On Implementing Robust and Resilient Mediators

Ittai Abraham
School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
ittaia@cs.huji.ac.il

Danny Dolev*
School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
dolev@cs.huji.ac.il

Joseph Y. Halpern†
Cornell University
Ithaca, NY 14850
halpern@cs.cornell.edu

January 26, 2007

Abstract

We consider games that have (k,t)-*robust* equilibria when played with a mediator, where an equilibrium is (k,t)-robust if it tolerates deviations by coalitions of size up to k and deviations by up to t players with unknown utilities. We matching lower bounds and upper bounds on the ability to implement such mediators using *cheap talk* (that is, just allowing communication among the players). The bounds depend on (a) the relationship between k, t and n, the total number of players in the system; (b) whether players know the exact utilities of other players; (c) whether there are broadcast channels or just point-to-point channels; (d) whether cryptography is available; and (e) whether the game has a $(k+t)$-*punishment strategy*; that is, a strategy that, if used by all but at most $k+t$ players, guarantees that every player gets a worse outcome than they do with the equilibrium strategy.

*Part of the work was done while the author visited Cornell university. The work was funded in part by ISF, ISOC, NSF, CCR, and AFOSR.

†Supported in part by NSF under grants CCR-0208535, ITR-0325453, and IIS-0534064, by ONR under grant N00014-01-10-511, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725, and by AFOSR under grant FA9550-05-1-0055.

The question of whether a problem in a multiagent system that can be solved with a trusted mediator can be solved by just the agents in the system, without the mediator, has attracted a great deal of attention in both computer science (particularly in the cryptography community) and game theory. In cryptography, the focus on the problem has been on *secure multiparty computation*. Here it is assumed that each agent i has some private information x_i. Fix functions $f_1, \ldots, f_n$. The goal is have agent i learn $f_i(x_1, \ldots, x_n)$ without learning anything about x_j for $j \neq i$ beyond what is revealed by the value of $f_i(x_1, \ldots, x_n)$. With a trusted mediator, this is trivial: each agent i just gives the mediator its private value x_i; the mediator then sends each agent i the value $f_i(x_1, \ldots, x_n)$. Work on multiparty computation [4; 7; 10] provides conditions under which this can be done. In game theory, the focus has been on whether an equilibrium in a game with a mediator can be implemented using what is called *cheap talk*—that is, just by players communicating among themselves (cf. [2; 3; 5; 8; 9]).

There is a great deal of overlap between the problems studied in computer science and game theory. But there are some significant differences. Perhaps the most significant difference is that, in the computer science literature, the interest has been in doing multiparty computation in the presence of possibly malicious adversaries, who do everything they can to subvert the computation. On the other hand, in the game theory literature, the assumption is that players have preference and seek to maximize their utility; thus, they will subvert the computation iff it is in their best interests to do so. In addition, the computer science literature has focused more on computational issues and resource-bounded players (although Urbano and Vila [8, 9] do consider implementation of mediators by resource-bounded players using cryptographic techniques) and on preserving secrecy of the inputs (although, for example, Heller [5] and Urbano and Vila [8, 9] also point out that their cheap talk protocols preserve secrecy of inputs).

In [1] (ADGH from now on), we argued that it is important to consider deviations by both rational players, who have preferences and try to maximize them, and players we can view as malicious, although it is perhaps better to think of them as rational players whose utilities are not known by the other players or mechanism designer. We considered equilibria that are (k, t)*-robust*; roughly speaking, this means that the equilibrium tolerates deviations by up to k rational players, whose utilities are presumed known, and up to t players with unknown utilities. We showed how (k, t)-robust equilibria with mediators could be implemented using cheap talk, by first showing how to implement secret sharing in a (k, t)-robust way using cheap talk. Our implementations depend on (a) the relationship between k, t and n, the total number of players in the system; (b) whether players know the exact utilities of other players; (c) whether there are broadcast channels or just point-to-point channels; (d) whether cryptography is available; and (e) whether the game has a $(k + t)$-*punishment strategy*; that is, a strategy that, if used by all but at most $k + t$ players, equilibrium strategy. We show that the possibility results proved by ADGH are in a sense optimal. The following is a high-level overview of our results:

- If $n > 3k + 3t$, then mediators can be implemented using cheap talk; no punishment strategy is required, no knowledge of other agents' utilities is required, and the cheap talk protocol has bounded running time that does not depend on the utilities.

- If $n \leq 3k + 3t$, then we cannot, in general, implement a mediator using cheap talk without knowledge of other agents' utilities . Moreover, even if other agents' utilities are known, we cannot, in general, implement a mediator without having a punishment strategy nor with bounded running time.

