

# Non-Traditional Irregular Interconnects for Massive Scale SoC

Christof Teuscher  
Los Alamos National Laboratory  
Los Alamos, NM 87545  
E-mail: christof@teuscher.ch

Anders A. Hansson  
Los Alamos National Laboratory  
Los Alamos, NM 87545  
E-mail: hansson@lanl.gov

**Abstract**— By using self-assembling fabrication techniques at the cellular, molecular, or atomic scale, it is nowadays possible to create functional assemblies in a mainly bottom-up way that involve massive numbers of interconnected components. However, such assemblies are often highly irregular, unreliable, and heterogeneous. A grand challenge for future and emerging electronics is thus to reliably and efficiently compute and communicate in such systems. The goal of this paper is to illustrate why non-traditional network-on-chip paradigms are promising for massive scale systems and what the limits are. We have previously shown that certain irregular 3D assemblies and interconnects have major advantages over regular 2D and 3D mesh fabrics in terms of latency, throughput, scalability, and the robustness against simple link failures. We present these results from a complex network perspective and look into the scaling properties of different interconnect topologies and routing algorithms in an abstract framework. We argue that only small-world topologies will scale up to massive scale systems. The long term goal in using irregular, fabrication-friendly, and non-traditional interconnects is to eventually be able to cheaply and easily assemble massive scale computing devices that are able to solve specific large-scale problems competitively with traditional top-down fabricated silicon technology.

## I. INTRODUCTION

Since the beginning of modern computer science some sixty years ago, we have seen a steady increase in complexity and performance in computing machines, both in the area of high-performance computing and embedded systems. This trend is essentially due to the enormous efforts and progress in miniaturization towards the deep sub-micron and nano-scale domain, which is reflected in Moore's law. However, we also steadily approach fundamental physical limits of miniaturization (e.g., [12]) and thus see a growing need for alternative computing paradigms and machines.

Because of the ongoing miniaturization, the importance of interconnects on integrated circuits has outrun the importance of transistors as a dominant factor of performance [15]. The reason is 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 as well as the operating frequency has dramatically increased. Global wires now span multiple clock domains in most cases, which requires to insert buffers and thus adds further delays. *Networks-on-chip* (NoCs) [7] offer promising solutions to many of today's *Systems-on-Chip* (SoC) interconnect and communication challenges. The NoC interconnects on-chip

modules with a communication network that is akin to the internet and allows for fast and reliable communication while sharing resources and offering scalability, abstraction levels, and modularity. However, compared to the Internet, NoC communication has different constraints and requirements, which generally results in techniques that are much simpler and more effective than for TCP/IP traffic [7].

The goal of this paper is to look beyond the horizon of today's top-down fabricated silicon technology and to take more of a bottom-up stance. As the progress in miniaturization will continue to slow down because of fundamental physical limits, we instead expect a massive increase in parallelism to keep up with the current pace of progress under Moore's law. This trend can already clearly been seen with the advent of on-chip multi-core systems. Massive scale parallelism is also particularly promising for bottom-up fabricated systems because it is generally easier and cheaper to achieve (e.g., bacteria that self-replicate or self-assembling carbon nanotubes) than in a top-down way. Because of the lack of precise control and the sheer scale, future and emerging computing devices that are assembled in a mainly bottom-up way are expected to be highly irregular, unreliable, and heterogeneous in structure and function. Clearly, this requires novel paradigms, tools, and methodologies to efficiently communicate and compute. Here, we look at self-assembled networks-on-chips and present the challenges and opportunities. Our work goes far beyond other alternative interconnects, such as proposed by [2], [3], [10], and [17], in terms of irregularity and scale, and currently also is more abstract. We believe that the NoC paradigm is very promising for future massive scale systems. The reasons are outlined in this paper, the remainder of which is as following: Section II outlines the challenges that need to be addressed. Section III highlights complex networks research and what we can learn from it for NoC. Our abstract framework is described in Section IV and two simple experiments in Section V. Section VI concludes the paper.

