

Modeling the power-reliability tradeoff in on-chip Networks

Claas Cornelius^{*}, Frank Sill⁺, Dirk Timmermann^{*}

^{*} Institute of Applied Microelectronics and Computer Engineering, University of Rostock, Richard-Wagner-Str. 31, 18119 Rostock-Warnemünde, Germany

⁺ Department of Electrical Engineering, Federal University of Minas Gerais (UFMG), Av. Antônio Carlos 6627, 31270-010 Belo Horizonte, Brazil

1. Motivation

Continuous scaling of transistor dimensions has enabled the exponential increase in performance and functionality of integrated circuits since its introduction. However, power and reliability issues have been rising similarly and endanger the future performance gain (Kim et al. 2003). To model all influencing parameters in different architectures and to explore the design space of miscellaneous complex systems requires tremendous computational effort because billions of transistors are expected to be integrated on a single die by the end of this decade.

Conventional bus-based or point-to-point connected communication structures are used to implement such complex systems. Though, they exhibit various drawbacks with the increasing number of modules and on-chip communication traffic, whereas scalability and communication bandwidth are the most limiting characteristics. Thus, this work targets Networks-on-Chip (NoC) which seem promising to become the dominating system paradigm to designing complex integrated systems (Benini et al. 2002). An NoC consists of independent computational resources that can be of any kind – e.g. general purpose processors, memories, I/Os or intellectual properties (IP). Those different resources are interconnected by an on-chip network that is constructed in a layered approach based on the OSI reference model (see Fig. 1 a). Such a system approach allows concurrent and separated computation as well as communication and already includes redundancy implicitly. Even though, all redundant components and additional features of NoC increase the reliability, they also rise the power consumption so that these two aspects have to be traded off against each other.

The first and intuitive approach to evaluate the described trade-off is to implement the system using hardware description languages (HDL) which enable fairly accurate results. High flexibility can additionally be achieved by

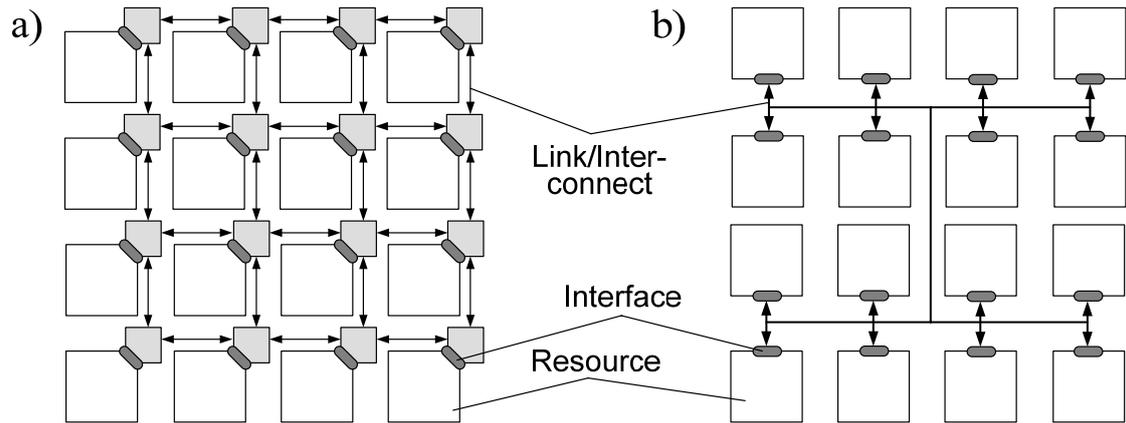


Fig. 1: Simple example of 16 resources in a regular 4x4 mesh interconnected with a) a network-on-chip and b) a single shared bus

using parametrizable components (Zeferino et al. 2004). However, simulation speed is a limiting factor when using hardware description languages such as VHDL or Verilog. Therefore, another approach was presented by Wiklund (Wiklund et al. 2004) using an event based simulator written in C++ whereas the system behavior is further abstracted by using input and traffic flow models represented in XML. This reduces the computational effort significantly and various results focusing on the communication parameters were presented. Nonetheless, power dissipation or reliability of the NoC was not considered. A further approach is to emulate the system on a field programmable gate array (FPGA) which enables maximum test speed (Genko et al. 2005). However, monitoring of local parameters or injecting errors to access reliability is rather difficult and system size is limited to current technology.

