

Power Network-on-Chip for Scalable Power Delivery

Inna Vaisband and Eby G. Friedman
Department of Electrical and Computer Engineering
University of Rochester
Rochester, New York 14627
{vaisband, friedman}@ece.rochester.edu

Abstract—Delivering high quality power to the on-chip circuitry with minimum energy loss is an essential component of integrated circuits. The quality of the power supply can be efficiently addressed with multiple power supplies and decoupling capacitors integrated on-chip close to the points-of-load. Distributed power delivery requires the co-design of hundreds of power converters with thousands of decoupling capacitors and billions of current loads within multiple power domains, significantly increasing the design complexity of existing power delivery systems. Efficient real-time management of the power budget in these complicated distributed power delivery systems is impractical with existing *ad hoc* approaches. The concept of a power network on-chip (PNoC) is introduced here as a systematic methodological solution for on-chip power delivery and management that provides enhanced power control and real-time management of resource sharing. A PNoC with four power domains is investigated based on the proposed architecture, and circuits for sensing, routing, and dynamic control of the on-chip power are described. The built-in modularity of the PNoC is exploited to apply dynamic voltage scaling, illustrating the scalability of the PNoC platform, while exhibiting power savings of up to 32%.

I. INTRODUCTION

The delivery of high quality power to the on-chip circuitry with minimum energy loss is an essential component of integrated circuits. To facilitate the integration of diverse functions, architectural, circuit, device, and material level power delivery solutions are required. The quality of the power supply can be efficiently addressed with point-of-load (POL) distributed power delivery [1], [2], which requires the on-chip integration of multiple power supplies. Distributed power delivery requires the co-design of hundreds of power converters with thousands of decoupling capacitors and billions of current loads within multiple power domains, significantly increasing the design complexity of power delivery systems. Per core

dynamic voltage and frequency scaling is a primary concern for efficiently managing a power budget, and requires the on-chip integration of compact controllers within hundreds of power domains and thousands of cores, further increasing the design complexity of these power delivery systems. While in-package and on-chip power integration has recently become a primary concern [3], [4], research remains focused on developing more compact and efficient power supplies. A methodology to design and manage in-package and on-chip power has to date not been a topic of emphasis. Thus, power delivery in modern ICs is currently dominated by *ad hoc* approaches. With the increasing number of power domains, the greater granularity of the on-chip supply voltages, and domain adaptive power requirements, the design of the power delivery process has greatly increased in complexity, and is impractical without a systematic methodology. The primary objective here is to provide a systematic methodology for on-chip power delivery and management.

To cope with design complexity in functional communication, the concept of "separation of different concerns" is used - a fundamental and cost effective approach for increasing performance and design productivity since the complexity of *ad hoc* solutions has become excessively high. Specifically, networks-on-chip (NoC) separate computation from communication, enhancing the performance, scalability, and control of the quality of service (QoS), while supporting heterogeneous integrated systems. Recently, a principle of separation of power conversion and regulation has been introduced [5] that addresses the issue of power efficiency in distributed power supply systems. Consistent with this separation principle, power should be primarily converted with a few power efficient switching supplies, delivered to on-chip voltage clusters, and regulated with linear low dropout (LDO) regulators within the individual power domains, as illustrated in Fig. 1. The concept of a power network on-chip (PNoC) is introduced here based on this separation principle. The analogy between a NoC and PNoC is illustrated in Fig. 2 with simplified NoC and PNoC models. Similar to a network-on-chip, a PNoC decreases the design complexity of power delivery systems, while enhancing the control of the quality of power (QoP) and dynamic voltage scaling (DVS), and providing a scalable platform for efficient power management.

The rest of the paper is organized as follows. The principles of the PNoC design methodology are described in Section II.

This research is supported in part by the Binational Science Foundation under Grant No. 2012139, the National Science Foundation under Grant No. CCF-1329374, and by grants from Qualcomm, Cisco Systems, and Samsung. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
SLIP '14, June 01-02 2014, San Francisco, CA, USA
Copyright 2014 ACM 978-1-4503-3053-4/14/06...\$15.00.
<http://dx.doi.org/10.1145/2633948.2633949>

