Distributed algorithmic mechanism design and network security

John Mitchell, Yoav Shoham, and Vanessa Teague

The Internet is a collection of autonomously administered systems that cooperatively share networking resources for the common good. The Internet is also subject to a range of computer security problems, ranging from distributed denial-of-service attacks to IP spoofing to network-based attacks on individual sites. We believe it may be possible to use mechanism design to improve network performance and security: mechanisms that provide incentives for cooperative behavior may also provide incentives to respond adaptively to malicious behavior, improving the performance, reliability, and security of the network.

There are two general reasons why mechanism-based incentives may lead to greater network security:

- Sites and domains can be given an incentive to adapt in ways that increase security,
- Accounting associated with mechanisms may provide useful instrumentation and monitoring that detect and isolate malicious behavior.

We see preliminary evidence for each of these possibilities in recent work. For example, algorithmic mechanisms have been used to optimize global system properties, such as minimizing the completion times of jobs scheduled on parallel machines [1]. In an algorithmic mechanism for BGP routing [2], nodes are billed for the traffic they send (to encourage others to tell the truth about their transit costs). This makes it much harder for a node to lie undetectably about the source of a packet.

We have augmented a well-studied scenario, multicast cost sharing based on a marginal cost mechanism [3], in an effort to understand how to implement a mechanism in an anarchic distributed setting. We assume that each node may autonomously choose to deviate from the specified protocol, ingenuously modifying any data in any computationally feasible way. In order to calculate pricing reliably, we added digital-signature-based authentication and some additional incentives to provide accurate data needed for routing decisions by other nodes. The result is a protocol that requires only local communication of a small number of messages, provides incentives for each node to reveal its true utility *and* prevents malicious behavior that would not reduce a node's welfare in the original model. If there are only honest and selfish agents then honestly following the protocol is an equilibrium. However, we have not been able to find a mechanism in which honest participation and adaptive routing around dishonest nodes is a dominant strategy.

Looking forward, we believe it will be fruitful to develop a general theory of systems comprising three forms of agent behavior: honest compliance, rational selfishness, and malicious disruption. Pragmatically speaking, it is sufficient to develop mechanisms that work properly when significant fractions of the population fall into the first two categories. Further, the goal

need not be to prevent malicious behavior, but to limit its consequences. At present, we do not know what assumptions will be sufficient to provide incentives for autonomous behavior that both maximizes a global objective such as network throughput and accurately identifies and isolates malicious behavior.

References

[1] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2001.

[2] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker. Incentive-compatible interdomain routing. *submitted*, 2002.

[3] Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001. Special issue on Internet Algorithms.