In the following, first results are discussed describing an on-chip network and its behavior in terms of power consumption, system throughput and reliability as a function of the number of IPs, i.e. resources. Thereto, the chosen approach and system scenario is described in the following section before the results are given and compared to a simple bus-based system design.

2. Approach

The selected technology scenario is mainly based on data given by the International Technology Roadmap for Semiconductors (SIA 2005) for a technology node expected in 2010. With this data as a foundation, analytical models were derived in a first step for a mesh-based system interconnected with an NoC and a simple bus-based system as the reference (see Fig. 1 a and b). These models describe system characteristics such as area, latency and frequency as well as power consumption with its different components:

- P_{dyn} : Dynamic components such as charging capacitances or glitching
- P_{leak} : Leakage currents during active state and stand-by
- P_{short} : Short circuit currents between the power rails

Additionally, network related parameters such as maximum and average data rate were modeled. Though, further interesting parameters can not be described analytically as they depend on the dynamic behavior of the system. This refers to accepted traffic or areas of increased packet congestion over time as well as temperature distribution or system reaction due to various types of failures and parameter variations.

Therefore, a simulator was developed in a second step that should ideally combine the advantages of gate level and network simulations, i.e. little computational effort and acceptable accuracy to evaluate and to explore the design space of different complex systems. Hence, simulation speed is the primary key demand. Thereto, a Transaction Level Model (TLM) was selected which allows modeling the system behavior functionally without any notion of time. Then, additional cycle accuracy was achieved by introducing delay values for the transactions that were extracted from extensive transistor and gate level implementations of the various components introduced in section 1. Concluding, such an approach clearly outperforms gate level simulations in terms of simulation speed and enables more realistic results than the analytical models. Admittedly, hardware emulations are still significantly faster, but they are limited to current technology and do not allow design space exploration for future NoCs where reliability and power consumption need to be considered thoroughly. Additionally, simulations enable extraction and monitoring of system parameters independently over time for the different components. Furthermore, system behavior in the presence of temporary failures such as soft errors and permanent malfunction of modules can be observed.

As the developed simulator requires the extraction of parameters from transistor and gate level simulations it is not possible to quickly evaluate other technology nodes. However, this can be achieved by simply adapting the data basis for the analytical models so that both approaches combined together allow valuable hints across design space and technology.

In the next section, results based on the analytical models are given that also relate to first results from the developed simulator. Though, accuracy and minor deviations of results between the analytical models and the simulations are still to be evaluated.

3. Results

The given results relate to a mesh-based core of computation resources interconnected by an on-chip network as shown in Fig. 1 a) and a simple bus reference model as shown in Fig. 1 b). The bus is the simplest approach of a shared media and could further be enhanced by a hierarchical or bridged approach. Nonetheless, it is only given as a reference and the results of these enhanced bus versions lie in between the simple bus and this NoC. Furthermore, it is assumed that the investigated systems require equal effort for the clocking and synchronization of computation resources and communication modules.

Lastly, the results were obtained for constant chip size which translates into varying size of the resources (i.e. IPs) depending on their number.

In Fig. 2 the power consumption of the NoC is normalized to the bus reference and is depicted over the number of IPs. It can be observed that the power dissipation is more than 2.5 times higher than the one of the bus (see squares) which relates pretty closely to the gate count of the communication structures and is dominated by the registers. However, in contrast to the bus where the complete bus length has to be charged/discharged even when communicating with its neighbor, communication in the NoC only requires those components being actively part of the transmission. Thus, power depends on the distance which can be exploited by an appropriate local application mapping where the workload is distributed locally first. The results for such a system where the workloads and thus also communication paths are not distributed uniformly but mainly locally show the decrease of power consumption by more than 20 % (see triangles). These values strongly depend on the degree of exploited locality.

