an autonomous end-to-end integrated machine learning solution allowing for fine-grained adaptation of HVR. We conduct comprehensive evaluations for the proposed techniques based on PARSEC benchmarks. The experimental results show that the the proposed adaptive 3-stage HVR reduces the total system energy dissipation by upto 23.9% and 15.7% on average compared with a conventional static two-stage voltage regulation using off- and on-chip switching VRs. Compared with the 3-stage static HVR, our runtime control reduces system energy by upto 17.9% and 12.2% on average. Furthermore, the proposed machine learning prediction offers upto 4.1% reduction of system energy.

## 5.2 Future Work

In this section, we discuss several potential directions to expand the presented work.

### 5.2.1 PDN Design with Stability Assurance

The first interesting direction is to apply the proposed HST-based PDN design methodology to investigate more complicated control mechanism for voltage regulation, such as novel centralized control for distributed on-chip regulation network. In the PDN system with centralized control, all the voltage regulators can respond globally to the workload dynamics, mitigating potential issues of distributed regulation such as imbalanced load sharing. By leveraging the proposed PDN design flow, such system can be designed targeting at improved regulation performance and guaranteed system stability. This will add great value to the existing PDN design.

Besides, the application scope of the current HST-based desgin flow is mainly for linear VRs such as LDOs. This is because the HST theoretical framework is built upon linear time-invariant (LTI) systems and thus the partitioned blocks are characterized with linear models. However, it is also possible to extend the proposed design methodology to various other VR types such as DC-DC buck converter and switching-capacitor (SC) based converter with approximated linear models with reasonable good accuracy. Investigating this direction will certainly make the entire design methodology more complete and general.

### 5.2.2 Workload-Aware Power Management

In terms of workload-aware power management, all the control decisions in the current work are made based on the workload estimation for the current control cycle through voltage/current sensor network. However, the computation of control variables and execution of control policies cause considerable latency in the close-loop control scheme and thus give rise to potential concerns such as performance degradation of PDN. Therefore, there is a desire to predict the workload ahead of time, for example, through machine learning based approach leveraging the statistical characteristics of a certain workload. If successful, it will enable proactive control to further improve the benefits of adaptive PDN.

# REFERENCES


[1] 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.

[2] T. Y. Man, K. M. Leung, and C. Y. L. *et al*, "Development of single-transistor-control LDO based on flipped voltage follower for SoC," *IEEE TCAS I: Regular Papers*, vol. 55, no. 5, pp. 1392–1401, 2008.

[3] C. Zhan and W.-H. Ki, "Output-capacitor-free adaptively biased low-dropout regulator for system-on-chips," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 57, no. 5, pp. 1017–1028, 2010.

[4] S. Kose, *High performance power delivery for nanoscale integrated circuits*. University of Rochester, 2012.

[5] P. Li, "Design analysis of ic power delivery," in *Proc. IEEE/ACM Conf. Cmput.-Aided Design*, pp. 664–666, 2012.

[6] I. Vaisband and E. G. Friedman, "Heterogeneous methodology for energy efficient distribution of on-chip power supplies," *IEEE Transactions on Power Electronics*, vol. 28, no. 9, pp. 4267–4280, 2013.

[7] L. Cheng, Y. Liu, and W.-H. Ki, "4.4 a 10/30mhz wide-duty-cycle-range buck converter with dda-based type-iii compensator and fast reference-tracking responses for dvs applications," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, pp. 84–85, IEEE, 2014.

[8] 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.

[9] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core dvfs using on-chip switching regulators," in *High Performance Computer Architecture, 2008. HPCA 2008. IEEE 14th International Symposium on*, pp. 123–134, IEEE, 2008.

[10] K. K. Rangan, G.-Y. Wei, and D. Brooks, "Thread motion: Fine-grained power management for multi-core systems," in *Proceedings of the 36th Annual International Symposium on Computer Architecture*, ISCA '09, pp. 302–313, ACM, 2009.

[11] E. Rotem, A. Mendelson, R. Ginosar, and U. Weiser, "Multiple clock and voltage domains for chip multi processors," in *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pp. 459–468, ACM, 2009.

