Theory of Networked Computation?

(FuDiCo 2007 Position Paper)

Joan Feigenbaum
http://www.cs.yale.edu/homes/jf

The increasing prominence of the Internet, the Web, and large data networks in general has wrought one of the most profound shifts in Computer Science since the field's inception. Traditionally, Computer-Science research focused primarily on understanding how best to design, build, analyze, and program computers. Research focus has now shifted to the question of how best to design, build, analyze, and operate networks. How can one ensure that a network created and used by many autonomous organizations and individuals functions properly, respects the rights of users, and exploits its vast shared resources fully and fairly?

Would it be fruitful for the Theory of Computation (ToC) community to address the full spectrum of research questions implicit in this grand challenge by developing a Theory of Networked Computation (ToNC)? ToC research has already evolved with and influenced the growth of the Web, producing interesting results and techniques in diverse problem domains, including search and information retrieval, network protocols, error correction, Internet-based auctions, and security. A more general Theory of Networked Computation could influence the development of new networked systems, just as formal notions of "efficient solutions" and "hardness" have influenced system development for single machines.

In this talk, I will attempt to justify the controversial position that there should be a "theory of networked computing" that is distinct (although not disjoint) from the existing "theory of distributed computing." I hope to stimulate discussion and to provoke at least one other workshop participant to articulate and defend the opposing position.

My case for the desirability of a theory of networked computing will be drawn from the results of the 2006 ToNC workshops. The rest of this position paper, which is excerpted from [FeMi], briefly summarizes the findings of those workshops; (much) more information about these findings can be found in [ToNC].

Research Goals

Workshop participants identified three broad, overlapping categories of ToNC-research goals:

- **Realizing better networks:** Numerous theoretical-research questions will arise in the design, analysis, implementation, deployment, operation, and modification of future networks.
- **Computing on networks:** Formal computational models of future networks will enable us both to design services, algorithms, and protocols with provable

properties and to demonstrate (by proving hardness results) that some networked-computational goals are unattainable.

- **Solving problems that are created or exacerbated by networks:** Not all of the ToNC-research agenda will involve new computational models. The importance of several established theoretical-research areas has risen dramatically as the use of networked computers has proliferated, and some established methods and techniques within these areas are not general or scalable enough to handle the problems that future networks will create. Examples of these areas include massive-data-set algorithmics, error-correcting codes, and random-graph models.

ToNC-research problems in the first category revolve around finding the right primitives and abstractions with which to study networks; they are exemplified in Goel's breakout-group summary slides from the first ToNC workshops [Goel]. Those in the second category revolve around efficient algorithms for "computing on a network" and hardness results showing that some networked-computational goals are unachievable; they are exemplified in Feigenbaum's talk [Feig] and Impagliazzo's breakout-group summary [Impa] at the second ToNC workshop. Those in the third category revolve around (single-machine) algorithmic questions about network design and network-generated data (as exemplified in Byers's breakout-group summary slides from the second ToNC workshop [Byer]) and mathematical questions about network modeling and network behavior (as exemplified in Kleinberg's breakout-group summary slides from the second ToNC workshop [Klei]).

Cross-Cutting Issues

Several cross-cutting, high-level issues are relevant to all three categories and arose repeatedly during plenary and breakout sessions at both workshops.

- **Incentive compatibility:** Perhaps the most important distinguishing feature of modern networks is that they are simultaneously built, operated, and used by multiple parties with diverse sets of interests and with constantly changing mixes of cooperation and competition. Formal models of networked computation and notions of hardness and easiness of computation will have to incorporate subnetwork autonomy and user self-interest in an essential way.
- **SPUR:** Achieving the broadest possible vision of "networked computation" will require substantial progress on Patterson's SPUR agenda [Patt]. In his words, "we have taken ideas from the 1970s and 1980s to their logical extreme, providing remarkably fast and cheap computing and communication (C&C) to hundreds of millions of people. … [F]or our new century, we need a new manifesto for C&C: ... Security, Privacy, Usability, and Reliability (SPUR)."
- **Build on success:** Although today's Internet may leave something to be desired with respect to security, privacy, usability, and reliability, it has far surpassed expectations with respect to several important design goals, *e.g.*, flexibility and

scalability. Are the new design criteria compatible with the (manifestly successful) old criteria, and, if not, what are our priorities?

- **"Clean slate":** The phrase "clean-slate design" has become a mantra in networking-research forums and in calls for proposals. Not surprisingly, many people have raised the question of whether anything that requires a "clean slate" could ever be brought to fruition in a world in which networked computation is pervasive and mission-critical. From a research perspective, the crucial point is that clean-slate *design* does not presume clean-slate *deployment*. Part of the ToNC agenda is the evaluation of new technologies, methods, algorithms, *etc.* from the perspective of incremental deployability and paths to adoption.
- **Diversity of "networks":** The scope of the networking research agenda is broader than "next-generation Internet," and thus the ToNC agenda must be broader as well. Interesting theoretical questions arise in the study of special-purpose networks (such as the DoD's Global Information Grid); of moderate-sized but functionally innovative networks; of sensor nets and other technologically constrained networks; of mobile networks; and of P2P and other application-layer networks.

References

[Byer] J. Byers, "ToNC: Summary of Algorithmic Foundations Working Group," http://www.cs.yale.edu/homes/jf/ToNC-Berkeley/Byers-Group4.ppt

[Feig] J. Feigenbaum, "Progress Report on the ToNC Initiative," http://www.cs.yale.edu/homes/jf/ToNC-Berkeley/Feigenbaum-Progress.ppt

[FeMi] J. Feigenbaum and M. Mitzenmacher, "Towards a Theory of Networked Computing," *SIGACT News* **37:4** (December 2006), pp. 22-26.

[GENI] Global Environment for Network Innovations, http://www.geni.net

[Goel] A. Goel, "Next-Generation Architecure," http://www.cs.yale.edu/homes/jf/ToNC-Princeton/Goel-Group2.ppt

[Impa] R. Impagliazzo, "Complexity-Theoretic Foundations of Networked Computing," http://www.cs.yale.edu/homes/jf/ToNC-Berkeley/Impagliazzo-Group1.ppt

[Karp] R. M. Karp, "Report of the SIGACT Committee on Funding for Theoretical Computer Science" (presented at the FOCS Business Meeting on October 24, 2005), *SIGACT News* **36:4** (December 2005), pp. 99-102.

[Klei] R. Kleinberg, "Analytical Foundations of Networked Computing," http://www.cs.yale.edu/homes/jf/ToNC-Berkeley/Kleinberg-Group5.ppt

[Patt] D. Patterson, "20th century vs. 21st century C&C: the SPUR manifesto," *Communications of the ACM* **48:3** (March 2005), pp. 15-16.

[ToNC] Towards a Theory of Networked Computation, http://www.cs.yale.edu/homes/jf/ToNC.html