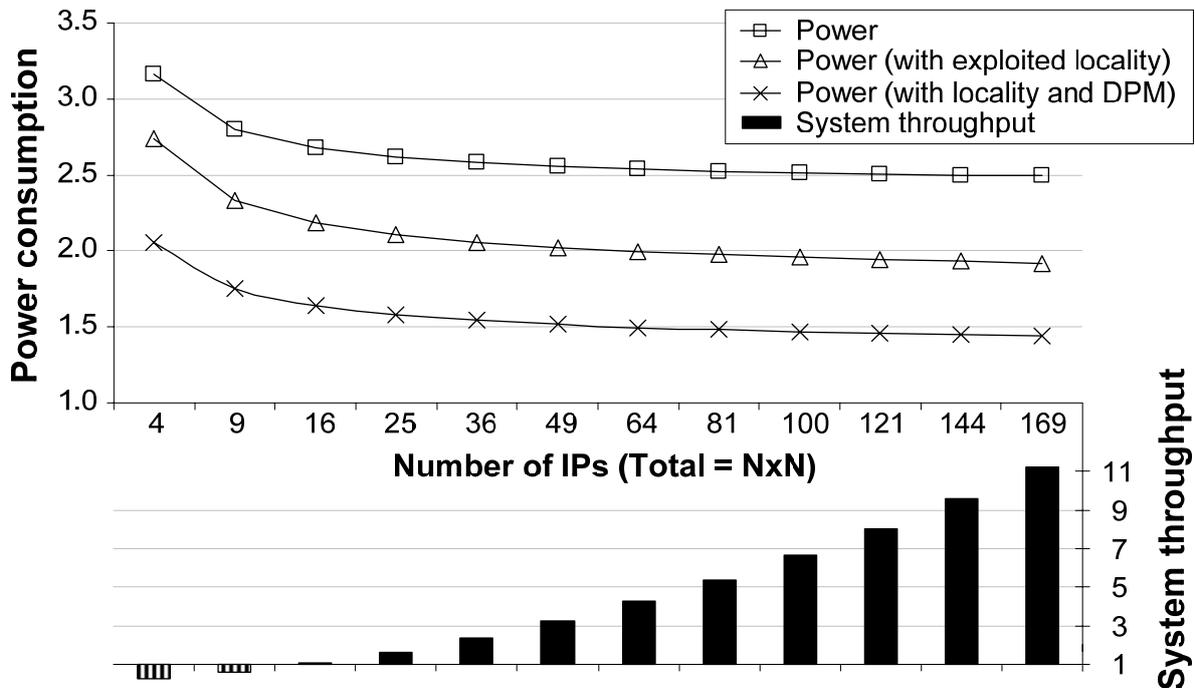


Fig. 2: Normalized power consumption and system throughput of the mesh-based network-on-chip in comparison to the reference bus modeled in a 45 nm technology

A further enhancement can be achieved by dynamic power management due to the same reasons as for the locality. That is, not the complete communication structure has to be used during a transmission. For instance when part of the communication structure is not used it can simply be turned off by clock- or power-gating techniques (e. g. sleep transistors). Finer grained approaches can be applied similarly (e. g. Dynamic Frequency/Voltage Scaling). Thus, with such measures the power overhead of the NoC can further be reduced to roughly 50 % compared to the bus (see crosses). The values for the exploited locality

and the dynamic power management do not consider the control logic itself which is assumed to be negligible with respect to the size of the communication structure especially in NoCs with a large amount of resources.

However, these power values are based on the same clock frequency for both system approaches which does not correspond to the system throughput. In Fig. 2 the system throughput of an NoC is also depicted and normalized to the bus. It can be observed that the system throughput of the NoC increases with the number of resources because an independent router is added – that allows concurrent communication – with each additional resource. Hence, the system throughput scales with the number of resources in contrast to the bus. With this in mind, the NoC is even more power efficient than the bus when the NoC is scaled to the same system throughput.

An exception has to be made for less than 10 resources where the system throughput of the bus outperforms the NoC. The reason is that the frequencies for such system size are very similar but the NoC suffers from a lower saturation point in the communication structure and little overhead due to the header of the packets.