[12] C. Isci, A. Buyuktosunoglu, C.-Y. Cher, P. Bose, and M. Martonosi, "An analysis of efficient multi-core global power management policies: Maximizing performance for a given power budget," in *Proceedings of the 39th annual IEEE/ACM international symposium on microarchitecture*, pp. 347–358, IEEE Computer Society, 2006.

[13] H. K. Krishnamurthy, V. Vaidya, S. Weng, K. Ravichandran, P. Kumar, S. Kim, R. Jain, G. Matthew, J. Tschanz, and V. De, "20.1 a digitally controlled fully integrated voltage regulator with on-die solenoid inductor with planar magnetic core in 14nm tri-gate cmos," in *Solid-State Circuits Conference (ISSCC), 2017 IEEE International*, pp. 336–337, IEEE, 2017.

[14] 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.

[15] 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.

[16] J. F. Bulzacchelli, Z. Toprak-Deniz, T. M. Rasmus, J. A. Iadanza, W. L. Bucossi, S. Kim, R. Blanco, C. E. Cox, M. Chhabra, C. D. LeBlanc, C. L. Trudeau, and D. J. Friedman, "Dual-

loop system of distributed microregulators with high dc accuracy, load response time below 500 ps, and 85-mv dropout voltage," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 4, pp. 863–874, 2012.

[17] B. Amelifard and M. Pedram, "Optimal design of the power-delivery network for multiple voltage-island system-on-chips," *IEEE Tran. Comput.-Aided Design Integr. Circuits Syst.*, vol. 28, pp. 888–900, Jun. 2009.

[18] E. Alon and M. Horowitz, "Integrated regulation for energy efficient digital circuits," *IEEE JSSC*, vol. 43, pp. 1795–1807, Aug. 2008.

[19] Z. Zeng, X. Ye, Z. Feng, and P. Li, "Tradeoff analysis and optimization of power delivery networks with on-chip voltage regulation," in *Proc. IEEE/ACM Conf. Cmput.-Aided Design*, pp. 831–836, Jun. 2010.

[20] I. Vaisband, B. Price, S. Köse, Y. Kolla, E. G. Friedman, and J. Fischer, "Distributed ldo regulators in a 28 nm power delivery system," *Analog Integrated Circuits and Signal Processing*, vol. 83, no. 3, pp. 295–309, 2015.

[21] V. Zyuban, J. Friedrich, D. M. Dreps, J. Pille, D. W. Plass, P. J. Restle, Z. T. Deniz, M. M. Ziegler, S. Chu, S. Islam, *et al.*, "Ibm power8 circuit design and energy optimization," *IBM Journal of Research and Development*, vol. 59, no. 1, pp. 9–1, 2015.

[22] S. Lai, B. Yan, and P. Li, "Localized stability checking and design of IC power delivery with distributed voltage regulators," *IEEE Tran. Comput.-Aided Design Integr. Circuits Syst.*, vol. 32, no. 9, pp. 1321–1334, 2013.

[23] F. Lima, A. Geraldes, T. Marques, J. Ramalho, and P. Casimiro, "Embedded cmos distributed voltage regulator for large core loads," in *Solid-State Circuits Conference, 2003. ESSCIRC'03. Proceedings of the 29th European*, pp. 521–524, IEEE, 2003.

[24] H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *Computer Architecture (ISCA), 2011 38th Annual International Symposium on*, pp. 365–376, IEEE, 2011.

[25] M. B. Taylor, "A landscape of the new dark silicon design regime," *IEEE Micro*, vol. 33, no. 5, pp. 8–19, 2013.

[26] Y. Zu, C. R. Lefurgy, J. Leng, M. Halpern, M. S. Floyd, and V. J. Reddi, "Adaptive guardband scheduling to improve system-level efficiency of the power7+," in *Microarchitecture (MICRO), 2015 48th Annual IEEE/ACM International Symposium on*, pp. 308–321, IEEE, 2015.

[27] Z. Hu, A. Buyuktosunoglu, V. Srinivasan, V. Zyuban, H. Jacobson, and P. Bose, "Microarchitectural techniques for power gating of execution units," in *Proceedings of the 2004 international symposium on Low power electronics and design*, pp. 32–37, ACM, 2004.

