You Can Get There From Here: Connectivity of Random Graphs on Grids

Steven P. Levitan
University of Pittsburgh
348 Benedum Hall
Pittsburgh, PA 15261
+1-412-648-9663
steve@ece.pitt.edu

ABSTRACT

Using random graphs, we address the need for non-local communication channels in emerging nanotechnology interconnection structures. In this model, a grid of $\sqrt{N} \times \sqrt{N}$ processing elements are connected by non-local wires, which are fabricated as a post processing self-assembly step by electro-chemical migration of functionalized carbon nanotubes. We show that under reasonable assumptions, 80% of the cells can be connected into a single spanning tree with only N wires deposited randomly across the surface of the substrate.

Categories & Subject Descriptors

B7.1 **[integrated circuits]:** Types and Design Styles – *advanced technologies*

General Terms Design, Reliability, Theory.

Keywords: Nanotechnology, Cellular Automata, Random Graphs, Carbon Nanotubes

1. INTRODUCTION

As process dimensions shrink towards the nanoscale, it is increasingly tempting to re-consider processing based on meshes of cellular automata [1]. In this model, large numbers of processing elements are interconnected via local nearest neighbor connections and work independently, but cooperatively, to solve a common problem. As powerful as this model is, it suffers from a serious scaling drawback. Global communication tasks (both one-to-many, and many-to-many) do not scale with the technology. Local communications imply "store and forward" protocols that scale as the diameter of the mesh. Therefore, delays grow as the mesh size scales up. Bus-based routing techniques and network-on-chip methodologies also suffer from the need to insert buffers to drive the relatively long distances across the chip [2].

Recently, distributed point-to-point networks have been re-examined as a way to provide short-cuts across the chip diameter. Many structures are possible including trees, hypercubes, and multi-dimensional arrays [3]. The risk with these networks is that their coherence depends on the reliability of the links between nodes. Definite structures, perhaps with explicit redundancy, are required to guarantee that connectivity is preserved in the presence of manufacturing defects or post fabrication failures.

While others are concerned with the utilization of long-range links to create "scale-free" networks-on-chip designs [4], the question we pose is: what is possible with a random interconnect? Or more precisely, how many random wires will it take to connect all the nodes into a common tree structure? And, how robust is this model?

DAC 2007, June 4–8, 2007, San Diego, California, USA

2. SIMULATIONS

We simulate the deposition of random links on a substrate. Given a $\sqrt{N} \times \sqrt{N}$ grid of lattice sites, links with a uniform length distribution (from 1 to $\sqrt{2N}$) are placed randomly on the grid. As links are inserted, their ends are checked for co-location with other linked sites. If the link ends on a site that is already linked, then a common (lower) link number is recorded. Thus, a forest of sites is grown. Figure 1 shows a 10x10 grid after seven links have been added. Note that site 5 in the upper left corner would, in fact, be labeled with a "1" indicating it is connected to a tree with three sites, labeled 1. Note links that cross, do not connect, e.g., 4-7 and 6-3. It is only ends that can make connections.

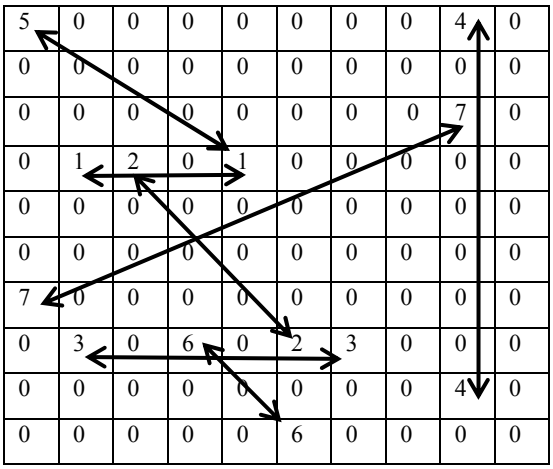

Figure 1: 10x10 Lattice after the addition of 7 links

Figure 2(a) shows, for 5 runs, the number of sites that are unlinked and the size of the single biggest tree as a function of time, or the number of links added sequentially. As shown at the top, after approximately 100 steps, the largest tree has

encompassed 80% of the sites and after about 125 steps only 10% of the sites are un-linked. The steep curve for the size of the largest spanning tree is reminiscent of the phase transition behavior associated with percolation [5]. This fact is even clearer as we scale up the size of the lattice to diameters of 100, and 500, (Figure 2(b, c)) where in each case the 80% mark is reached near step N. These simulations were carried out with periodic boundary conditions and uniform length distributions. Figures 3 and 4 show details of 20 runs for 25x25 and 50x50 lattices, with Gaussian length, l, distributions ($\bar{l} = \sqrt{N}/4, \sigma = 1$) and clipping boundary conditions; neither of these two conditions appreciatively affected the location or shape of the maximum spanning tree size curve. Further, the histograms of snapshots of minimum path lengths and node degrees after kN wire insertions show that after $2N$ insertions, average (minimum) path lengths are less than $log_2(N)$ and external node degree is > 2. These values compare favorably with unbounded random graphs [5]. This fact also leads us to believe that this phenomenon is very robust. Therefore, it should be straightforward to build systems with a surfeit of randomly placed wires to ensure a large fraction of the computational elements can be recruited to do useful work.

