

PRADA: Combating Voltage Noise in the NoC Power Supply Through Flow-Control and Routing Algorithms

Prabal Basu Rajesh JayashankaraShridevi Koushik Chakraborty Sanghamitra Roy
USU BRIDGE LAB, Electrical and Computer Engineering, Utah State University
prabalbasu1989@yahoo.com, jsrajesh34@gmail.com, {koushik.chakraborty, sanghamitra.roy}@usu.edu

ABSTRACT

Network-on-Chip (NoC) has become the *de-facto* standard for on-chip communication in MPSoCs. The growing NoC power footprint, increase in the transistor current, and high switching speed of the logic devices, exacerbate the peak power supply noise (PSN) in the NoC power delivery network (PDN). Hence, preserving power supply integrity in the NoC PDN is critical. In this work, we propose PRADA (PSN-aware Runtime Adaptation)—a collection of a novel flow-control protocol (PAF) and an adaptive routing algorithm (PAR), to mitigate PSN in NoCs. Our best scheme achieves 14% and 12% improvements in the regional peak PSN and energy efficiency, with an average of 4.6% performance overhead and marginal area and power footprints.

1. INTRODUCTION

Supply voltage integrity is a growing concern in modern multiprocessor system-on-chips (MPSoCs). The varying current demand due to the simultaneous switching of the logic devices, creates a noise in the power delivery network (PDN), resulting in a drop in the effective supply voltage. This power supply noise (PSN) has a detrimental effect on the performance, reliability and energy efficiency of various system components. As current and upcoming MPSoCs are embracing Network-on-Chips (NoCs) as their de-facto standard for on-chip communication, PSN will negatively impact fault-free communication on them.

In this paper, we uncover a key circuit-architectural insight: *simultaneous and sudden rise in traffic loads within proximal regions in a NoC, can lead to a significant voltage noise*. We also demonstrate that existing NoC flow-control protocols and congestion aware routing algorithms are unable to mitigate the PSN problem effectively. Figure 1 shows the improvement in the regional peak PSN with a representative congestion aware DBAR routing scheme [7], compared to deterministic Dimension Order (DOR) XY routing. Both the routing schemes are used, along with wormhole flow-control. DBAR shows average peak PSN improvements of only 0.1-1%, across all the benchmarks. Some regions show worse peak PSN with DBAR, as DBAR allows simultaneous and large change in activity in proximal routers, causing damaging noise in the NoC PDN.

Our contributions in this paper are as follows:

- We propose a couple of runtime solutions, collectively referred as PRADA (PSN-aware Runtime Adaptation), to mitigate PSN in NoCs. PRADA comprises a novel PSN-Aware Flow-control (PAF) and an adaptive PSN-Aware Routing (PAR) algorithm (Section 2).
- Our best scheme can reduce the regional peak PSN by




Figure 1: Improvement in the regional peak PSN with DBAR compared to DOR. Green regions represent PSN improvement, while red regions represent PSN degradation.

14% and improve the energy efficiency by 12.2% compared to a representative routing scheme (DBAR), with a nominal 4.6% average performance overhead and marginal area/power overheads (Section 4).

To the best of our knowledge, our work is the first of its kind to investigate voltage noise aware flow-control and routing algorithms for a NoC.

2. PSN AWARE RUNTIME ADAPTATIONS

PRADA aims to dampen high simultaneous current loads in proximal regions, by dynamically altering their respective *flit acceptance potentials* and proactively dispersing the flit routes in the network.

2.1 Design Challenges

(a) *Performance impact*: Run-time adaptations to mitigate PSN should have a low performance overhead.

(b) *Deadlock avoidance*: Throttling the *flit acceptance potential* of a router can create buffer back-pressure in the upstream routers. Under a high flit injection rate, the back-pressure can grow so large that it may lead to a network deadlock. It is important to guarantee freedom from deadlocks in PRADA.

(c) *Scalability*: An adaptive PSN improvement technique should scale with the size of the communication fabric. It is imperative to minimize its implementation overhead so as to sustain its efficacy in future exascale computing.

