

Hybrid Network on Chip (HNoC): Local Buses with a Global Mesh Architecture

Payman Zarkesh-Ha
University of New Mexico
Department of ECE
Albuquerque, NM, USA
+1 (505) 277-6724

George B. P. Bezerra
University of New Mexico
Department of CS
Albuquerque, NM, USA
+1 (505) 277-3411

Stephanie Forrest
University of New Mexico
Department of CS
Albuquerque, NM, USA
+1 (505) 277-7104

Melanie Moses
University of New Mexico
Department of CS
Albuquerque, NM, USA
+1 (505) 277-9140

payman@ece.unm.edu gbezerra@cs.unm.edu forrest@cs.unm.edu melaniem@unm.edu

ABSTRACT

Network on chip (NoC) is often implemented with packet-based communication rather than bus connections between cores. Although NoC is a good solution for long-distance communication, local buses are more efficient for short-distance connections. In this paper, we propose a hybrid network on chip (HNoC) fabric that uses *local buses* for nearest-neighbor communication and the *standard NoC topology* for global interconnection. Local buses carry all the nearest-neighbor traffic, reducing traffic on the global network, which results in increased throughput and reduced energy consumption.

Based on a communication probability density (CPD) function derived from Rent's rule, it is shown that in a 25-core chip multiprocessor, HNoC can remove up to 78% of the traffic from the global NoC topology, which results in 4.6x higher throughput and a 58% reduction in energy consumption compared to a conventional NoC topology.

Categories and Subject Descriptors

B.7.2 [Integrated Circuits]: Design aids – *simulation*; C.4 [Performance of Systems]: Modeling techniques.

General Terms

Design, Performance, Theory.

Keywords

Stochastic model, hybrid network on chip, throughput, network traffic, local buses, global mesh.

1. INTRODUCTION

Microprocessor performance has increased exponentially over the last four decades as advancing semiconductor technology has vastly increased the quantity and improved the speed of on-chip transistors available to circuit designers. Traditionally, computer

designers took advantage of these resources to improve uniprocessor structures because of its simpler programming model compared to systems with distributed structures [1]. However, power consumption and wire delay have recently limited the continued scaling of uniprocessor systems making chip multiprocessor architectures more appealing [2]. In addition, network on chip (NoC) has become the emerging paradigm for communication within large chip multiprocessor systems to overcome scalability, power, delay, and other issues with global interconnects.

Many different NoC topologies for multiprocessor systems have been proposed and studied by researchers. For instance, in [3], Balfour and Dally published a comprehensive analysis of several possible NoC topologies, such as: mesh, torus, fat tree, concentrated mesh, concentrated torus, and tapered fat tree. Of these networks, the best topology in terms of energy and communication time was concentrated mesh, a type of mesh topology that uses larger-radix routers to cluster four processors at each mesh node and contains express channels around the perimeter of the network. Figure 1 illustrates three examples of NoC topologies.

To further optimize NoC performance and energy, a hybrid optical/electrical NoC architecture was recently proposed [4-6], where a photonic network and an electronic network coordinate to provide the system with high bandwidth communications. The optical circuit switching network handles long-lived bulk data transfers, whereas the secondary lower-bandwidth electronic packet switching network accommodates collectives and short-lived data exchanges. Although in this approach a high bandwidth optical network and a low cost electrical network are hybridized together to improve energy and performance, the communication is still purely packet based.

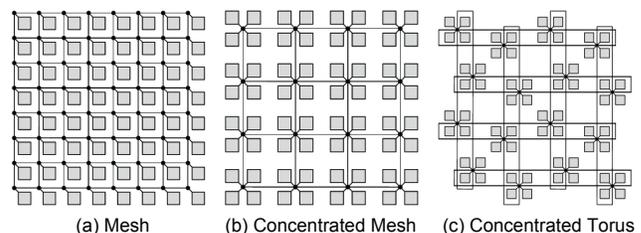


Figure 1. Three examples of NoC topologies [3].

