This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCAD.2017.2778080, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems

Workload-Aware Adaptive Power Delivery System Management for Many-Core Processors

Haoran Li, *Student Member, IEEE,* Jiang Xu, *Member, IEEE,* Zhe Wang, *Student Member, IEEE,* Rafael K.V. Maeda, Peng Yang, *Student Member, IEEE,* Zhongyuan Tian

***Abstract*—The power delivery system (PDS), which plays a crucial role in guaranteeing the proper functionality of computing systems, has been a serious constraint on performance due to its significant power loss, especially for high-performance many-core processors. As the PDS design is usually optimized to provide power to the target chip at its best performance level, the energy efficiency can be notably degraded when the workload of the processors is highly dynamic. Therefore, dynamically adjusting the PDS to adapt to the run-time chip workloads are expected to be beneficial. In addition, the introduction of on-chip voltage regulators (VRs) has also significantly broadened the design space of PDS. In this work, we present a workload-aware quantized power management (QPM) scheme to dynamically manage the PDS in order to improve system energy efficiency. VRs at different stages are scheduled as online or offline based on chip workload estimation and prediction. The simulation results show that with the proposed scheme, a hybrid PDS with on/off-chip VRs can achieve 74.6% overall efficiency on average, 12.7% higher than a conventional PDS with one off-chip VR. The proposed scheme also shows its potential advantage in improving system energy efficiency with large-scale many-core processors and with more advanced processor technology nodes.**

***Index Terms*—power delivery system, power management, voltage regulator, many-core processors, energy efficiency**

I. Introduction

With the design trend towards integrating more cores onto a single die, parallel computing capability of multi-core or even many-core processors has been well exploited to improve system performance. However, as the constraints of power consumption and heat dissipation become increasingly critical in the Dark Silicon Age, the maximum performance is severely throttled by the system energy efficiency[1].

The power delivery system (PDS), which is responsible for supplying sufficient power from the external power source to all the functional units, plays a crucial role in maintaining system functionality and reliability. PDS usually includes DC-DC voltage regulators (VRs), power delivery networks (PDNs), decoupling capacitors (Decap), power management logic units, etc. Voltage regulators, which convert higher or lower input voltages to the required supply voltage, are critical elements of the PDS. There are several types of VRs. Linear regulators can be simply viewed as resistive voltage dividers, providing stable voltage regulation with a simple structure and little area costs. However, their inherent power loss and fixed voltage regulation ratio limits their application. Switched-capacitor regulators only require capacitors as passive components and are able to provide multiple levels of output voltages, but they suffer from poor ability to handle high current load[2]. The most efficient technologies are buck converters, which can easily maintain high efficiency of 80%–90% in a large current load and conversion ratio range[3].

In conventional solutions, VRs are put on a printed circuit board (PCB) with bulky inductors and capacitors. However, integrating VRs on-chip has been increasingly preferred by designers[4][5][6][7][8][9]. Firstly, on-chip VRs with small passive components and MHz-level frequencies are able to achieve fast and fine-grain voltage control. Secondly, the small area overhead makes it easy to realize multi-rail on-chip delivery paths and multiple voltage domains. Besides, on-chip VRs lower the current flowing into the package, potentially alleviating the I^2R power loss and the need for package power pins. Although issues such as lower energy efficiency compared to off-chip VRs are encountered, the introduction of on-chip VRs has greatly broadened the design space of PDS.

Energy efficiency has been an critical issue in PDS design, taking significant proportion of total power losses. The intrinsic power losses in VRs and the delivery loss due to the parasitics of the PDN taking a large proportion of system's total power losses. To make things worse, the "dark silicon" power/thermal constraint limits the maximum performance of modern many-core systems, and power management techniques such as dynamic voltage frequency scaling (DVFS), clock gating or power gating have been extensively applied to improve the energy efficiency, resulting in highly dynamic run-time workloads. As the PDS is usually designed and optimized off-line at the system average performance level, significant PDS energy efficiency degradation will be encountered under light/heavy-load scenarios. Traditionally, VRs at the same stage will supply power evenly to the next stage or target chip, while implementing multiple on-chip and off-chip VRs provides a chance of intelligently managing PDS to further improve system energy efficiency. In this paper, we will explore the potential of applying adaptive power management to dynamically improve the energy efficiency based on the run-time workload.

The rest of the paper is organized as follows. Section II reviews related works. In Section III we present the modeling of different power delivery system paradigms constructed by on/off-chip VRs and PDNs, as well as the design optimization of different power delivery systems. In Section IV, we present

This work is partially supported by GRF16235616 and FSGRF15EG04 of Hong Kong.

The authors are with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong, China. e-mail: {hr.li, jiang.xu}@ust.hk.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCAD.2017.2778080, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems

our quantized power management scheme, and in Section V, we simulate and evaluate the performance of the proposed scheme. Some modeling details of power losses are explained in the Appendix.