3. DISCUSSION

While one can be intimidated by O(N) additional wires as a fabrication requirement, it is important to understand the fabrication methodology envisioned. For instance, carbon nanotubes make good wires and are easy to fabricate in bulk [6]; what is hard is placing them at specific sites. This work shows that random placement (i.e., electro-chemical attachment by functionalizing the ends of the tubes and having them migrate to lattice sites [7]) is sufficient to guarantee that a large fraction of computational nodes will be interconnected in a dense network.

Future work will be to: (1) find an analytical model for the growth of the spanning trees; (2) develop communication protocols for the nodes to self organize; and (3) consider how this technique would scale for 3D networks compatible with 3D VLSI chip fabrication methodologies.

4. ACKNOWLEDGEMENTS

This work supported, in part, by the National Science Foundation grants C-CR0306325 and CCF-0541150. The author would also like to thank D. Chiarulli and C. Teuscher for helpful discussions.

5. REFERENCES

[1] von Neumann, J. *The Theory of Self-reproducing Automata*, A. Burks, ed., Univ. of Illinois Press, Urbana, IL 1966.

[2] Ogras, U.Y., et al "It's a Small World After All: NoC Performance Optimization via Long-Range Link Insertion," *IEEE Transactions on VLSI, 14*(7), 2006.

[3] Leighton, F.T. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes,* Morgan Kauffman, San Mateo, CA, 1992.

[4] Teuscher, C. "Small-World Power-Law Interconnects for Nanoscale Computing Architectures," *Sixth IEEE Conference on Nanotechnology, Volume 1*, 379 – 382, 17-20 June 2006.

[5] Newman, M.E.J. and Watts, D.J., Scaling and percolation in the small-world network model, Physical Review E, 60(6) 7332-7342, (December 1999).

[6] Raychowdhury, A., et al "A circuit model for carbon nanotube interconnects: comparative study with Cu interconnects for scaled technologies," *ICCAD-2004.* pp. 237- 240, 7-11 Nov. 2004.

[7] Saleem, G., et al Nanotube electronics: Large-scale assembly of carbon nanotubes, *Nature 425*, 36-37 (4 September 2000).

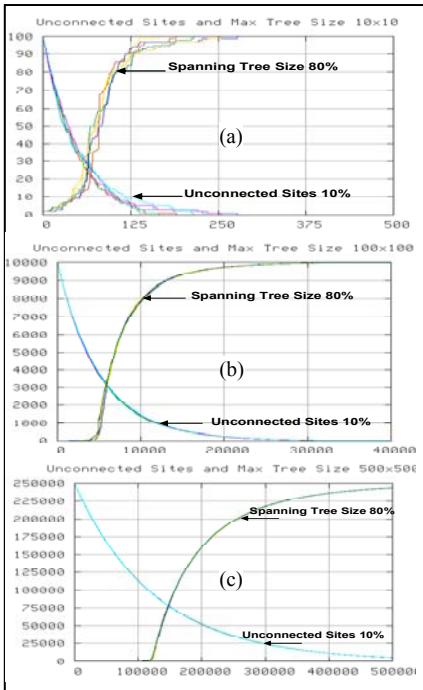

Figure 2: Size Scaling

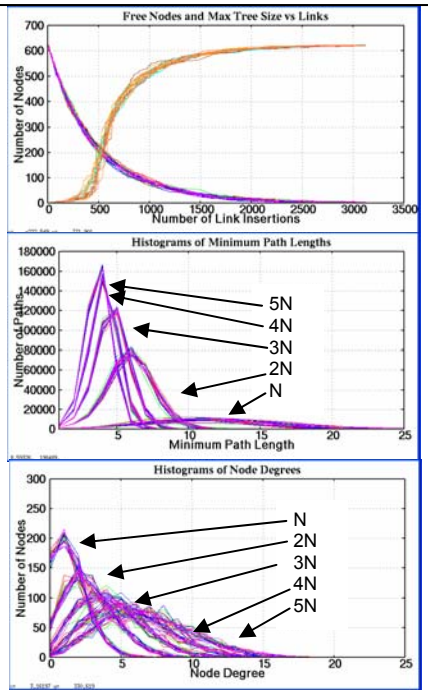

Figure 3: 25x25 lattice details

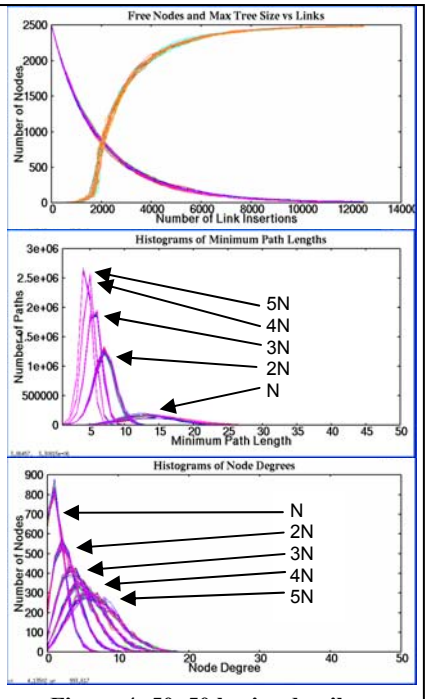

Figure 4: 50x50 lattice details