[28] S. Kim, S. V. Kosonocky, and D. R. Knebel, "Understanding and minimizing ground bounce during mode transition of power gating structures," in *Proceedings of the 2003 International Symposium on Low Power Electronics and Design*, pp. 22–25, ACM, 2003.

[29] C. R. Lefurgy, A. J. Drake, M. S. Floyd, M. S. Allen-Ware, B. Brock, J. A. Tierno, and J. B. Carter, "Active management of timing guardband to save energy in power7," in *Microarchitecture (MICRO), 2011 44th Annual IEEE/ACM International Symposium on*, pp. 1–11, IEEE, 2011.

[30] A. Grenat, S. Pant, R. Rachala, and S. Naffziger, "5.6 adaptive clocking system for improved power efficiency in a 28nm x86-64 microprocessor," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, pp. 106–107, IEEE, 2014.

[31] C. Tokunaga, J. F. Ryan, C. Augustine, J. P. Kulkarni, Y.-C. Shih, S. T. Kim, R. Jain, K. Bowman, A. Raychowdhury, M. M. Khellah, *et al.*, "5.7 a graphics execution core in 22nm cmos featuring adaptive clocking, selective boosting and state-retentive sleep," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, pp. 108–109, IEEE, 2014.

[32] F. C. Lee and Y. Yu, "Input-filter design for switching regulators," *IEEE Transactions on Aerospace and Electronic Systems*, no. 5, pp. 627–634, 1979.

[33] E. G. Ciprut, Albert ; Friedman, "On the stability of distributed on-chip low dropout regulators," in *Circuits and Systems (MWSCAS), 2017 60th International Midwest Symposium on*, IEEE, 2017.

[34] I. Vaisband and E. G. Friedman, "Stability of distributed power delivery systems with multiple parallel on-chip LDO regulators," *IEEE Trans. Power Electronics*, Oct. 2015.

[35] S. Lai, B. Yan, and P. Li, "Stability assurance and design optimization of large power delivery networks with multiple on-chip voltage regulators," in *Proc. IEEE/ACM Conf. Cmput.-Aided Design*, pp. 247–254, 2012.

[36] S. Lai, *Modeling, design and optimization of IC power delivery with on-chip regulation*. PhD thesis, Texas A&M University, 2014.

[37] J. R. Forbes and C. J. Damaren, "Hybrid passivity and finite gain stability theorem: stability and control of systems possessing passivity violations," *IET Control Theory & Applications*, vol. 4, no. 9, pp. 1795–1806, 2010.

[38] X. Zhan, P. Li, and E. Sánchez-Sinencio, "Distributed on-chip regulation: Theoretical stability foundation, over-design reduction and performance optimization," in *Proceedings of the 53rd Annual Design Automation Conference*, p. 54, ACM, 2016.

[39] X. Zhan, P. Li, and E. Sánchez-Sinencio, "Taming the stability-constrained performance optimization challenge of distributed on-chip voltage regulation," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2018.

[40] X. Zhan, J. Riad, P. Li, and E. Sánchez, "Design space exploration of distributed on-chip voltage regulation under stability constraint," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, no. 99, pp. 1–5, 2018.

[41] 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.

[42] 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*, no. 9, pp. 2538–2551, 2017.

[43] H. Li, J. Xu, Z. Wang, P. Yang, R. K. Maeda, and Z. Tian, "Adaptive power delivery system management for many-core processors with on/off-chip voltage regulators," in *Proceedings of the Conference on Design, Automation & Test in Europe*, pp. 1265–1268, European Design and Automation Association, 2017.

[44] W. Godycki, C. Torng, I. Bukreyev, A. Apsel, and C. Batten, "Enabling realistic fine-grain voltage scaling with reconfigurable power distribution networks," in *Microarchitecture (MICRO), 2014 47th Annual IEEE/ACM International Symposium on*, pp. 381–393, IEEE, 2014.

[45] 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.