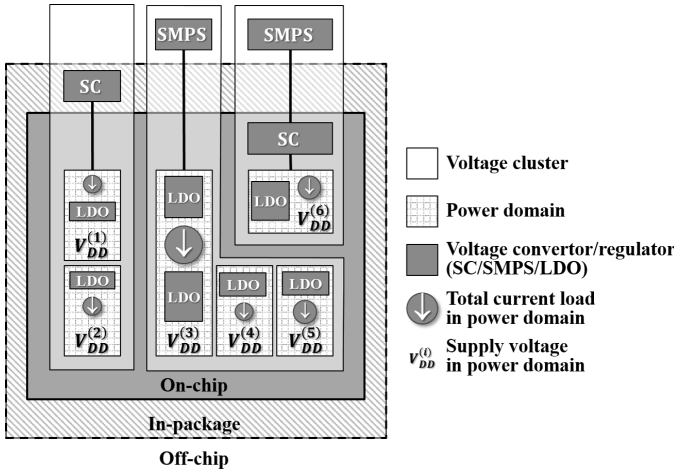


Fig. 1. Heterogeneous power delivery with two switching mode power supplies (SMPS), two switched capacitor (SC) voltage converters, seven low dropout regulators (LDO), and six power domains grouped into three voltage clusters.




Fig. 2. On-chip networks based on the approach of separation of functionality, (a) network-on-chip (NoC), and (b) power network-on-chip (PNoC).

In Section III, the performance of the PNoC architecture is compared with existing approaches based on the evaluation of several test cases. Design and performance issues of the PNoC architecture are also discussed. Some concluding remarks are offered in Section IV.

II. POWER NETWORK-ON-CHIP ARCHITECTURE

The key concept in systemizing the design of power delivery is to convert the power off-chip, in-package, and/or on-chip with multiple power efficient but large switching power supplies, deliver the power to on-chip voltage clusters, and regulate the power with hundreds of linear low dropout (LDO) regulators at the point-of-load [5]. A power network-on-chip is proposed here as a novel systematic solution to on-chip power delivery that leverages distributed point-of-load power delivery within a fine grained power management framework. The PNoC architecture is a mesh of power routers and locally powered loads, as depicted in Fig. 2. The power routers are connected through power switches, distributing current to those local loads with similar voltage requirements. A PNoC is illustrated in Fig. 3 for a single voltage cluster with nine locally powered loads and three different supply voltages, $V_{DD,1}$, $V_{DD,2}$, and $V_{DD,3}$. The power network configuration

is shown at two different times, t_1 and t_2 .

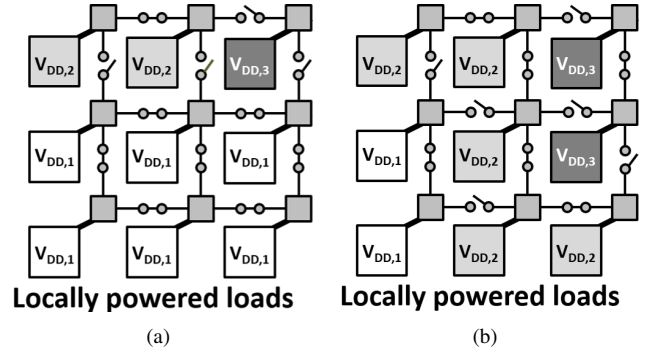


Fig. 3. On-chip power network with multiple locally powered loads and three supply voltage levels (a) PNoC configuration at time t_1 , and (b) PNoC configuration at time t_2 .

The concept of a power network-on-chip is proposed here to virtually manage the power in SoCs through specialized power routers, switches, and programmable control logic, while supporting scalable power delivery in heterogeneous ICs. A PNoC is comprised of physical links and routers that provide both virtual and physical power routing. This system senses the voltages and currents throughout the system, and manages the POL regulators through power switches. Based on the sensed voltages and currents, a programmable unit makes real-time decisions to apply a new set of configurations to the routers per time slot, dynamically managing the on-chip power delivery process. Novel algorithms are required to dynamically customize the power delivery policies through a specialized microcontroller that routes the power. These algorithms satisfy real-time power and performance requirements.

A PNoC composed of power routers connected to global power grids and locally powered loads is illustrated in Fig. 4. Global power from the converters is managed by the power

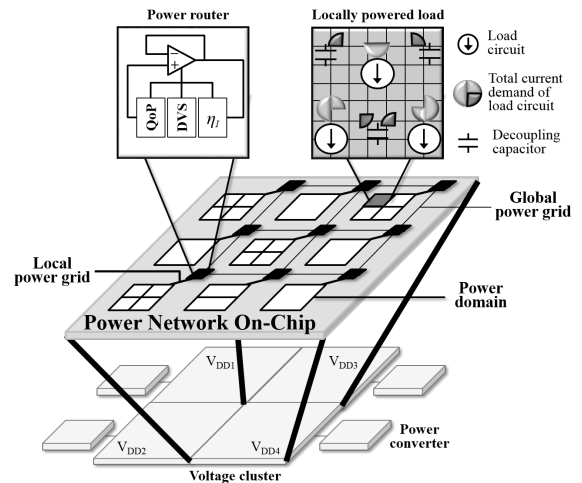


Fig. 4. On-chip power network with routers distributing the current over the power grid to the local loads.