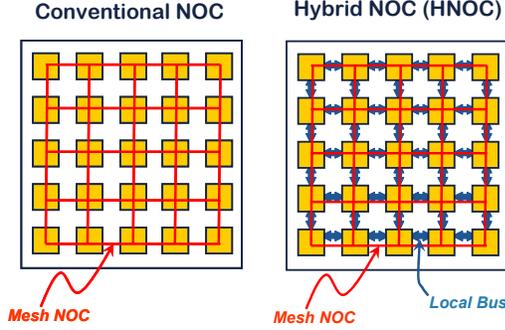


Figure 2. HNoC versus conventional NoC

In this paper, we propose a hybrid network-on-chip (HNoC) that uses standard NoC topology (e.g., mesh) for packet-based global interconnections along with local buses for nearest-neighbor communications as shown in Figure 2. Unlike the hybrid optical/electrical NoC architecture that is purely packet-based, our HNoC uses local buses to transmit data directly to the nearest neighbors in a parallel fashion, which eliminates the need for serializer, router, and deserializer. Moreover, since the local bus interconnects are short, they inherently exhibit lower loss and therefore can provide higher bandwidth and consume less power.

Section 2 gives details of the proposed HNoC architecture. In Section 3, we use a communication probability density (CPD) model derived from Rent's rule to assess traffic patterns for conventional NoC. In Section 4, the Rent's rule-based CPD is used to predict the performance and energy consumption of HNoC versus conventional NoC. Some simulation results using a commercial NoC simulator are then presented in Section 5. Finally, some discussions and conclusions are provided in Section 6 and 7, respectively.

2. HYBRID NoC (HNoC) TOPOLOGY

Today's chip multiprocessor interconnects use packet-switching networks to connect the cores. Inter-processor messages are broken into packets, which are routed through the NoC switches. However, as the number of cores increases, performance and power consumption of NoC degrade significantly due to higher communication costs [6]. To address this issue, we first review the communication cost in a packet-switching network. To transfer data from one core to another through NoC fabric, first it is packetized, then sent to the transmitting router, passed through the network wiring channel, delivered to the receiving router, and finally depacketized. This leads to an inefficient communication for short distance data transmission.

Figure 2 illustrates the proposed HNoC topology, where nearest-neighbor communications are carried by local bus interconnects instead of the mesh NoC. Local bus interconnects are direct connections between neighboring cores dedicated for direct data exchange without any packetizing overhead. The energy consumption and latency of short distance communication through local buses are therefore much smaller than those through NoC fabric. Note that the HNoC topology is different from Concentrated Mesh shown in Figure 1, where there is no local bus for neighboring cores.

To assess the benefits of HNoC versus conventional NoC fabric, we next review the application of Rent's rule to digital systems. A multiprocessor system can be seen as a digital circuit, where a logic gate is analogous to a processor core. By empirical observation and from the Rent's rule based wire-length distribution, we find that the majority of interconnections are short. This is true because circuit designers tend to connect blocks that are closer together. Similarly, programmers try to map tasks with high communication rates onto processors that are closer together in a multiprocessor system. It can therefore be expected that the majority of communications within a multiprocessor NoC system are to close neighbors. Using Rent's rule, we incorporate a communication probability distribution (CPD) model to quantify the advantage of HNoC versus conventional NoC topologies.

3. COMMUNICATION PROBABILITY DISTRIBUTION (CPD) MODEL

A statistical analysis of several NoC traffic patterns was presented by Soteriou et. al. [7] in 2006. Then Greenfield et al. [8] applied the principle of Rent's rule to the analysis of NoC architecture and showed that for a VLSI design consisting of many blocks wired together, if one replaced these wires with a NoC, then the Rent's terminal-exponent of the former may match the bandwidth-exponent of the latter [8], or:

$$B = bN^p, \quad (1)$$

where B is the communication bandwidth, N is the number of nodes, and b and p are the Rent's coefficient and exponent, respectively.

It was also shown in [8] that Rent's rule can be used to characterize NoC architectures similar to the way it characterizes interconnects in VLSI designs. For instance, NoC hop-length distributions can be directly derived from VLSI wire-length distributions [8]. Later, Heirman et al. [9] validated the application of Rent's rule to NoC architectures by analyzing the SPLASH-2 benchmarks. They further analyzed temporal behavior of network traffic using Rent's rule and confirmed that Rent's rule and all of its applications to VLSI can be applied to NoC topologies.