2.2 Design of PAF

The design of PAF involves a hierarchical approach to dictate the *Maximum Current Load (MCL)*¹ across the NoC, while ensuring a minimal performance impact.

2.2.1 Hierarchical MCL Allocation

High concurrent switching of proximal regions is avoided by carefully adjusting the MCL allocated to each region. To realize MCL allocation principles at different granularities, we define a metric *Flit Acceptance Potential (FLAP)*. For a given

¹We define MCL of an integrated circuit in an epoch (few cycles) as the highest possible amount of current that the circuit can draw from the power supply, in that epoch.

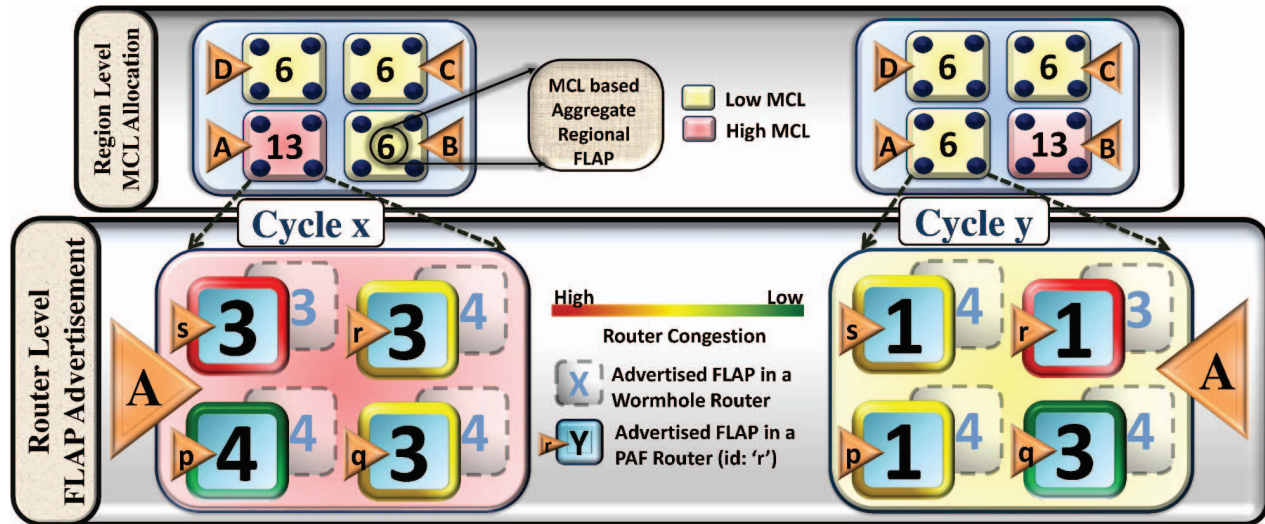


Figure 2: Overview of the PAF flow-control protocol.

input channel of a router, the *FLAP* is set to 1 when it can receive an incoming flit (otherwise it is set to 0). For a router, the *FLAP* indicates the aggregate *FLAP* of its input channels. Similarly, the *FLAP* of a particular region represents the aggregate *FLAP* of the routers in that region.

At any given time, the *FLAP* of a router employing wormhole flow control in a 2D mesh with four input channels is 4, when all of its input channels can receive at least one flit. PAF allocates variable MCL to each region by dynamically throttling their *FLAP*s, irrespective of the space availability in the input channel's buffers.

MCL allocation is a hierarchical process that can be applied at multiple spatial granularities. The allocated MCL for the large region is distributed among the sub-regions, ensuring that proximal sub-regions are not simultaneously allocated with high MCLs. At the lowest granularity, each router's *FLAP* is managed in a manner that is consistent with the MCL allocation of the entire sub-region.

2.2.2 Illustrative Example

Figure 2 depicts the PAF technique using a 4x4 2D-mesh NoC, divided into 4 regions (A,B,C,D), each comprising 4 routers. In cycle *x*, PAF allocates a high MCL to region A and low MCLs to the proximal regions (B,C,D). To ensure a fair provisioning, PAF redistributes the MCL allocation in cycle *y*, so that region B is allocated with a high MCL, while its proximal regions are allocated with low MCLs.

