

Topology-Unaware Routing in Irregular Self-Assembled Networks-on-Chip: An Explorative Case Study

Christof Teuscher
Los Alamos National Laboratory
CCS-3, MS-B256
Los Alamos, NM, 87545, USA
christof@teuscher.ch

ABSTRACT

Networks-on-chips (NoCs) offer promising solutions to many of today's on-chip interconnect and communication challenges. While traditional NoC are designed in a top-down way, interconnecting functional blocks with either a regular or a fully customized topology, bottom-up self-assembled conductive nanowires have the potential to create highly complex interconnect fabrics in a very simple way at virtually no cost. However, such complex networks are usually inherently irregular, heterogeneous, and unreliable, and the principal challenge thus consists in developing appropriate paradigms that allow for efficient and reliable communication. Here, we address this challenge on an architectural level within a simple, hybrid, and abstract framework of functional IP blocks that are irregularly interconnected by a nanowire fabric. We have previously shown that certain irregular 3D assemblies and interconnects have major advantages over regular 2D and 3D mesh fabrics in terms of communication performance and the robustness against failures. We extend this body of work and compare the communication properties of simple, local, and topology-unaware routing strategies on several network-on-chip interconnect topologies. We further analyze the scalability and the robustness of these approaches. The results underline the importance of decentralized routing in such assemblies, but also highlight the limits. Our contributions are relevant for the wiring and communication solutions of future bottom-up self-assembled Avogadro-scale systems, and are part of the general effort to explore alternative design paradigms and architectural trade-offs between block granularity and interconnect complexity.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Computer-Communication Networks [distributed networks, network topology, packet-switching networks]

General Terms

Design

Keywords

network-on-chip, nanowire, routing, interconnect topology

1. INTRODUCTION AND MOTIVATION

For many years, people have investigated alternative computing substrates and architectures with the goal to go beyond Moore's Law, however, there is a lack of consensus on what type of technology and computing architecture holds most promises to keep up the current pace of progress. In this paper, we will primarily focus on self-assembled nanoscale electronics based on nanowires or nanotubes because these fabrication technologies have become quite mature on the physical level. It is, however, still unclear how to develop higher-level computational architectures in a reliable way, although a number of promising approaches have been explored in detail in the past.

Building a scalable computing architecture on top of a potentially very unreliable physical substrate is a non-trivial task, which is guided by a number of major trade-offs in the design space [14], such as the number and the characteristics of the resources available, the required performance, the energy consumption, and the reliability. The lack of systematic understanding of these issues and of clear design methodologies makes the process still more of an art than of a scientific endeavor and the appearance of novel and non-standard physical computing devices (e.g., [17]) generally only aggravates these difficulties.

In recent years, the importance of interconnects on electronic chips has outrun the importance of transistors as a dominant factor of performance [6, 8, 11]. The reasons are twofold: (1) the transistor switching speed for traditional silicon is much faster than the average wire delays and (2) the required chip area for interconnects has dramatically increased. The 2003 ITRS roadmap [1] listed a number of critical challenges for interconnects and states that “[i]t is now widely conceded that technology alone cannot solve the on-chip global interconnect problem with current design methodologies.” The major challenges are related to delays of non-scalable global interconnects and reliability in general, which leads to the observation that simple scaling will no longer satisfy performance requirements as feature sizes continue to shrink [8].

The goal of this paper is to compare several simple and local routing strategies with regard to their ability to scale up

to larger system sizes. We hypothesize that it will be very hard to self-assemble complicated computing elements and therefore seek for the most simplest solution that would still allow to obtain a reasonable performance. In order to compare the results, we have chosen three representative assemblies and interconnect architectures: (1) a regular 2D mesh (i.e., cellular-automata-like) arrangement and interconnect (Fig 1, A), (2) a regular 3D mesh (i.e., cellular-automata-like) arrangement and interconnect (Fig 1, B), and (3) an irregularly arranged and interconnected but realistic assembly of components (Fig 1, C). More details will be given in Section 2. Routing is crucial in our framework because it essentially is the only free parameter once the components are arranged and the connections established.

The motivation for investigating irregularly arranged and interconnected assemblies of simple components, that could potentially be built easily and cheaply by self-assembly, can be summarized by the following observations: (1) it is unclear whether a precisely regular and homogeneous arrangement of components is needed and possible on a multi-billion-component or even an Avogadro-scale assembly of nanoscale components [17] and (2) “[s]elf-assembly makes it relatively easy to form a random array of wires with randomly attached switches” [20].