## II. THE CHALLENGES

From an abstract point of view, we can consider a computing machine as a graph composed of  $N$  nodes and  $E$  edges, where the nodes represent processing units and the edges communication links between these units. Unless only a single processing unit is used and no communication is

necessary with other nodes at all, details aside, we can assume that the more efficient the physical communication network in terms of throughput, delay, area requirements, and energy consumption, the more efficient a given algorithmic problem can be solved. Since the interconnect has become more important than the transistors in integrated circuits, we can justify a research approach that focuses on the physical topology and communication aspects first. Scalability is a crucial measure with the general trend of growing overall system complexity. While systems-on-chip were composed of some tens of modules a few years ago, the latest Xilinx Virtex<sup>TM</sup>-5 *Field Programmable Gate Arrays* (FPGAs) offer up to 330,000 logic cells that are to be interconnected and need to communicate efficiently among themselves. We expect that future systems will far exceed these numbers and that it will be possible to assemble circuits involving peta- ( $10^{15}$ ) or even Avogadro-numbers ( $6 \times 10^{23}$ ) of interconnected components. We formulate the grand research challenges in this area—on which we focus in this paper—as following: (1) what interconnect topologies can be built and (2) what communication schemes can be used that allow to efficiently and reliably communicate in such massive scale systems? We further argue that systems at this scale will have to be built in a mainly bottom-up way, which will inevitably result in irregularly assembled systems—at least to some extent—because we currently lack control over these mechanisms. However, if we can harness this interconnect complexity for our purposes, we could potentially build cheaper and more efficient computing machinery.

### III. LESSONS FROM COMPLEX NETWORK THEORY

Modern complex network theory (see e.g., [1] for an overview) provides us with a powerful toolbox to analyze network topologies and properties. For our purposes, the *average path length*  $L$  of a network, defined as the average number of nodes in the shortest path for all possible pairs of network nodes, is the most relevant measure because it is directly related to the latency a message observes while being send through the network, which in turn is directly related to the communication efficiency. This is obviously a very simplified view and does not take into account the processing element's delays, wire delays, message lengths, routing algorithms, and possible congestion. Yet, we are solely interested in the average path length as the system size increases because that measure of the physical network topology will play the most significant part (i.e., a lower limit) for massive scale systems.

Between a regular, locally interconnected mesh network and a completely random Erdős-Rényi type topology, there are other classes of graphs [1], such as for example *small-world* and *scale-free* graphs. Networks with the small-world property (which both random and scale-free networks possess as well) have a very short average path length, bounded by a polynomial in  $\log(N)$  (where  $N$  is the number of nodes), which makes them particularly interesting for efficient communication. In turns out that most real-world networks




Fig. 1. Average path length scaling of different interconnect topologies. Small-world graphs show a scaling bounded by  $\log(N)$ . Note that these scaling laws do not take into account realistic assumptions, but they merely serve as a baseline.

are small-world graphs, including the internet and electronic circuits that have been designed in a top-down way [11].

In reality, it is fair to assume that local connections have a lower cost (e.g., in terms of the wire delays and area requirements) than long-distance connections. Physically realizing small-world networks with uniformly distributed long-distance connections is thus unrealistic, and distance, i.e., the wiring cost, needs to be taken into account, a perspective that recently gained increasing attention [13], [16]. In summary: there is a design trade-off between (1) the physical realizability and (2) the communication characteristics for a physical network topology. A locally and regularly interconnected mesh topology is in general easy to build and only involves minimal wire and area cost, but it offers poor global communication characteristics (i.e., high latencies) and scales-up poorly with system size. On the other extreme, a random Erdős-Rényi random topology or a fully connected network scales-up well and has a very short-average path length, but it is not physically plausible because it involves costly long-distance connections. Small-world topologies with a distance-dependent distribution of the connectivity are located in a unique spot in the design space and also offer two other highly relevant properties [14], [19]: (1) efficient navigability and thus potentially efficient routing, and (2) robustness against random link removals. For these reasons, we can conclude that small-world graphs are the most promising interconnects for massive scale systems.

As an illustration, Figure 1 shows the average path length scaling as a function of system size  $N$  of several typical network topologies. As one can see, the logarithmic scaling in the average path length for small-world networks makes them very attractive—and the only practical choice—for massive scale systems.

