Rational Secret Sharing, Revisited

[Extended Abstract]

S. Dov Gordon
gordon@cs.umd.edu

Jonathan Katz*
jkatz@cs.umd.edu

Dept. of Computer Science
University of Maryland
College Park, MD 20742

ABSTRACT

We consider the problem of secret sharing among n rational players. This problem was introduced by Halpern and Teague (STOC 2004), who claim that a solution is *impossible* for $n = 2$ but show a solution for the case $n \geq 3$. Counter to their claim, we show a simple protocol for the case of $n = 2$ players. Our protocol extends to the case $n \geq 3$, where it is both simpler than the Halpern-Teague solution and also offers a number of other advantages. We also show how to avoid the continual involvement of the dealer, in either our own protocol or that of Halpern-Teague.

Our techniques extend to the case of rational players trying to securely compute an arbitrary function, under certain assumptions on the utilities of the players.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Systems**]: General—*security and protection*; E.3 [**Data**]: Data Encryption

General Terms

Security

Keywords

Game theory, secret sharing, secure computation

1. INTRODUCTION

The classical problem of *t-out-of-n secret sharing* [10, 1] involves a "dealer" D who wishes to entrust a secret s to a group of n players $P_1, \ldots, P_n$ so that (1) any group of t or more players can reconstruct the secret without further intervention of the dealer, yet (2) any group of fewer than t players has no information about the secret. As an example, consider the scheme due to Shamir [10]: say the secret s lies in a field $\mathbb{F}$, with $|\mathbb{F}| > n$. The dealer chooses a random polynomial $f(x)$ of degree at most $t-1$ subject to the constraint $f(0) = s$, and gives the "share" $f(i)$ to player P_i (for $i = 1, \ldots, n$). Any set of t players can recover $f(x)$ (and hence s) by interpolation; furthermore, no set of fewer than t players has any information about s.

The implicit assumption above is that at least t players are willing to cooperate and pool their shares[1] when it is time to recover the secret; equivalently, at least t players are *honest* and hence at most $n - t$ players are *malicious*. Halpern and Teague [4] consider a scenario in which no one is (completely) honest, but instead all that is guaranteed is that at least t players are *rational* (as before, up to $n - t$ players may refuse to cooperate altogether). Shamir's protocol may no longer succeed in this scenario [4]. Specifically, if all players prefer to learn the secret above all else, and otherwise prefer that fewer parties learn the secret, then no player has any incentive to reveal their share. Consider P_1: if strictly fewer or greater than $t-1$ other players reveal their shares, nothing changes whether P_1 reveals his share or not. On the other hand, if *exactly* $t-1$ other players reveal their shares then P_1 learns the secret regardless (using his share), while P_1 can prevent other players from learning the secret by not publicly revealing his share. Thus, although for $t < n$ having all players reveal their shares is a Nash equilibrium, it is a weakly dominating strategy for each player *not* to reveal its share (and this is the equilibrium likely to be reached). Note that for $t = n$, having all players reveal their shares is not even a Nash equilibrium.

Does there exist *any* protocol for reconstructing the secret in which it is in players' best interests to follow the protocol? Generalizing the above, Halpern and Teague rule out any protocol terminating in a *fixed* number of rounds. This leaves open the possibility of *probabilistic* protocols without a fixed upper bound on their round complexity, and indeed Halpern and Teague show such a protocol for $n \geq 3$ parties. In contrast, they claim that a solution is *impossible* for $n = 2$ even if probabilistic protocols are allowed.

Our results: We revisit the question of rational secret sharing, in the model of [4]. As perhaps our most surprising

*Research supported by NSF Trusted Computing grants #0310499 and #0310751; NSF CAREER award #0447075; and US-Israel Binational Science Foundation grant #2004240.

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.

[1]We assume adversarial behavior is limited to refusal to cooperate. Reporting an incorrect share is easily prevented by having the dealer sign the shares.

result, we show a simple, probabilistic protocol for $n = 2$ parties to reconstruct a shared secret, thus disproving the claim of Halpern and Teague. Interestingly, their *proof* appears to be correct; the problem is that their *assumptions* about the types of protocols that might be used are too restrictive. By relaxing their assumptions in a reasonable way, we are able to circumvent their impossibility result.

Our protocol generalizes in a straightforward way to the case of $n \geq 3$ and arbitrary t. Although Halpern and Teague also claim a general solution of this sort, our solution is much simpler and has a number of other advantages.