By using an abstract, yet physically plausible and fabrication-friendly nanoscale computing framework, we have shown elsewhere [15, 16] that self-assembled interconnect fabrics with small-world (or small-world-like) [19] properties have major advantages in terms of performance and robustness over purely regular and nearest-neighbor connected fabrics, such as mesh topologies. While there is ample evidence of the superior communication characteristics of small-world and power-law over locally connected topologies, most abstract models are not physically plausible and are thus of limited significance for real-world implementations. For example, it is very unrealistic to assume a uniform rewiring probability (as in the original Watts-Strogatz model [19]) over all possible nodes. Spatial aspects of small-world topologies and wiring-cost perspectives have only recently gained more attention [9, 10, 13].

Here we are interested in experimentally exploring topology-unaware routing schemes that could be used for bottom-up self-assembled networks-on-chip architectures [7] (e.g., based on nanowires), and which therefore possess a potentially very irregular structure. The framework will be briefly described in Section 2, while Section 3 describes for simple experiments.




Figure 1: Regular and locally interconnected 2D (A) and 3D (B) ensembles of processing nodes (PNs), switch nodes (SNs), and interconnections. (C) A 3D random ensemble (RE).

2. DESCRIPTION OF THE FRAMEWORK

The framework has been described in detail elsewhere [16] and we will therefore only briefly summarize the main characteristics of the network-on-chip-like approach. Also, the physical plausibility and realizability of such an assembly—which we hypothesize could be bottom-up self-assembled from conductive nanowires—has been discussed in [16], however, several questions remain open and are part of ongoing research.

2.1 Node and Link Model

The basic system-on-chip-like architecture (see also Figure 1) that we use is composed of (1) P programmable computing elements, called *processing nodes* (PNs), and (2) of an associated switch-based interconnect fabric, which is itself composed of (3) S *switch nodes* (SNs) and (4) bi-directional point-to-point interconnects among them. Both processing and switch nodes can be considered as simple modules of a large-scale system that need to communicate efficiently among each other. Each switch node can only transmit in parallel messages on C different virtual channels to its neighbors according to a specific routing scheme. No further information processing is done in the switch nodes. We assume that they can temporarily store a limited number of M messages. The processing nodes, on the other hand, simply send and receive messages according to a specific traffic scheme. Since we are mainly interested in interconnect and routing issues in this paper, we do not further specify or limit the processing nodes’ computing capacity.

2.2 Network Topologies

Most real networks, such as brain networks, electronic circuits, the Internet, and social networks share the so-called *small-world* (SW) property [19]. Compared to purely locally interconnected networks (such as the mesh interconnect), small-world networks have a very short average distance between any pair of nodes, which makes them particularly interesting for efficient communication.

The classical Watts-Strogatz small-world network [19] is built from a regular lattice with only nearest neighbor connections. Every link is then rewired with a *rewiring probability* p to a randomly chosen node. Thus, by varying p , one can obtain a fully regular ($p = 0$) and a fully random ($p = 1$) network topology. The rewiring procedure establishes “shortcuts” in the network, which significantly lower the average distance (i.e., the number of edges to traverse) between any pair of nodes. In the original model, the length distribution of the shortcuts is uniform since a node is chosen randomly. If the rewiring of the connections is done proportional to a power law, $l^{-\alpha}$, where l is the wire length, then we obtain a *small-world power-law network*. The exponent α affects the network’s communication characteristics [10] and navigability [9], which is better than in the uniformly generated SW network.

In a real network, it is fair to assume that local connections have a lower cost (in terms of resources required and delay) than long-distance connections. Physically realizing small-world networks with uniformly distributed long-distance connections is thus not realistic and distance, i.e., the wiring cost, needs to be taken into account [13].

In this paper, we will compare the following three reference network topologies in order to quantitatively compare and evaluate the routing protocols:

- *2DCA*: 2D (unfolded) regularly arranged and locally interconnected mesh (von Neumann neighborhood), see Figure 1 (A);
- *3DCA*: 3D (unfolded) regularly arranged and locally interconnected mesh, see Figure 1 (B);
- *3DRERealistic*: 3D *random ensemble* (RE), small-world, power-law, $\alpha = 1.8$, upper limit $k_{max} = 10$ on the number of connections per node, independently of the average connectivity $SN_k = 4$, see Figure 1 (C).