### IV. THE FRAMEWORK

The abstract framework has been described in detail elsewhere [18], [19] and we will therefore only briefly summarize the main characteristics of the network-on-chip-like approach.

We use a high-level description at this point, however, future simulations will take into account additional low-level physical and device details.

The architecture 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 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 generate messages according to a simple traffic scheme. Since we are mainly interested in interconnect issues in this paper, we do not further specify or limit the processing nodes' computing capacity.

Here, we will consider the following three reference network topologies (see [19] for more topologies and details):

- **2DCA**: 2D (unfolded) regularly arranged and locally interconnected mesh (von Neumann neighborhood);
- **3DCA**: 3D (unfolded) regularly arranged and locally interconnected mesh;
- **3DRERealistic**: 3D *random ensemble/multitude* (RE), small-world, power-law [16],  $\alpha = 1.8$ , upper limit  $k_{max} = 10$  on the number of connections per node.

The physical plausibility and realizability of the *3DRERealistic* assembly—which we hypothesize could be bottom-up self-assembled from conductive nanowires or nanotubes—has been discussed in [19], however, several questions remain open and are part of ongoing research. Our approach consists in using a hybrid assembly (as others explore too, e.g., [9]), where the computing and switch blocks will still be traditional silicon, while the interconnect is made up from self-assembled nanowires. Nanowires can be grown in various ways using diverse materials, such as metals and semiconductors. We have chosen a novel way to grow conductive nanowires, which Wang et al. [21] at LANL have previously demonstrated: Ag nanowires can be fabricated on top of conducting polyaniline polymer membranes via a spontaneous electrodeless deposition (self-assembly) method. This will allow to densely interconnect silicon components in a simple and cheap way with specific distance-dependent wire-length distributions. Figure 2 shows an example of nanowires grown in such a way.

## V. SCALABILITY EXPERIMENTS

The goal of this first experiment is to illustrate how the three different topologies as described above perform as the system size scales up. 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  $L$ . For illustrative purposes, only shortest path routing was used in this experiment. As Figure 3 shows, the locally connected topologies, i.e., the 2D and 3D CA, scale up much worse with




Fig. 2. Ag nanowires grown on a polyaniline polymer membrane. Image courtesy by Hsing-Lin Wang, LANL.

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 3 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.




Fig. 3. 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.

There exists a large number of different routing strategies and flavors. Many of the studied solutions in the literature cannot be applied here because of the network's irregularity and the lack of global coordinate and routing information, which for example prevents us from applying simple XY routing, as used in [4] for irregular topologies. We are particularly interested in routing based on local information only (e.g., [6], [22]) because of the relevance for massive scale systems, where global path information is either unavailable or too costly to keep in each node. We compare the following very simple routing strategies to obtain a baseline (see also [20] for more details): (1) shortest path routing, (2) ant routing [5], (3) random wandering, (4) congestion-gradient routing [6], and (5) broadcast.

Figure 4 illustrates how these 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, while ant routing shows still a very good performance. Broadcasting naturally comes at a high communication overhead cost. In [8], Elsaesser has also presented a scalable and robust algorithm for randomized broadcasting in irregular networks. Random and congestion-gradient-based routing show a less favorable scaling behavior. 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 broadcast are practical and reasonably efficient solutions, which do not need any global system information.




Fig. 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.

## VI. CONCLUSION

Most of today's systems-on-chip are composed of a few up to a hundred building blocks, such as I/O blocks, processing elements, memory, which are of a rather high complexity. On the other hand, FPGAs have large numbers of very simple building blocks. For both of these types of architectures, there is an increasing trend towards NoC (-like) interconnects with packet-oriented traffic because NoCs offer various advantages over rigid ad-hoc networks. While the current design approaches and fabrication technologies are suited for gigascale systems, we argued that novel interconnect approaches, which offer a better scalability and which can be fabricated more easily, are needed for peta-scale or bigger systems. A physical interconnect topology will have to have the small-world property in order to scale up to massive scale systems and distance-dependent wire length distributions to be physically plausible. It remains to be experimentally shown that such a topology can be self-assembled, which is part of ongoing research at LANL. If this can be done successfully, irregular bottom-up interconnects would offer high performance and inherent robustness at a very low fabrication cost. Future work

will also focus on a rigorous design methodology that will allow to map a given circuit on such an irregular device.