In the Halpern-Teague solution, the dealer is involved periodically throughout the entire lifetime of the protocol. We also show how to remove this involvement of the dealer, in either our own protocol or that of Halpern and Teague.

Our results extend to the case of rational players computing an *arbitrary* function, under certain assumptions regarding their utilities. See the full version [3] for more details.

Related work: There has been much recent interest in bridging cryptography and game theory [6, 5, 7]. While prior work [6, 5] offers solutions to the problem considered here, we focus on simplicity and efficiency rather than generality. Our work also makes weaker physical assumptions than that of [6, 5]: specifically, we assume *simultaneous broadcast* (equivalently, we do not allow *rushing*) rather than "*secure envelopes*." See [3] for further discussion.

Recent and independent work [8] shows how to use essentially the same ideas shown here to obtain a stronger and more general result.

2. MODEL

We review the model of Halpern and Teague, filling in some details they omit. We have a dealer D holding a secret s, and n players. A protocol proceeds in a sequence of *iterations*, where each iteration consists of multiple *rounds*. At the beginning of each iteration, D distributes some information (privately) to the n players; at this point, any set of fewer than t players should have no information about s. The dealer is not involved during an iteration. Instead, some set of at least t players run the protocol amongst themselves by simultaneously broadcasting messages in a series of rounds. (Halpern-Teague additionally allow private communication between the players but we do not need this.) At the end of an iteration, the protocol either terminates or proceeds to the next iteration. We assume the dealer follows the protocol as specified. To rule out trivial protocols, we require that if at least t players follow the protocol in each iteration, the secret is eventually reconstructed.

Let $\vec{\sigma} = (\sigma_1, \ldots, \sigma_n)$ denote a vector of (possibly randomized) strategies used by the players. A *protocol* corresponds to the above *game* along with a prescribed strategy vector $\vec{\sigma}$. As in [4], we are interested in strategy vectors corresponding to a Nash equilibrium that survives iterated deletion of weakly-dominated strategies. See [9, 4] for definitions.

Let $\mu_i(\vec{\sigma})$ denote the utility of P_i for the strategy vector $\vec{\sigma}$. For a particular outcome o of the protocol, we let $\delta_i(o)$ be a bit denoting whether or not P_i learns the secret, and let $\mathsf{num}(o) = \sum_i \delta_i(o)$. We assume that for all i:

- $\delta_i(o) > \delta_i(o') \Rightarrow \mu_i(o) > \mu_i(o')$.
- If $\delta_i(o) = \delta_i(o')$, $\mathsf{num}(o) < \mathsf{num}(o') \Rightarrow \mu_i(o) > \mu_i(o')$.

That is, players first prefer outcomes in which they learn the secret; if this is held fixed, players prefer outcomes in which the fewest other players learn the secret. Let $U_i(\vec{\sigma})$ denote the expected value of the utility of P_i under strategy vector $\vec{\sigma}$. We assume rational players wish to maximize this value.

3. PROTOCOLS

We provide a high-level overview of the Halpern-Teague solution for 3-out-of-3 secret sharing. (Details of their proposed generalization for $n > 3, t \geq 3$ are in [4, 3].) At the beginning of each iteration, the dealer runs a fresh invocation of the Shamir scheme and sends the appropriate share to each player. During an iteration, each P_i flips a biased coin c_i with $\Pr[c_i = 1] = \alpha$. Players then securely compute $p = \bigoplus c_i$ such that it is impossible to cheat or to learn information about the $\{c_i\}$ values of the other parties. If $p = c_i = 1$, player P_i broadcasts his share. If all shares are revealed, the secret is reconstructed and the protocol ends. If $p = 1$ and no shares are revealed, all players terminate the protocol. In any other case, players proceed to the next iteration. Assuming players act honestly, the expected number of iterations until the protocol terminates is α^{-3}.

For a quick sketch as to why this works, assume P_1, P_2 follow the protocol and consider whether P_3 should deviate. One can show that there is no incentive for P_3 to change the distribution of c_3. If $p = 0$ or $c_3 = 0$, there is clearly no incentive for P_3 to deviate. When $p = c_3 = 1$, player P_3 does not know whether $c_1 = c_2 = 1$ (which occurs with probability $\frac{\alpha^2}{\alpha^2+(1-\alpha)^2}$) or $c_1 = c_2 = 0$ (which occurs with the remaining probability). If P_3 does not broadcast its share it runs the risk of having the protocol terminate without learning the secret. Setting α appropriately based on P_3's utility function, it is not in P_3's best interest to deviate.