For the *random ensemble* (RE)—formerly also called *random multitude* (RM)—as depicted in Figure 1 (C), both processing and switch nodes are randomly arranged in 3D space. Both the arrangement and the wire topology are inspired by self-assembled nanoscale electronics [15, 16]. To make comparisons with the CA grid-like architectures easier, we assume that each processing node is connected to the nearest switch node by a single connection of 0.01 unit length only. The switch nodes are connected among themselves by a small-world power-law network [13] with average connectivity $SN_k = 4$, $\alpha = 1.8$, and $k_{max} = 10$. The choice of these values was guided by the experiments. The idea of this restriction is to make the topology more physically plausible because it is not plausible in our framework to assume that certain nodes can establish a larger number of connections. Oshida and Ihara [12] have investigated random networks-on-chip with the scale-free property, which are based on hubs that are highly connected.

2.3 Routing and Traffic Models

There exists a large number of different routing strategies and flavors, which allow to send packets along pre-specified or dynamically chosen paths in a given network from a source to a destination. In this paper, we deal with packet routing only and the switch nodes do therefore have to know where to route a packet that they receive. The path information can be obtained dynamically or statically, based on the available information on the network’s topology. Efficient search and information transfer on complex networks while avoiding congestion and optimizing throughput and latency at the same time are of great importance in real-world systems [2, 4, 5]. It has also been shown that small-world and scale-free networks offer great communication characteristics, are efficient to navigate [9, 18], and reduce congestion. Routing based on local information only (e.g., [4, 18]) is of particular interest for large-scale systems, where global path information is either unavailable or too costly to keep in each node.

In the paper, we explore the following routing strategies: (1) shortest path routing (which is a topology-aware approach, but we include it here solely for comparisons), (2) ant routing, (3) random wandering, (4) congestion-gradient routing, and (5) broadcast. Shortest path routing is optimal if the traffic is low and no congestion occurs, but every node needs to store a routing table that can get considerably large for large networks. In random wandering, the switch node that receives a message simply sends it to a randomly chosen neighbor. This is very simple to implement and robust against link and node failures, but very inefficient for larger system sizes. Ant routing [3] is a decentralized and agent-based (i.e., the “ants”) approach to find the shortest path in a network. Other parameters, such as congestion can be

considered easily. Congestion-gradient routing implements Danila et al.’s [4] algorithm, which considers the congestion of neighboring nodes for the local routing decision. Finally broadcast is an extremely simple strategy where each node simply sends the message to all neighbors that haven’t received it yet.

We decided to adopt a very simple—and admittedly not very realistic—traffic model, that is however still widely used to evaluate networks: uniform random traffic. Every processing node injects a message into the network to a randomly chosen destination processing node with probability p_I at each time step. Thus, if $p_I = 1$, a message will be generated during every time step by every processing node. For our experiments presented here, we used a low injection rate of $p_I = 0.1$, which prevented the network from jamming.

3. EXPERIMENTS

In the following sections, four simple experiments shall be described. All simulations were written in Matlab.

3.1 System Scalability

The goal of this first experiment is to illustrate how the three different topologies as described in Section 2.2 perform as the system size scales up. Scalability is a critical issue for nanoscale systems because it is generally very easy to build systems that involve huge numbers of components, e.g., Avogadro-scale systems. If the communication fabric doesn’t allow to efficiently send data across such an assembly, it will be impossible to solve tasks efficiently and thus to stay competitive with conventional design approaches.

For all three assemblies we have varied the system size and measured the average number of hops, which is proportional to the average path length L , i.e., the number of edges in the shortest path between two nodes, averaged over all pairs of nodes, as for example used in [19]. The different system sizes ranged from $N = S = 9$ to 125. For illustrative purposes, only shortest path routing was used in this experiment.

As Figure 2 shows, the locally connected topologies, i.e., the 2D and 3D CA, scale up worse with system size than the 3D random ensemble. The average path length of small-world graphs scales up logarithmically with the number of nodes, which Figure 2 confirms for the random ensemble. A fully (i.e., globally) connected network would obviously be most efficient, however, it does not represent a physically plausible solution for a self-assembled nanowire network.