The allocated MCL translates to a regional *FLAP*, which is distributed among the routers of a region. For example, in cycle *x*, a regional *FLAP* of 13 is distributed among the routers of region A. Router *p* advertises a *FLAP* of 4, while the other routers (*q,r* and *s*) advertise 3 *FLAP*s each.

2.2.3 Optimizations of PAF

The generic PAF technique needs multiple optimizations to efficiently tackle the design challenges (Section 2.1).

Minimizing Performance Impact: We explore a few complementary approaches to retain a high performance in PAF.

- **Judicious *FLAP* Management:** To avoid a large flit delay in a given region, PAF allows intermittent high and low *FLAP*s in a router. For example, in contrast to cycle *x*,

router *q* advertises more *FLAP* (3) in cycle *y* compared to the other routers.

- **Topological Awareness:** PAF can be adapted based on the network topology and expected traffic pattern. For example, we can allocate greater *FLAP*s to the central routers of a mesh, to meet their high resource demand.
- **Congestion Awareness:** We explore two variants of PAF. **PAF-Static:** This is a congestion agnostic variant that statically allocates high and low *FLAP*s to the regional routers based on a round-robin fairness policy. **PAF-Cong:** This variant manages the *FLAP* of a router based on the relative congestion of the network buffers. The least congested router of a region is allocated with a high *FLAP*. However, the other routers are allocated with low *FLAP*s to avoid high simultaneous switching. The aggregate *FLAP*s of the routers is consistent with the allocated MCL of the region. For example, in cycle *y* in Figure 2, the least congested router *q* advertises more *FLAP* (3) compared to the other routers, each of which, advertises 1 *FLAP*. The aggregate *FLAP*s of the routers match the allocated MCL based regional *FLAP*.

Avoiding Deadlock: Repeated blocking of the flits at the same input channel of a router in successive cycles can cause a deadlock situation. To ensure freedom from deadlock, PAF adopts a round-robin fairness scheme to restrict flit reception across all the input channels of a router. Moreover, PAF uses deterministically routed escape VCs, allowing all the possible turns in the network without a deadlock situation.

Scalability: PAF is a hierarchical technique that uses local network information at the smallest regional granularity to ascertain the *FLAP*s of the routers. As the size of the smallest region remains the same even for a larger NoC, PAF can scale efficiently with the network size.

2.3 PAF Aware Adaptive Routing Algorithm

Dynamically throttling the *FLAP* of a router may cause an intermittent upsurge in the local PSN due to an increased resource contention. We propose PAR (*PSN-Aware Routing*), a PAF cognizant routing algorithm, to circumvent this upsurge, by steering the flit towards an unthrottled downstream path. PAR, primarily makes the routing decision based on

the relative regional congestion information, aggregated solely along the minimal paths. If the chosen output channel has a throttled FLAP, PAR reroutes the flit to an orthogonal output channel, strictly maintaining the minimal path constraint. This strategy reduces local current spike and PSN by relieving router contention, but may occasionally increase the network latency by routing some flits towards more congested downstream paths. In a scenario, where both the minimal paths are blocked due to throttled FLAPs, the flit adheres to the initial channel assignment and waits in the upstream router for another cycle. PAR incurs no additional circuit overhead as it utilizes the same information required for PAF.

2.4 Implementation

The implementation of PRADA involves FLAP management and congestion management in the regional routers.

- **FLAP Management:** Reception of flits in a router is managed by sending a *credit_valid* signal to the upstream router. We use the *credit_valid* signal, along with a statically managed, low overhead, round-robin logic, to ascertain the FLAP of a router. Additionally, we feed the *credit_valid* signal with one of the output bits of a simple one-hot encoded ring counter, to sporadically restrict an incoming flit.
- **Congestion Management:** We create a low-bandwidth monitoring network to propagate the congestion information among the adjacent routers in a region. The monitoring network involves an *aggregation* and a *propagation* module at the router's low overhead port preselection logic [5]. The *aggregation* module combines the weighted congestion values from the downstream routers and the *propagation* module transmits the congestion information to the adjacent routers of a region.