- If $n > 2k + 3t$, then mediators can be implemented using cheap talk if there is a punishment strategy (and utilities are known) in finite expected running time that does not depend on the utilities .

- If $n \leq 2k + 3t$, then we cannot, in general, implement a mediator, even if there is a punishment strategy and utilities are known.

- If $n > 2k + 2t$ and we can simulate broadcast then, for all ϵ, we can ϵ-implement a mediator (intuitively, there is an implementation where players get utility within ϵ of what they could get by deviating) using cheap talk, with bounded expected running time that that does not depend on the utilities in the game or ϵ.

- If $n \leq 2k + 2t$, we cannot, in general, ϵ-implement a mediator using cheap talk even if we have broadcast channels. Moreover, even if we assume cryptography, polynomially-bounded players, a public-key infrastructure (PKI), and broadcast channels, we cannot, in general, ϵ-implement a mediator using cheap talk with expected running time that does not depend on the utilities in the game or ϵ; if there is a punishment strategy, then we still cannot, in general, ϵ-implement a mediator using cheap talk with expected running time independent of the utilities in the game.

- If $n > k + 3t$ then, assuming cryptography and polynomially-bounded players, we can ϵ-implement a mediator using cheap talk; if n is also greater than $2k + 2t$ (in particular, if $t \geq k$), then the running time is independent of ϵ and the utilities in the game; if $k + 3t < n \leq 2k + 2t$ and there is a punishment strategy, then the running time depends on the utilities in the game, but not on ϵ; if $k+3t < n \leq 2k+2t$ and there is no punishment strategy, then the running time depends on both ϵ and the utilities.
- If $n \leq k+3t$, then even assuming cryptography and polynomially-bounded players, we cannot, in general, ϵ-implement a mediator using cheap talk.
- If $n > k+t$ then, assuming cryptography, polynomially-bounded players, and a PKI, we can ϵ-implement a mediator. Moreover, if there is a punishment strategy, the expected running time does not depend on ϵ.

The proofs of the impossibility results bring out deep connections between implementing mediators and various agreement problems, such as Byzantine agreement [6].

It is reasonable to ask at this point whether mediators are of practical interest. After all, if three companies negotiate, they can just hire an arguably trusted mediator, say an auditing firm such as PricewaterhouseCoopers. The disadvantage of this approach in a setting like the internet, with constantly shifting alliances, there are always different groups that want to collaborate; a group may not have the time and flexibility of hiring a mediator, even assuming they can find one they trust. Another concern is that our results simply shift the role of what has to be trust elsewhere. It is certainly true that our results assume point-to-point communication that cannot be intercepted. If $n \leq k + 3t$, then we must also assume the existence of a public-key infrastructure. Thus, we have essentially shifted from trusting the mediator to trusting the PKI. In practice, individuals who want to collaborate may find point-to-point communication and a PKI more trustworthy than an intermediary, and easier to work with.

References

[1] I. Abraham, D. Dolev, R. Gonen, and J.Y. Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *Proc. 25th ACM Symp. Principles of Distributed Computing*, pages 53–62, 2006.

[2] I. Barany. Fair distribution protocols or how the players replace fortune. *Mathematics of Operations Research*, 17:327–340, 1992.

[3] E. Ben-Porath. Cheap talk in games with incomplete information. *J. Economic Theory*, 108(1):45–71, 2003.

[4] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *Proc. 19th ACM Symp. Theory of Computing*, pages 218–229, 1987.

[5] Y. Heller. A minority-proof cheap-talk protocol. Unpublished manuscript, 2005.

[6] L. Lamport, R. Shostak, and M. Pease. The Byzantine Generals problem. *ACM Trans. on Programming Languages and Systems*, 4(3):382–401, 1982.

[7] A. Shamir, R. L. Rivest, and L. Adelman. Mental poker. In D. A. Klarner, editor, *The Mathematical Gardner*, pages 37–43. Prindle, Weber, and Schmidt, Boston, Mass., 1981.

[8] A. Urbano and J. E. Vila. Computational complexity and communication: Coordination in two-player games. *Econometrica*, 70(5):1893–1927, 2002.

[9] A. Urbano and J. E. Vila. Computationally restricted unmediated talk under incomplete information. *Economic Theory*, 23(2):283–320, 2004.

[10] A. Yao. Protocols for secure computation (extended abstract). In *Proc. 23rd IEEE Symp. Foundations of Computer Science*, pages 160–164, 1982.