routers, delivered to individual power domains, and regulated

within the locally powered loads. These locally powered loads combine all of the current loads located within a specific on-chip power domain with the decoupling capacitors that supply the local current demand within that region. To support DVS, power switches in the proposed PNoC are dynamically controlled, (dis)connecting power routers within individual voltage clusters in real time. Current loads powered at similar voltage levels therefore draw current from all of the connected power routers, lessening temporary current variations. Similar to a mesh based clock distribution network [6], the shared power supply lessens the effect of the on-chip parasitic impedances, enhancing the voltage regulation and quality of power.

The power routers and local current loads are described in the following subsections. Different PNoC topologies and specific design objectives are also considered.

A. Power routers

The efficient management of the energy budget is dynamically maintained by the power routers. Each power domain is controlled by a single power router. A router topology ranges from a simple linear voltage regulator, shown in Fig. 5(a), to a complex power delivery system, as depicted in Fig. 5(b), with sensors, dynamically adaptable power supplies, switches, and a microcontroller. The proposed structure features real-time voltage/frequency scaling, adaptable energy allocation, and precise control over the on-chip QoP. With the PNoC routers the power is managed locally based on specific local current and voltage demands, decreasing the dependency on remotely located loads and power supplies. Thus, the scalability of power delivery process is enhanced with PNoC approach.

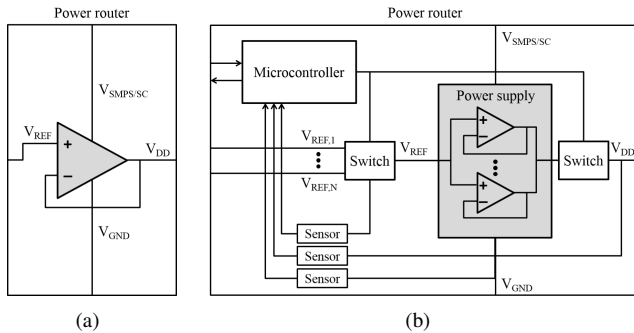


Fig. 5. Power routers for PNoC a) Simple topology with linear voltage regulator. b) Advanced topology with dynamically adaptable voltage regulator and microcontroller.

B. Locally powered loads

Locally powered loads with different current demands and power budgets can be efficiently managed with this proposed PNoC. The local power grids provide a specific voltage to the nearby load circuits. The highly complex interactions among the multiple power supplies, decoupling capacitors, and load circuits need to be considered, where the interactions among the nearby components are typically more significant. The effective region for a point-of-load power supply is the

overlap of the effective regions of the surrounding decoupling capacitors [2], [7]. Loads within the same effective region are combined into a single equivalent locally powered load regulated by a dedicated LDO. All of the LDO regulators within a power domain are controlled by a single power router. A model and closed-form expressions of the interactions among the power supplies, decoupling capacitors, and current loads are required to efficiently partition an IC with billions of loads into power domains and locally powered loads.

III. CASE STUDY

To evaluate the performance of the power router, a PNoC with four power routers is considered, supplying power to four power domains. IBM power grid benchmarks [8] model the behavior of the individual power domains. To simulate a dynamic power supply in PNoC, the original IBM voltage profiles have been scaled to generate the target power supplies between 0.5 volts and 0.8 volts. Target voltage profiles with four voltage levels (0.8 volts, 0.75 volts, 0.7 volts, and 0.65 volts) within a PNoC are illustrated in Fig. 6. The number of




Fig. 6. Preferred and supplied voltage levels in PNoC with four power domains.

power domains with each of the four supply voltages changes dynamically based on the transient power requirements of the power domains.

A schematic of a PNoC with four power domains (I, II, III, and IV) and four power routers (PR_I, PR_{II}, PR_{III}, and PR_{IV}) is presented in Fig. 7. Each of the power routers is composed of an LDO with four switch controlled reference voltages to support dynamic voltage scaling. In addition, the power routers feature adaptive RC compensation and current boost networks controlled by load sensors to provide quality of power control and optimization, as shown in Fig. 8. The adaptive RC compensation network is comprised of a capacitive block connected in series with two resistive blocks, all digitally controlled. These RC impedances are digitally configured to stabilize the LDO within the power routers under a wide range of process variations. The proposed current boost circuit is composed of a sensor block that follows the output voltage




Fig. 7. Proposed PNoC with four power domains and four power routers connected with control switches.