3. METHODOLOGY

Our methodology can be classified into PSN estimation and performance evaluation.

3.1 Power Supply Noise Estimation

Dahir et al. recently proposed a MATLAB based PSN estimation tool for NoC [3]. We re-implement the tool in C++ and integrate it with Booksim2.0 [6], to tightly couple the stages of architectural evaluation and PSN estimation. We collect the following data for accurate PSN estimation.

- **Interconnect RLC Parameters:** We compute the R,L,C values of the grid interconnect for the 32 nm technology node using the ASU PTM interconnect model [11].
- **Router Pipeline Energies:** We use the recently proposed DSENT 0.91 [10] tool to evaluate the energy of the router pipeline stages, using the router microarchitectural parameters for the 32 nm technology node.
- **Traffic and Router Activity Dump:** We instrument Booksim 2.0, in order to dump various router activities at each cycle, by running PARSEC benchmarks [1] on an 8x8 regular 2D mesh NoC. To mimic the traffic generated by multiple co-scheduled applications in an MP-SoC, we superimpose heavy random traffic (with a flit injection rate of 0.15) on top of the original application induced traffic of the PARSEC benchmarks.

3.2 Performance Evaluation

Following are the various performance evaluation metrics studied in this work.

3.2.1 Regional Peak PSN

We divide an 8x8 mesh NoC into 16 regions, each containing 4 routers, and assign minimum operating voltage at the regional granularity, to ensure fault-free communication. We evaluate the regional peak PSN of the comparative schemes.

3.2.2 Average Network Latency

We use Booksim2.0 as our architectural simulator to run network simulations (for 1 million cycles) of the comparative schemes using real workloads. We report the performance overhead of the comparative schemes in terms of overall average network latency.

3.2.3 Energy Delay Product

Mitigating the peak supply noise reduces the minimum voltage guardband required for fault-free operation. As a result, all the routers in the network can operate at a reduced supply voltage and consume less energy. We analyze the improvement in router energy using DSENT, and estimate the energy efficiency using Energy Delay Product (EDP).

3.2.4 Area and Power

We modify the RTL of the open source Stanford Verilog model of a modern virtual channel NoC router [2] to implement the PRADA techniques. The router is assumed to be a part of a 2D mesh topology with 5-input/output ports and 8 VCs per port. We synthesize the augmented router RTL with the TSMC 45nm library using Synopsys Design Compiler and calculate the area and power overheads.

4. EXPERIMENTAL RESULTS

In this section, we analyze the efficacy and overheads of various comparative schemes (Section 4.1).

4.1 Comparative Schemes

Table 1 shows the various schemes we explore in this study.

Schemes	Flow-Control	Routing Algorithm
<i>Baseline</i>	Wormhole	DBAR
<i>PAF-SD</i>	PAF-Static	DBAR
<i>PAF-SP</i>	PAF-Static	PAR
<i>PAF-CD</i>	PAF-Cong	DBAR
<i>PAF-CP</i>	PAF-Cong	PAR

Table 1: Comparative schemes.

4.2 Regional Peak PSN Comparison

Figure 3 shows the percentage improvement in regional peak PSN of various comparative schemes, with respect to the baseline. We notice that PAF-SP and PAF-CP, show more pronounced improvements, as PAR can mitigate local PSN by reducing the intermittent upsurge in resource contention. The respective maximum regional PSN improvements observed in all the schemes are 8.1%, 13.2%, 7.8% and 14%, with respective average PSN improvements as 4.7%, 5.7%, 5% and 5.8%. Some regions show slightly worse peak PSN compared to the baseline, due to occasional increase in local congestion, incurred by PAF.

4.3 Performance Overhead

Figure 4 shows the network latency overheads of the comparative schemes, with respect to the baseline. PAF-SP and




Figure 3: Percentage improvement in regional peak PSN with respect to the baseline for comparative schemes. Each small square represents a region consisting of 4 routers. Reddish and greenish regions represent worse and improved peak PSN, respectively.




Figure 4: Performance overhead (lower is better).




Figure 5: EDP improvement (higher is better).