Recently, Bezerra et al. [10] presented a closed form model, *communication probability Distribution* (CPD), as another derivative of Rent's rule for multiprocessor systems. CPD is the probability that a processor communicates with another processor at distance d in a chip multiprocessor system. The communication probability of distance d for an $N \times N$ multiprocessor system is determined by [10] as:

$$CPD(d) = \frac{\Gamma f(d)}{d} \left\{ \begin{array}{l} [1 + d(d-1)]^p - [d(d-1)]^p \\ + [d(d+1)]^p - [1 + d(d+1)]^p \end{array} \right\}, \quad (2)$$

where p is the Rent's exponent, Γ is the normalization coefficient such that $\sum_{d=1}^{2N-2} CPD(d) = 1$, and $f(d)$ is given by:

$$f(d) = \begin{cases} \frac{d^3}{3} - 2d^2N + \frac{d}{3}(6N^2 - 1), & 1 \leq d < N \\ -\frac{d^3}{3} + 2d^2N - \frac{d}{3}(12N^2 - 1) + \frac{2}{3}N(4N^2 - 1), & N \leq d \leq 2N - 2 \end{cases} \quad (3)$$

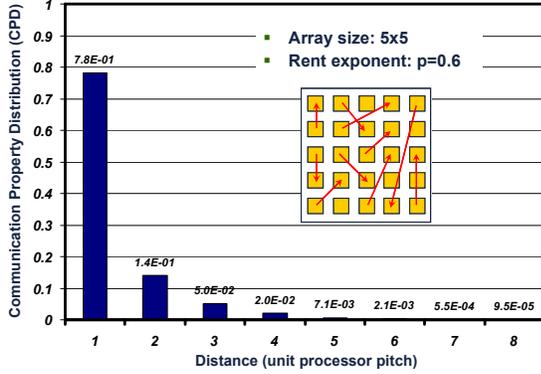


Figure 3. Example of communication probability distribution

An example of the CPD for 5×5 arrays of multiprocessors is shown in Fig. 3, for the Rent’s exponent of 0.6. In this example, 78% of the communication is between nearest neighbors.

The communication probability given in (2) can be considered simply as the normalized hop-distance distribution proposed in [8] or the normalized wire-length distribution proposed by [11]. However, unlike the hop-distance distribution, CPD can provide substantial information about the NoC system with only the Rent’s exponent, p , and independent of the bandwidth Rent’s coefficient, b , because it is cancelled out by normalization.

Figure 4 illustrates the CPD distributions using various Rent’s exponents ranging from 0.1 to 0.9. As shown, Rent’s rule predicts that the majority of the communications in a 5×5 array of multiprocessors are within the nearest neighbors.

4. HNoC VERSUS CONVENTIONAL NoC

In Section 3 we showed that the nearest-neighbor traffic can be directly transmitted through local buses rather than the main NoC fabric. In this Section we use the Rent’s rule based CPD to quantify the advantages of HNoC versus conventional NoC.

4.1 Throughput Analysis

Consider the 5×5 array of multiprocessors shown in Fig. 3. If 78% of the communication can be moved to local buses, the NoC will be responsible for only 22% of the traffic. Therefore, the

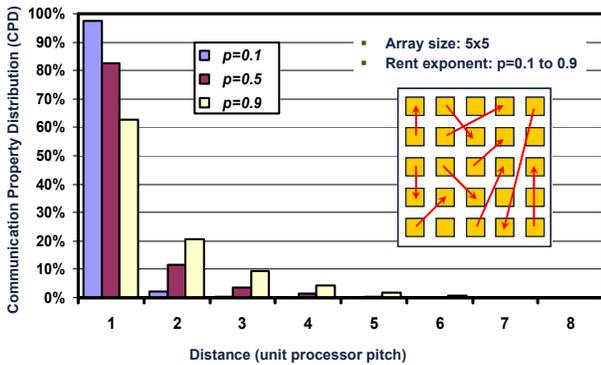


Figure 4. CPD for various Rent’s exponents, p

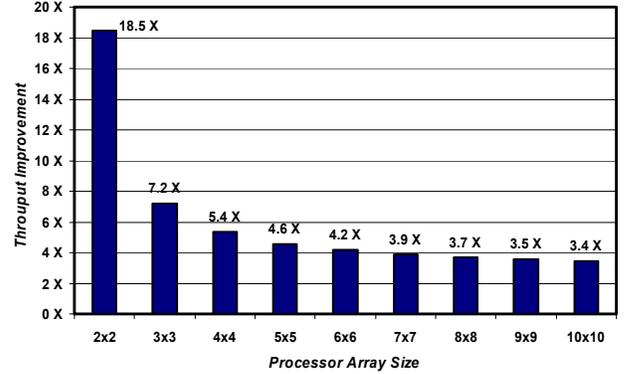


Figure 5. Predicted maximum throughput improvement in HNoC versus conventional NoC for different array sizes, assuming that $p=0.6$

throughput of the NoC can potentially improve by 4.6x for the maximum injection rate. In general, the rate of throughput improvement is determined by:

$$\frac{HNoC_{Throughput}}{NoC_{Throughput}} \approx \frac{1}{1 - CPD(1)}, \quad (4)$$

where CPD is given by (2) and (3).