3.2 Congestion-Gradient Routing

Danila et al.’s [4] congestion-gradient routing considers the neighboring node’s queue sizes for the local routing decision. A single parameter β allows to set the routing strategy from fully random ($\beta = 0$) to a rigid congestion-gradient flow ($\beta \approx 10$). They have shown that only a small amount of congestion awareness significantly improves the transport capacity and that an optimal β exists. However, they have only used random and scale-free networks. We have implemented the algorithm and Figure 3 illustrates the results for our three physically plausible network topologies. Shortest path (SP) and broadcast (BC) routing are shown for comparison (they are independent of β). As one can see, the throughputs for 2DCA, 3DCA, and 3D random ensemble (both with limited and unlimited number of links per node) all increase significantly for small β and outperform the throughput achieved with shortest path routing. The




Figure 2: Scaling of the number of hops as a function of the system size $N = S$ for shortest path routing. The $\alpha = 1.8$ (realistic) random ensemble (RE) shows approximately a logarithmic scaling behavior. Average over 10 networks for REs and over 7 runs for CAs.

2DCA shows the most dramatic increase. Similar to Danila et al.'s observation, there is an optimum for β , which depends on the network topology. If β is increased above this optimum, the throughput starts to slowly decrease again because of so called transport traps [4]. Obviously, the latencies (not shown here) significantly increase as well as β increases because the messages take less optimal routes.

3.3 Routing Scalability

Figure 4 illustrates how the different routing schemes scale up as the system size of the 3D random ensemble increases. Not surprisingly, broadcasting and shortest path routing scale up best (i.e., the number of hops only slowly increases with system size), while ant routing shows still a very good performance. Random and congestion-gradient based routing show a less favorable scaling behavior. Note that ant routing has not been optimized for each of the system sizes considered here, which therefore leaves room for improvement. This experiment shows that random and congestion-gradient routing are impractical solutions for larger system sizes, unless a hierarchical network structure is introduced, which allows to partition the system into smaller sub-systems. Both ant routing and simple broadcast are practical and reasonably efficient solutions, which do not need any global system information.

3.4 Robustness against Link Removals

Figure 5 shows the average number of hops as a function of the number of randomly removed links in the network. For both shortest path and ant routing, which are included for a baseline comparison, the routing tables were established *after* deleting the links. Broadcast is essentially unaffected by the link deletion up to the number of removed links (i.e., 40) that we measured. On the other hand, the number of hops of random and congestion-gradient routing starts to drop after about 10 removed links already because the messages have to take alternative and longer paths. Even though this




Figure 3: Average throughput as a function of β . $N = S = 27$, average values over 10 networks with standard error bars, traffic injection rate $p_I = 0.5$. Shortest path (SP) and broadcast (BC) routing (independent of β) are shown for comparison. The dash-dotted lines indicate the average values.

is certainly an overly simplistic failure model, it shows that broadcasting is inherently robust, which makes intuitively sense.




Figure 4: Average number of hops for different routing schemes as the system size scales up. 3D random ensemble with small-world power-law topology, $N = S$, average values over 20 networks.

4. CONCLUSION

We have experimentally and pragmatically investigated several topology-unaware routing schemes for both regular and irregular, system-on-chip-like computing architectures for self-assembled nanoscale electronics. The small-world topology with a power-law decaying distribution of shortcut lengths of the 3D random ensemble investigated is both physically plausible, could likely be built very economically by self-assembling processes, possess great communication characteristics, and is inherently robust against link failures. While regular and local-neighborhood interconnects are easier and more economical to build than interconnects




Figure 5: Robustness against random link removals for a 3D random ensemble. $N = S = 27$, average values over 20 networks.

with lots of global or semi-global long-distance connections, we have seen in [16] that they are not as efficient for global communication, which is very important and directly affects how efficient problems can be solved with such architectures.