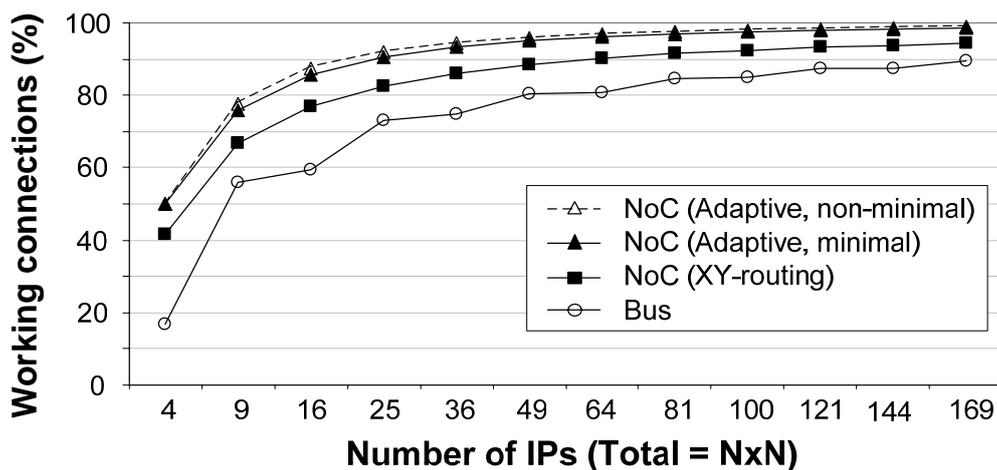


Fig. 3: Reliability in terms of the average number of working connections in a cohesive connected system for a single error (depicted for various NoC routing schemes and the reference bus)

Further results of the comparison are depicted in Fig. 3 where reliability is shown in terms of working connections for different routing schemes. The routing scheme describes the algorithm that is used to establish a path between source and sink in the network. XY-routing is a very simple scheme where data packets are routed along the X-dimension first and along the Y-dimension when the column of the sink is reached. Thus, this scheme is deterministic but cannot adapt the path when a permanent error occurs. However, the percentage of working connections in the system is much higher than that of the bus (see squares) due to the independent and redundant communication modules in the NoC. Even better results can be achieved with adaptive routing schemes. Such a scheme (here greedy) adapts the path for the data transmission based on

additional information from the network. Thus, packet congestions, thermal hot spots or unreliable paths can be avoided. The results in Fig. 3 (see triangles) exhibit enhanced reliability which comes at the price that packets might have to be reordered at the sink. A third approach that is depicted in Fig. 3 is also an adaptive approach but the communication path is not necessarily the shortest and minimal one. This slightly increases the number of working connections in the presence of a single error but comes at the price of increased power consumption due to a higher hop count.

It should be noted that the shown results are the average number of working connections. The worst case scenario for the NoC results in only little less working connections due to the redundancy of its communication structure. In contrast thereto, the worst case results of the bus are significantly lower. This can be understood by considering that the complete bus size can possibly be cut in half dividing the system into two independent parts.

4. Conclusion

A combined approach of analytical models and a simulator was presented that allows fairly accurate and fast evaluations of different systems across design space and technology. First results were also presented and prove the chosen approach to be an option for simulations requiring relatively little computational effort. Nonetheless, final error estimation and verification of the simulator still needs to be done. However, figures of performance, power, system throughput and reliability were extracted for a mesh-based system interconnected by a network-on-chip and a bus. The comparison of the results exhibited that on-chip networks promise to be both a power-efficient and reliable system approach.

5. References

- Kim N. S. et al. (2003) Leakage current: Moore's law meets static power. *IEEE Computer*, Vol. 36, No. 12, pp. 668-75.
- Benini L. and de Micheli G. (2002) Networks on chips: a new SoC paradigm. *IEEE Computer*, Vol. 35, No. 1, pp. 70-78.
- Zeferino C., Kreutz M. and Susin A. (2004) RASoC: A router soft-core for networks-on-chip. In *Proc. of DATE*, Vol. 3, pp. 198-203.
- Wiklund D., Sathe S. and Liu D. (2004) Network on chip simulations for benchmarking. In *Proc. of IWSOC*, pp. 269-274.
- Genko N. et al. (2005) A complete network-on-chip emulation framework. In *Proc. of DATE*, Vol. 1, pp. 246-251.
- SIA: Semiconductor Industry Association (2005) International Technology Roadmap for Semiconductors. Published online: <http://www.itrs.net/>.