II. Related Works

Research has shown that dynamic adjustment of VRs has the potential to improve energy efficiency adaptively. Zumel *et al.* [10] presented a technique for multi-phase buck converters to dynamically change the number of phases in order to reduce the power losses at light load, which is also called phase shedding[11]. In[12], Zhou *et al.* proposed a mode-hopping strategy such that VRs are switched from continuous-conduction mode (CCM) to discontinuous-conduction mode (DCM) at light workload to improve the energy efficiency. In addition, while pulse-width-modulation (PWM) with fixed frequency is the normal control mode for buck converters, some research has focused on pulse-frequency-modulation(PFW) to improve light-load efficiency. Luo *et al.* [13] proposed a skipe cycle modulation(SCM) mode for switching-mode VRs, which regulates the output voltage by different frequencies proportional to the output power, and the light-load efficiency was improved. Abu Qahouq *et al.*[14] presented pulse sliding (PSL), which modulates switching frequencies based on current tracking.

Those circuit-level approaches mainly focused on the internal controls of a single VR, which introduce significant design and control complexity. Some research has focused on system-level PDS optimal design. Wang *et al.* [4] presented an integrated and analytical model for the entire PDS called PowerSoC, providing a platform to explore the design space and evaluate the characteristics of the PDS. They also discovered that hybrid paradigms with both on- and off-chip VRs reach a better balance between energy efficiency and overhead. In [15], Amelifard *et al.* constructed a hybrid PDS for systems with multiple voltage domains and DVFS capability. Combinatorial problems are formulated and solved to select the optimal voltage regulators to achieve specfic design targets. Vaisband *et al.* [16] proposed a heterogeneous methodology employing off-chip VRs and multiple on-chip linear regulators, achieving up to 93% energy efficiency with the optimized PDS.

These above works focus more on static design analysis and optimization, lacking adaptability to the run-time system workloads. How to adaptively and efficiently manage the whole PDS to improve system overall energy efficiency is also worth exploring. Yan *et al.* [17] proposed a hybrid on/off-chip VR scheme called AgileRegulator redeeming Dark Silicon to explore both fine-grain and coarse-grain power phases. The proposed shceme divides processors into clusters and for each cluster used one on-chip VR to power up one core at a time. Lee *et al.* [18] proposed an optimization scheme targeting smartphone platform which involves both VR static switch sizing and dynamic switch modulation to improve light-load VR energy efficiency. Sinkar *et al.* [19] proposed a cost-effective technique to support per-core voltage domains in multi-core processors with LDOs altered from power-gating devices. Toprak-Deniz *et al.* [20] implemented a distributed integrated voltage regulator module(iVRM) system for POWER8 microprocessor. Multiple LDOs, which are controlled by a single VR controller, regulate each power domain, where energy savings is achieved with fine-grained DVFS. Vaisband *et al.* [21] proposed a recursive clustering algorithm that dynamically controls a heterogeneous PDS, which includes off-chip/in-package switching-mode VRs and abundant on-chip LDOs. Most of the previous proposed schemes are applied to multiple on-chip LDOs, which are limited by poor energy efficiency, small load capacity and fixed conversion ratios. In this paper, we construct the PDS with a cascade of both on- and off-chip buck converters, and realize intelligent scheduling based on dynamic chip workload estimation or prediction.

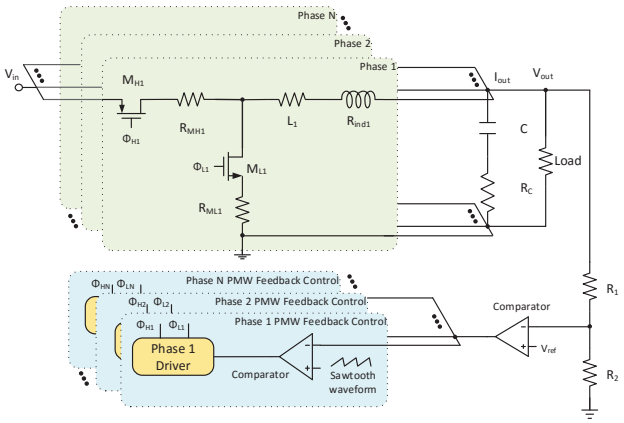

Fig. 1. Schematic diagram of a multi-phase buck converter with PWM feedback control

III. Modeling of Power Delivery Systems

A. Modeling of Buck Voltage Regulators

As key components of a PDS, VRs have a huge impact on the system's overall energy efficiency and performance. Being able to maintain high efficiency over a large workload range and voltage conversion ratios, buck converters are suitable for many-core processors step-down voltage regulation. Fig.1 gives the schematic view of a multi-phase buck converter with PWM feedback control. It is usually composed of MOS power transistors, inductors, output capacitors, feedback control circuits, etc.. Pairs of MOS power transistors work as inverters with fixed frequency and continuously charge and discharge the LC low-pass filter to generate a stable output voltage. N phases are each interleaved with a $360°/N$ phase shift each to reduce the peak current in each inductor and counteract the output voltage ripple. Compensation techniques such as Type-III compensators are employed to maintain loop stability and extend loop gain bandwidth. The voltage conversion ratio is adjusted by changing the duty cycle.

Theoretically, buck converters working in continuous-conduction mode (CCM) can reach 100% conversion efficiency. However, considering the effect of parasitics, their energy efficiency degradation can be significant and should be carefully investigated, especially for on-chip buck converters which utilize integrated power transistors and passive components[5][22][7]. There are five main sources of power losses in a buck converter: the switching loss of the power