The improvement rate for various array sizes is shown in Figure 5 assuming $p=0.6$. As shown in this figure, even in a large array size of 10×10 (i.e. 100 processors), the predicted improvement rate is 3.4x.

Note that the improvement rates given in (4) and shown in Figure 5 are for an ideal case, where it is assumed that the local buses impose no overhead. Equation (4) presents an *ultimate* HNoC benefit without making too many assumptions, which may be application or design dependent. In practice, however, depending on the design and application, the local bus overhead will impact the performance of HNoC as will be presented in Section 5.

4.2 Energy Consumption Analysis

Similarly, energy consumption can be reduced by introducing local buses. Again, consider the 5×5 array of multiprocessor shown in Figure 3. The CPD shown in Figure 3 can be used to compute the energy reduction rate in HNoC. Using the NoC energy model presented in [10] and assuming that the power consumption of routers is dominant, the energy consumption of the HNoC against conventional NoC for the same throughput can potentially be reduced by factor of 2.4x. In general, for the same throughput the energy reduction rate in an array of $N \times N$ microprocessors is approximated by [10]:

$$\frac{NoC_{Energy}}{HNoC_{Energy}} \approx \frac{\sum_{d=1}^{2N-2} d \cdot CPD(d)}{\sum_{d=2}^{2N-2} d \cdot CPD(d)}, \quad (5)$$

where the CPD is given by (2) and (3).

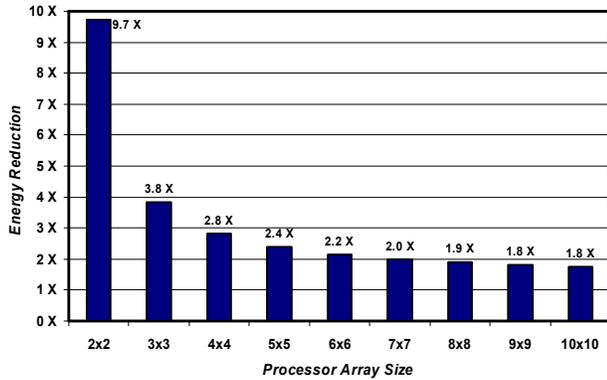


Figure 6. Energy reduction in HNoC compared to conventional NoC for different array sizes when $p=0.6$

Using (5) and assuming that $p=0.6$, the projection of the best case energy reduction in HNoC versus conventional NoC is shown in Figure 6. As shown in this figure, the energy reduction rate is about 1.8x for large array size of 10×10 .

Similar to the throughput improvement analysis, the power improvement model given in (5) is for an ideal case, where it is assumed that the local buses impose no overhead. In practice, however, depending on the design and application, the local bus power overhead will impact the power consumption of HNoC as will be presented in Section 5.

4.3 Scalability Analysis

It is expected that the number of cores will increase with technology scaling. HNoC is scalable and provides better power and performance even in a chip with many cores. Although the advantage of HNoC over conventional NoC shrinks as the number of cores grows, the throughput and energy improvements are considerable even for large multiprocessor systems. Based on Figure 5, the rate of throughput improvement of HNoC over NoC stays at about 3.5x, and based on Figure 6, the rate of energy saving in HNoC over NoC stays at about 1.8x for up to 100 processors.

5. SIMULATION RESULTS

To verify our analysis in Section 4, the proposed HNoC was implemented in a system simulator, Orion 2 [12] and compared to the conventional NoC. The system parameters for this simulation are shown in Table 1.

As shown in Figure 7, the benefit of HNoC is negligible when the injection rate is low. Once the injection rate increases beyond 0.2 packets per cycle, traffic starts to saturate the throughput of conventional mesh NoC. However, HNoC can handle most of the additional traffic using the local buses and continue to provide more throughput. Figure 7 shows that at the injection rate of 0.6 packets per cycle, HNoC provides 2.6x more throughput than conventional mesh NoC.

Without the local bus overhead, equation (4) predicts a 3.7x improvement in throughput. However, due to the limitation of local bus bandwidth and the injection rate constraint in this test system, the HNoC throughput improvement is reduced to 2.7x.