Fig. 8. Power router with voltage regulator, load sensor, and adaptive networks.

at the drain of transistor M_p (see Fig. 8), and a current boost block that controls the current through the differential pair within the LDO. When a high slew rate transition at the output of the LDO occurs, the boost mode is activated, raising the tail current of the LDO differential pair. Alternatively, during regular mode, no additional current flows into the differential pair, enhancing the power efficiency of the LDO.

The power routers are connected with controlled switches to mitigate load transitions in domains with similar supply voltages. To model the RLC parasitic impedances of the package and power network, non-ideal LDO input and output impedances (PN_I , PN_{II} , PN_{III} , and PN_{IV}) are considered, as shown in Fig. 7. The load current characteristics are listed in Table I for each of the four domains. The proposed PNoC is

TABLE I
OUTPUT LOAD IN A PNoC WITH FOUR POWER DOMAINS.

Domain	I	II	III	IV
Minimum output current [mA]	10	75	20	20
Maximum output current [mA]	50	10	30	20
Output transition time [ns]	50	50	10	N/A

simulated in SPICE. Simulation results are presented in Fig. 9, exhibiting a maximum error of 0.35%, 2.0%, and 2.7% for, respectively, the steady state dropout voltage, load regulation due to the output current switching, and load regulation due to dynamic PNoC reconfigurations. Good correlation with the required power supply in Fig. 6 is demonstrated. The power savings in each of the power domains range between 21.0% to 31.6% as compared to a system without dynamic voltage scaling. These average power savings show that the




Fig. 9. Voltage levels in PNoC with four power domains.

PNoC architecture can control the power supplies in real-time, optimizing the power efficiency of the overall power delivery system.

IV. SUMMARY

To address the issues of power delivery complexity and quality of power, a power delivery system should provide a scalable modular architecture that supports integration of additional functional blocks and power features (e.g., DVS, adaptive RC compensation, and efficiency optimization with adaptive current boost) without requiring the re-design of the power delivery system. The architecture should also support heterogeneous circuits and technologies.

The concept and architecture of an on-chip power network is described in this paper. The primary objectives of an on-chip PNoC based power delivery system are 1) to exploit the concept of on-chip networks for systematic power delivery in SoCs to reduce design complexity while increasing scalability, 2) to provide a methodology that separates power conversion and regulation for efficiently enhancing the quality of power, 3) to enable the application of local power routing through a specialized microcontroller for on-chip power management, and 4) to utilize small area power supplies as point-of-load voltage regulators.

REFERENCES

- [1] S. Kose, S. Tam, S. Pinzon, B. McDermott, and E. G. Friedman, "Active Filter Based Hybrid On-Chip DC-DC Converters for Point-of-Load Voltage Regulation," *IEEE Transactions on Very Large Scale Integration (VLSI) Circuits*, Vol. 21, No. 4, pp. 680–691, April 2013.
- [2] S. Kose and E. G. Friedman, "Distributed On-Chip Power Delivery," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, Vol. 2, No. 4, pp. 704–713, December 2012.
- [3] J. Kim *et al.*, "Chip-Package Hierarchical Power Distribution Network Modeling and Analysis Based on a Segmentation Method," *IEEE Transactions on Advanced Packaging*, Vol. 33, No. 3, pp. 647–659, August 2010.
- [4] P. Hazucha *et al.*, "A 233-MHz 80%-87% Efficient Four-Phase DC-DC Converter Utilizing Air-Core Inductors on Package," *IEEE Journal of Solid-State Circuits*, Vol. 40, No. 4, pp. 838–845, April 2005.

- [5] 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, September 2013.
- [6] P. J. Restle *et al.*, "A clock distribution network for microprocessors," *IEEE Journal of Solid-State Circuits*, Vol. 36, No. 5, pp. 792–799, May 2001.
- [7] M. Popovich, M. Sotman, A. Kolodny, and E. G. Friedman, "Effective Radii of On-Chip Decoupling Capacitors," *IEEE Transactions on Very Large Scale Integration (VLSI) Circuits*, Vol. 16, No. 7, pp. 894–907, July 2008.
- [8] S. R. Nassif, "Power Grid Analysis Benchmarks," *Proceedings of the IEEE/ACM Asia and South Pacific Design Automation Conference*, pp. 376–381, January 2008.