[46] S. B. Nasir, Y. Lee, and A. Raychowdhury, "Modeling and analysis of system stability in a distributed power delivery network with embedded digital linear regulators," in *Quality Electronic Design (ISQED), 2014 15th International Symposium on*, pp. 68–75, IEEE, 2014.

[47] S. Ben-Yaakov, "Behavioral average modeling and equivalent circuit simulation of switched capacitors converters," *IEEE Transactions on Power Electronics*, vol. 27, no. 2, pp. 632–636, 2012.

[48] G. E. Dullerud and F. Paganini, *A course in robust control theory: a convex approach*, vol. 36. Springer Science & Business Media, 2013.

[49] M. Green and D. J. Limebeer, *Linear robust control*. Courier Corporation, 2012.

[50] A. G. J. MacFarlane and I. Postlethwaite, "The generalized nyquist stability criterion and multivariable root loci," *International Journal of Control*, vol. 25, no. 1, pp. 81–127, 1977.

[51] J. M. Maciejowski, "Multivariable feedback design," *Electronic Systems Engineering Series, Wokingham, England: Addison-Wesley*, vol. 1, 1989.

[52] G. A. Gray and T. G. Kolda, "Algorithm 856: Appspack 4.0: Asynchronous parallel pattern search for derivative-free optimization," *ACM Trans. Math. Softw.*, vol. 32, no. 3, pp. 485–507, 2006.

[53] M. S. Gupta, J. L. Oatley, R. Joseph, G.-Y. Wei, and D. M. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *Design, Automation & Test in Europe Conference & Exhibition, 2007. DATE'07*, pp. 1–6, IEEE, 2007.

[54] T. Xu, P. Li, and S. Sundareswaran, "Decoupling capacitance design strategies for power delivery networks with power gating," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 20, no. 3, p. 38, 2015.

[55] E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill, "Fivr- fully integrated voltage regulators on 4th generation intel® core socs," in *Applied Power Electronics Conference and Exposition (APEC), 2014 Twenty-Ninth Annual IEEE*, pp. 432–439, IEEE, 2014.

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

[57] 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.

[58] M. Ware, K. Rajamani, M. Floyd, B. Brock, J. C. Rubio, F. Rawson, and J. B. Carter, "Architecting for power management: the ibm® power7 approach," in *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*, pp. 1–11, IEEE, 2010.

[59] X. Liu, S. Sun, X. Li, H. Qian, and P. Zhou, "Machine learning for noise sensor placement and full-chip voltage emergency detection," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 36, no. 3, pp. 421–434, 2017.

[60] T. Wang, C. Zhang, J. Xiong, and Y. Shi, "Eagle-eye: a near-optimal statistical framework for noise sensor placement," in *Proceedings of the International Conference on Computer-Aided Design*, pp. 437–443, IEEE Press, 2013.

[61] H. Lin and P. Li, "Relevance vector and feature machine for statistical analog circuit characterization and built-in self-test optimization," in *Proceedings of the 53rd Annual Design Automation Conference*, p. 11, ACM, 2016.

[62] 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.

[63] 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.

[64] 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.

[65] C. Bienia and K. Li, *Benchmarking modern multiprocessors*. Princeton University, 2011.

[66] C.-H. Chan, Y. Zhu, I.-M. Ho, W.-H. Zhang, U. Seng-Pan, and R. P. Martins, "16.4 a 5mw 7b 2.4 gs/s 1-then-2b/cycle sar adc with background offset calibration," in *Solid-State Circuits Conference (ISSCC), 2017 IEEE International*, pp. 282–283, IEEE, 2017.

[67] B. Verbruggen, J. Craninckx, M. Kuijk, P. Wambacq, and G. Van-Der-Plas, "A 2.6 mw 6b 2.2 gs/s 4-times interleaved fully dynamic piplined adc in 40nm digital cmos," in *Digest of Technical Papers of the Solid State Circuits Conference*, pp. 398–400, IEEE, 2010.

[68] H. Akkary, R. Rajwar, and S. T. Srinivasan, "Checkpoint processing and recovery: Towards scalable large instruction window processors," in *Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture*, p. 423, IEEE Computer Society, 2003.

[69] 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.