Table 1. System Parameters used in simulations

System Parameters	Values
Number of Cores	64
Die Size	1cm x 1cm
Technology Node	45 nm
Clock Frequency	1 GHz
Flit Size	64 bits
Packet Size	5 flits
Ren's Exponent, p	0.60

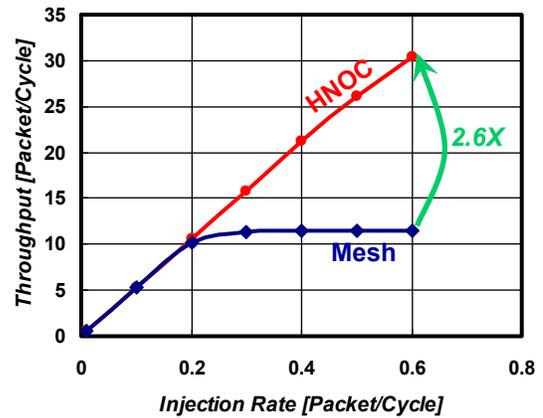


Figure 7. Simulation results for HNoC throughput

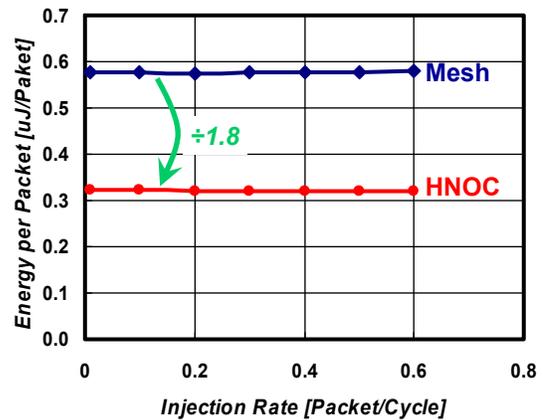


Figure 8. Simulation results for HNoC energy per packet

Figure 8 illustrates the energy consumption per packet in HNoC and conventional mesh NoC. As shown in this figure, HNoC dissipates 1.8x less energy per packet than conventional mesh NoC within a large range of injection rates.

Similarly, without the local bus overhead, equation (5) predicts a 1.9x improvement in energy, which is close to the experimental results. This means that in this test system the local bus energy overhead is negligible compared to the mesh NoC energy consumption.

6. DISCUSSION

Rent's rule arises in digital systems because EDA tools optimize placement and routing in order to reduce wiring requirements and minimize the number of long wires. Similar to EDA tools, compilers will soon be designed such that locality becomes a primary objective and it is evident that a "smart compiler" with optimized program mapping and task assignment will be required to achieve the best benefit of NoC architecture.

The purpose of the analysis presented in this paper is to support the concept of traffic localization as previously suggested by some researchers [13-15]. Once traffic localization is obtained, the proposed HNoC architecture can significantly improve the energy usage and performance of the system by directing the local communications through the low-latency, high-bandwidth, and low-power local buses and leaving the global communications to the standard NoC topology.

In practice, however, achieving this locality may be challenging. Even when the algorithm exhibits localized communication the system needs to be able to map it such that the neighboring threads (from a communication point of view) are mapped onto neighboring network cores. Moreover, even though some applications communicate in a localized fashion at each point in time, sometimes a thread's neighbors can change over time, which may require runtime re-mapping and significant data movement [9]. Therefore, in such a short period of time, when there is only a large burst of long-distance communications, the local buses in HNoC may not be able to provide significant support. However, on average and over time, HNoC will indirectly support long-distance communication by removing the local communication traffic from the mesh NoC, leaving the mesh NoC fully dedicated to long-distance traffic.

7. CONCLUSIONS

A hybrid network on chip (HNoC) fabric that uses local buses for nearest-neighbor communication and the standard NoC topology for global interconnection was described. It is shown that the local buses can carry all the nearest-neighbor traffic, reducing traffic on the global network, which results in increased throughput and reduced energy consumption.

Based on a CPD function derived from Rent's rule, it was shown that HNoC can significantly improve the throughput and reduce the NoC energy consumption. To achieve the benefit of HNoC, compilers must take locality as a primary objective, similar to EDA tools used in VLSI designs.

