Agents' Privacy in Distributed Algorithmic Mechanisms
(Position Paper)

Joan Feigenbaum
Yale University
jf@cs.yale.edu

Noam Nisan
Hebrew University
noam@cs.huji.ac.il

Vijay Ramachandran
Yale University
vijayr@cs.yale.edu

Rahul Sami
Yale University
sami@cs.yale.edu

Scott Shenker
ICSI
shenker@icsi.berkeley.edu

May 4, 2002

In traditional theoretical computer science (TCS), computational agents are typically assumed either to be *obedient* (*i.e.*, to follow the prescribed algorithm) or to be *adversaries* who "play against" each other. On the other hand, the *strategic* agents in economic theory are neither obedient nor adversarial. Although one cannot assume that they will follow the prescribed algorithm, one can assume that they will respond to incentives. Thus, the economics literature traditionally stressed incentives and downplayed computational complexity, and the TCS literature traditionally did the opposite. The emergence of the Internet as a standard platform for distributed computation has radically changed this state of affairs: Ownership, operation, and use by many self-interested, independent parties give the Internet the characteristics of an economy as well as those of a computer.

To date, the TCS community's work on *distributed algorithmic mechanism design* (DAMD) has focused on "truthful" mechanisms. The overall approach, which is consistent with the approach in the economics literature, is to design mechanisms that are *incentive-compatible* in the technical sense that strategic agents cannot improve their welfare by lying about their private inputs. The well known Revelation Principle ensures that, if there is a mechanism that achieves a given economic design goal, then there is a truthful mechanism that achieves the same goal. The premise of this overall approach is that agents will voluntarily reveal their private information if it can be proven that lying does them no good in the situation addressed by this particular mechanism-design exercise.

We question this premise. Revelation of private information may be an agent's best possible strategy for the particular game at hand, but it may be unacceptable in the broader context. For example, in one formulation of the interdomain-routing mechanism-design problem [FPSS02], the agents are Internet domains (or *autonomous systems*, ASs), and an agent's private input is the cost it incurs when carrying transit traffic. Revealing true transit costs may reveal details about an AS's internal network that it wants to keep private for reasons that have nothing to do with transit-traffic revenues. Moreover, the real mechanism-design goal is not to convince agents to reveal their private inputs but rather to compute a global optimum that depends on these inputs (in the [FPSS02] formulation, a set of lowest-cost interdomain routes).

The theory of *secure, multiparty function evaluation* (SMFE), developed by the cryptologic-research community, shows that such global optima can often be computed in such a way that nothing about agent A's private input need be revealed to agent B (except what is logically implied by the result and agent B's private input). Existing SMFE protocols cannot be applied off-the-shelf to interdomain routing, or to DAMD problems generally, because they make assumptions about the fraction of obedient agents that are incompatible with the mechanism-design framework, and they have unacceptable network complexity. However, they provide a logical starting point for the design of privacy-friendly solutions.

The agent-privacy issue was raised in an early paper of Nisan [Nis99], as was the potential applicability of SMFE techniques. However, few specific DAMD problems have thus far been addressed from an agent-privacy point of view. We believe that privacy-friendly DAMD is an important research direction and that interdomain routing is a compelling example in which it is needed.

References

[FPSS02] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. "A BGP-based Mechanism for Lowest-Cost Routing," to appear in *Proceedings of the 2002 ACM Symposium on Principles of Distributed Computing*.

[Nis99] N. Nisan. "Algorithms for Selfish Agents," in *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, LNCS 1563, pages 1-17, Springer, Berlin, 1999.

[NR01] N. Nisan and A. Ronen. "Algorithmic Mechanism Design," *Games and Economic Behavior* **35** (2001), pages 166-196.