3.1 Our Solution

Halpern and Teague implicitly assume that the dealer is restricted to sending valid Shamir shares to the players, and their impossibility proof for $n = 2$ therefore focuses only on what happens *during* an iteration. Removing this restriction circumvents the impossibility result for $n = 2$ (and also drastically simplifies things for the case of general n [3]).

The idea is as follows: with some probability the dealer shares the *actual* secret, and with the remaining probability the dealer shares a "bogus" secret. No player can tell which is the case given the share he receives. Then, players simply pool their shares and reconstruct the shared value. If this is the "actual" secret, the protocol terminates; otherwise, players continue to the next iteration. (We have to provide the players with a way to detect whether a reconstructed value is the actual secret or not; this is quite easy to do.)

We focus on the case $n = 2$ but it is easy to see that our idea generalizes to arbitrary t, n. Say the dealer holds a secret s that lies in a *strict subset* S of a field $\mathbb{F}$ (if s lies in a field $\mathbb{F}'$, this is achieves by taking a larger field $\mathbb{F}$ containing $\mathbb{F}'$ as a subfield). At the beginning of each iteration, with probability β the dealer generates a random sharing of s, and with probability $1-\beta$ the dealer generates a random sharing of an arbitrary element $\hat{s} \in \mathbb{F} \setminus S$. In a given iteration, the players simply broadcast their shares. If in any iteration some player does not broadcast their share, both players immediately terminate the protocol. Otherwise, both shares were broadcast and the players either reconstruct the secret

$s \in S$ and terminate the protocol successfully, or reconstruct a value $\hat{s} \in \mathbb{F} \setminus S$ and proceed to the next iteration.

To see why this works, note first that a player cannot tell from its share whether the dealer has shared the "real" secret s or the "bogus" secret $\hat{s}$. Assume P_1 acts honestly and consider whether P_2 has any incentive to deviate. The only possible deviation is for P_2 to refuse to broadcast his share. In this case, it learns the secret (while P_1 does not) with probability β, but with probability $1 - \beta$ it will never learn the secret. Setting β appropriately depending on P_2's utility, it is not in P_2's best interest to deviate.

The above shows that following the protocol is a Nash equilibrium. It is possible to additionally prove that it survives iterated deletion of weakly dominated strategies [3].

4. DISCUSSION AND EXTENSIONS

We view the main import of our result as a demonstration that rational secret sharing is, in fact, possible when $n = 2$; this serves as an illustration of the sensitivity of an impossibility result to the precise model under consideration. Our approach also has various other advantages as compared to the Halpern-Teague solution; chief among these may be its *simplicity*. See [3] for additional points of comparison.

In the full version [3], we show how to remove the need for the dealer to be involved at the beginning of each iteration so that, as in standard secret sharing, the dealer need only be involved *once*, at the beginning of the protocol. We also show how our results may be extended to the case of secure computation of arbitrary functions (à la [2]) by parties assumed only to be *rational* (i.e., without making the assumption that any parties are completely honest [2]).

5. REFERENCES

[1] G.R. Blakley. Safeguarding Cryptographic Keys. *National Computer Conference*, AFIPS Press, 1979.

[2] O. Goldreich, S. Micali, and A. Wigderson. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. STOC '87.

[3] S.D. Gordon and J. Katz. Rational Secret Sharing, Revisited. Available at http://eprint.iacr.org/2006/142.

[4] J. Halpern and V. Teague. Rational Secret Sharing and Multiparty Computation. STOC 2004.

[5] S. Izmalkov, S. Micali, and M. Lepinski. Rational Secure Function Evaluation and Ideal Mechanism Design. FOCS 2005.

[6] M. Lepinski, S. Micali, C, Peikert, and A. Shelat. Completely Fair SFE and Coalition-Safe Cheap Talk. PODC 2004.

[7] M. Lepinski, S. Micali, and A. Shelat. Collusion-Free Protocols. STOC 2005.

[8] A. Lysyanskaya and N. Triandopoulos. Rationality and Adversarial Behavior in Multi-Party Computation. Crypto 2006, to appear.

[9] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[10] A. Shamir. How to share a secret. *Comm. ACM*, 22(11): 612–613 (1979).