## ACKNOWLEDGMENT

The authors would like to thank Elshan A. Akhadov and Hsing-Lin Wang. We gratefully acknowledge the support of the U.S. Department of Energy through the LANL/LDRD Program for this work

## REFERENCES

- [1] R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, January 2002.
- [2] K. Banerjee and N. Srivastava. Are carbon nanotubes the future of VLSI interconnections? In *Proceedings of the 43<sup>rd</sup> Annual Conference on Design Automation (DAC'06)*, pages 809–814, 2006.
- [3] D. Bertozzi and L. Benini. Xpipes: A network-on-chip architecture for gigascale systems-on-chip. *IEEE Circuits and Systems Magazine*, 4(2):18–31, 2004.
- [4] E. Bolotin, I. Cidon, R. Ginosar, and A. Kolodny. Efficient routing in irregular topology NoCs. Technical Report CCIT #554, Israel Institute of Technology, Technion Department of Electrical Engineering, Sep 2005.
- [5] G. Di Caro and M. Dorigo. AntNet: Distributed Stigmergetic Control for Communications Networks. *Journal of Artificial Intelligence Research*, 9:317–365, 1998.
- [6] B. Danila, Y. Yu, S. Earl, J. A. Mash, Z. Toroczka, and K. E. Bassler. Congestion-gradient driven transport on complex networks. *Physical Review E*, 74:046114, 2006.
- [7] G. de Micheli and L. Benini, editors. *Networks on Chips: Technology and Tools*. Morgan Kaufmann, 2006.
- [8] R. Elsaesser. On the communication complexity of randomized broadcasting in random-like graphs. In *Proceedings of the 18<sup>th</sup> Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 148–157, New York, NY, USA, 2006. ACM.
- [9] D. K. Ferry. Nanowires in nanoelectronics. *Science*, 319(5863):579–580, 2008.
- [10] A. Gayasen, V. Vijaykrishnan, and M. J. Irwin. Exploring technology alternatives for nano-scale FPGA interconnects. In *Proceedings of the 42<sup>nd</sup> Annual Conference on Design Automation (DAC'05)*, pages 921–926, 2005.
- [11] R. Ferrer i Cancho, C. Janssen, and R. V. Sole. Topology of technology graphs: Small world patterns in electronic circuits. *Physical Review E*, 64:046119, 2001.
- [12] L. B. Kish. End of Moore's law: Thermal (noise) death of integration in micro and nano electronics. *Physics Letters A*, 305:144–149, 2002.
- [13] J. K. Kleinberg. Navigation in a small world. *Nature*, 406:845, 2000.
- [14] J. K. Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the 32<sup>nd</sup> Annual ACM Symposium on Theory of Computing*, pages 163–170, 2000.
- [15] J. D. Meindl. Interconnect opportunities for gigascale integration. *IEEE Micro*, 23(3):28–35, 2003.
- [16] T. Petermann and P. De Los Rios. Physical realizability of small-world networks. *Physical Review E*, 73:026114, 2006.
- [17] L. K. Scheffer. CAD implications of new interconnect technologies. In *Proceedings of the 44<sup>th</sup> Design Automation Conference (DAC'07)*, pages 576–581, 2007.
- [18] C. Teuscher. On irregular interconnect fabrics for self-assembled nanoscale electronics. In *Proceedings of the 2<sup>nd</sup> IEEE International Workshop on Default and Fault Tolerant Nanoscale Architectures, NANOARCH'06*, pages 60–67, Boston, MA, USA, June 17 2006. arXiv:cond-mat/0606584.
- [19] C. Teuscher. Nature-inspired interconnects for emerging large-scale network-on-chip designs. *Chaos*, 17(2):026106, 2007.
- [20] C. Teuscher. Topology-unaware routing in irregular self-assembled networks-on-chip: An explorative case study. Second International Conference on Nano-Networks (Nano-Net 2007), 2007.
- [21] H.-L. Wang, W. Li, E. Akhadov, and Q. Jia. Electrodeless deposition of metals and metal nanoparticle using conducting polymers. *Chemistry of Materials*, 19:520, 2007.
- [22] 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.