8. ACKNOWLEDGMENTS

P. Zarkesh-Ha acknowledges the support of the US Department of Energy, Office of Science, under Grant DE-SC0002113. S. Forrest acknowledges the support of the National Science Foundation (grants CCF 0621900, CCR-0331580, SHF-0905236), Air Force Office of Scientific Research MURI grant FA9550-07-1-0532, and the Santa Fe Institute. M. Moses acknowledges the support of a grant from Microsoft Research.

The authors would like to thank fruitful discussions with Bill Loh at Verdura Systems, and Jim Koford, Venkatesh Akella, Mathew Wojko, and Douglas Boyle at Novarus Logic. Finally,

the in depth review of the SLIP technical committee members and their constructive comments are greatly appreciated.

9. REFERENCES

- [1] Theodoros Konstantakopoulos, Jonathan Eastep, James Psota, and Anant Agarwal, "Energy Scalability of On-Chip Interconnection Networks in Multicore Architectures," MIT CSAIL Technical Report, November 2007.
- [2] M. Horowitz and W. Dally, "How scaling will change processor architecture," Proceedings of the *International Solid-State Circuits Conference (ISSCC)*, pp. 132-133, Feb. 2004.
- [3] James Balfour and William J. Dally, "Design Tradeoffs for Tiled CMP On-Chip Networks," Proceedings of the 20th *ACM International Conference on Supercomputing (ICS)*, June 2006.
- [4] P. C. Luca, P. Partha, X. Yuan, "Networks-on-Chip in Emerging Interconnect Paradigms: Advantages and Challenges," *ACM/IEEE International Symposium on Networks-on-Chip*, pp. 93-102, June 2009.
- [5] S. Kamil, A. Pinar, D. Gunter, M. Lijewski, L. Oliker, J. Shalf, "Reconfigurable Hybrid Interconnection for Static and Dynamic Scientific Applications", *ACM International Conference on Computing Frontiers*, 2007.
- [6] G. Hendry, S. Kamil, A. Biberman, J. Chan, B. Lee, M. Mohiyuddin, A. Jain, K. Bergman, L.P. Carloni, J. Kubiatowicz, L. Oliker, and J. Shalf, "Analysis of Photonic Networks for a Chip Multi-Processor Using Scientific Applications," *Proceedings of the Third International Symposium on Networks-on-Chip (NoCS)*, pp. 104-113, June 2009.
- [7] V. Soteriou, H. Wang, and L.S. Peh, "A Statistical Traffic Model for On-Chip Interconnection Networks," *IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, pp. 104-116, September 2006.
- [8] D. Greenfield, A. Banerjee, J.-G. Lee, and S. Moore, "Implications of rent's rule for NoC design and its fault-tolerance," *International Symposium on Networks-on-Chip*, pp. 283 - 294, June 2007.
- [9] W. Heirman, J. Dambre, D. Stroobandt, and J. Campenhout, "Rent's rule and parallel programs: Characterizing network traffic behavior," *International Workshop on System Level Interconnect Prediction*, pp. 87-94, April 2008.
- [10] G. Bezerra, S. Forrest, M. Moses, A. Davis, and P. Zarkesh-Ha, "Prediction of NoC Energy Consumption using Rent's rule based Communication Probability Distribution," submitted to *International Workshop on System Level Interconnect Prediction*, June 2010.
- [11] J. A. Davis, V. K. De and J. D. Meindl, "A Stochastic Wire-length Distribution for Gigascale Integration (GSI): Part I: Derivation and Validation," *IEEE Transaction on Electron Devices*, pp. 580.589, March 1998.
- [12] A. Kahng, B. Li, L. Peh, and K. Samadi. Orion 2.0: A fast and accurate noc power and area model for early-stage

design space exploration. In *Design, Automation, and Test in Europe*, pp. 423-428, June 2009.

[13] P. Pande, C. Grecu, M. Jones, A. Ivanov, and R. Saleh: "Effect of Traffic Localization on Energy Dissipation in NoC-based Interconnect," *IEEE International Symposium on Circuits and Systems*, pp 1774-1777, July 2005.

[14] J. Hu and R. Marculescu, "Energy-aware Mapping for Tile-based NoC Architectures under Performance Constraints,"

Asia and South Pacific Design Automation Conference, pp. 233-239, Jan. 2003.

[15] E. Nilsson, M. Millberg, J. Oberg, and A. Jantsch, "Load Distribution with the Proximity Congestion Awareness in a Network on Chip," *IEEE Design Automation and Test in Europe Conference and Exhibition*, pp. 1126-1127, Dec. 2003.