PAF-CP incur slightly more overheads, compared to the other schemes, as PAR sometimes takes more congested downstream paths in the network. We also notice that PAF-SD performs slightly better than PAF-CD due to PAF-Static’s inherent fairness in FLAP allocation. There is a maximum performance degradation of 5.7% (Ferret in PAF-CP) with an average degradation of 4.6%, across all the schemes.

4.4 Energy Efficiency Comparison

Figure 5 shows the improvement in energy efficiency of the comparative schemes, in terms of EDP. We notice that both PAF-Static and PAF-Cong incur better EDP, when used along with PAR routing (PAF-SP and PAF-CP). We observe a maximum EDP improvement of 12.2% (Swaptions in PAF-SP), with an average improvement of about 10%, across all the schemes. PAF-SP shows maximum improvements in EDP, among all the schemes.

4.5 Area and Power Footprint

PAF-Static incurs marginal area and power overheads of 0.10% and 0.16%, respectively. Due to the larger footprints of the congestion management unit, PAF-Cong incurs more area (1.42%) and power (2.38%) overheads.

5. RELATED WORK

Works related to our effort of reducing peak voltage noise can be categorized in the following two domains.

Understanding voltage noise in a NoC: Penolaazi et al. pro-

pose a power model for the *Nostrum* NoC to accurately estimate power fluctuations for a NoC load [9]. Recently, Dahir et al. have developed a dedicated tool for NoC PSN analysis based on their detailed workload model [3]. However, their work does not aim at dampening peak PSN in NoCs.

Flow control and Routing techniques: Traditionally, flow-control techniques have been developed to improve the communication efficiency and fault tolerance in NoCs. In [8], Michelogiannakis et al. propose *elastic buffers* to improve the peak throughput and average latency in a NoC. Jafri et al. propose an adaptive flow control, which can dynamically adapt to varying loads to maximize performance and minimize the energy consumption. Further, abundant congestion aware schemes have been developed to improve communication efficiency under high loads [4, 7]. *But, to the best of our knowledge, no previous work explores the use of flow-control and routing algorithm for peak noise mitigation in NoCs.*

6. CONCLUSION

In this work, we propose PRADA, a collection of a novel flow-control protocol (PAF) and an adaptive routing algorithm (PAR), to improve the peak PSN in NoCs. Our best scheme improves the regional peak PSN by 14% and the EDP by ~12% with marginal overheads.

7. ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation grants (CAREER-1253024, CCF-1318826, CNS-1421022, and CNS-1421068). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

8. REFERENCES

- [1] PARSEC. <http://parsec.cs.princeton.edu/>.
- [2] BECKER, D. Open Source NoC Router RTL, August 2012.
- [3] DAHIR, N. AND OTHERS Modeling and Tools for Power Supply Variations Analysis in Networks-on-Chip. *TC* 63, 3 (2014), 679–690.
- [4] EBRAHIMI, M. AND OTHERS CATRA- congestion aware trapezoid-based routing algorithm for on-chip networks. *Proc. of DATE* (2012), 320–325.
- [5] GRATZ, P. AND OTHERS Regional congestion awareness for load balance in networks-on-chip. In *HPCA* (2008), pp. 203–214.
- [6] JIANG, N. AND OTHERS A detailed and flexible cycle-accurate Network-on-Chip simulator. In *ISPASS* (2013), pp. 86–96.
- [7] MA, S. AND OTHERS DBAR: an efficient routing algorithm to support multiple concurrent applications in networks-on-chip. In *Proc. of ISCA* (2011), pp. 413–424.
- [8] MICHELOGIANNAKIS, G. AND OTHERS Elastic-buffer flow control for on-chip networks. In *HPCA* (2009), pp. 151–162.
- [9] PENOLAZZI, S., AND JANTSCH, A. A High Level Power Model for the *Nostrum* NoC. In *Proc. of DSD* (2006), pp. 673–676.
- [10] SUN, C. AND OTHERS DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling. In *NOCS* (2012), pp. 201–210.
- [11] ZHAO, W., AND CAO, Y. New Generation of Predictive Technology Model for Sub-45 nm Early Design Exploration. *Electron Devices, IEEE Transactions on* (2006), 2816–2823.