In this paper, our simple routing case study has highlighted the importance of topology-unaware and decentralized routing schemes as well their limits. Scalability is a major point of concern and while topology-unaware routing is commonly less efficient because it is not based on global information (which may be impossible to obtain or too costly to store locally), it tends to be more robust. Broadcasting and ant routing both represent very interesting alternatives. Ant routing is completely decentralized and can be as good as shortest path routing, but only after a self-configuration phase. Broadcasting is extremely simple and efficient. The drawback is the additional traffic generated, which is, however, also the origin of its robustness against failures. We can summarize that none of the routing schemes explored here scales up to larger system sizes, at least not under the assumption that the switch nodes are simple and do not allow to store large routing tables. In general, a high number of hops and large latencies are the price we pay for decentralized and robust solutions. The results also suggest—as with traditional silicon electronics—that the communication fabric is most efficient if a hierarchical approach is used. Despite these drawbacks, we believe that communication and computation in random self-assemblies of components and interconnections is a very appealing paradigm, both from the perspective of fabrication as well as of performance and robustness, but further theoretical and experimental investigations are necessary to explore the numerous design trade-offs involved.

5. REFERENCES

- [1] International technology roadmap for semiconductors (ITRS). Semiconductor Industry Association, <http://www.itrs.net/reports.html>, 2003.
- [2] S. Carmi, R. Cohen, and D. Dolev. Searching complex networks efficiently with minimal information. *Europhysics Letters*, 74(6):1102–1108, 2006.
- [3] G. Di Caro and M. Dorigo. AntNet: Distributed Stigmergetic Control for Communications Networks. *Journal of Artificial Intelligence Research*, 9:317–365, 1998.
- [4] B. Danila, Y. Yu, S. Earl, J. A. Mash, Z. Toroczkai, and K. E. Bassler. Congestion-gradient driven transport on complex networks. *Physical Review E*, 74:046114, 2006.
- [5] B. Danila, Y. Yu, J. A. Marsh, and K. E. Bassler. Optimal transport on complex networks. *Physical Review E*, 74:046106, 2006.
- [6] J. A. Davis, R. Venkatesan, A. Kaloyeros, M. Beylansky, S. J. Souri, K. Banerjee, K. C. Saraswat, A. Rahman, R. Reif, and J. D. Meindl. Interconnect limits on gigascale integration (GSI) in the 21st century. *Proceedings of the IEEE*, 89(3):305–324, 2001.
- [7] G. de Micheli and L. Benini, editors. *Networks on Chips: Technology and Tools*. Morgan Kaufmann, 2006.
- [8] R. Ho, K. W. Mai, and M. A. Horowitz. The future of wires. *Proceedings of the IEEE*, 89(4):490–504, 2001.
- [9] J. K. Kleinberg. Navigation in a small world. *Nature*, 406:845, 2000.
- [10] B. Kozma, M. B. Hastings, and G. Korniss. Diffusion processes on power-law small-world networks. *Physical Review Letters*, 95:018701, 2005.
- [11] J. D. Meindl. Interconnect opportunities for gigascale integration. *IEEE Micro*, 23(3):28–35, 2003.
- [12] N. Oshida and S. Ihara. Packet traffic analysis of scale-free networks for large-scale network-on-chip design. *Physical Review E*, 74:026115, 2006.
- [13] T. Petermann and P. De Los Rios. Physical realizability of small-world networks. *Physical Review E*, 73:026114, 2006.
- [14] M. R. Stan, P. D. Franzon, S. C. Goldstein, J. C. Lach, and M. M. Ziegler. Molecular electronics: From devices and interconnect to circuits and architectures. *Proceedings of the IEEE*, 91(11):1940–1957, 2003.
- [15] C. Teuscher. On irregular interconnect fabrics for self-assembled nanoscale electronics. In *Proceedings of the 2nd IEEE International Workshop on Defect and Fault Tolerant Nanoscale Architectures, NANOARCH'06*, pages 60–67, Boston, MA, USA, June 17 2006. arXiv:cond-mat/0606584.
- [16] C. Teuscher. Nature-inspired interconnects for emerging large-scale network-on-chip designs. *Chaos*, 17(2):026106, 2007.
- [17] J. Tour, W. L. Van Zandt, C. P. Husband, S. M. Husband, L. S. Wilson, P. D. Franzon, and D. P. Nackashi. Nanocell logic gates for molecular computing. *IEEE Transactions on Nanotechnology*, 1(2):100–109, 2002.
- [18] W.-X. Wang, B.-H. Wang and C.-Y. Yin, Y.-B. Xie, and T. Zhou. Traffic dynamics based on local routing protocol on scale-free network. *Physical Review E*, 73:026111, 2006.
- [19] D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.
- [20] V. V. Zhirnov and D. J. C. Herr. New frontiers: Self-assembly in nanoelectronics. *IEEE Computer*, 34:34–43, January 2001.