Scalable Rational Secret Sharing

[Extended Abstract]

Varsha Dani
Dept. of Computer Science
University of New Mexico
varshadani@gmail.com

Mahnush Movahedi
Dept. of Computer Science
University of New Mexico
movahedi@cs.unm.edu

Yamel Rodriguez
Dept. of Computer Science
University of New Mexico
yamel@cs.unm.edu

Jared Saia[*]
Dept. of Computer Science
University of New Mexico
saia@cs.unm.edu

ABSTRACT

We consider the classical secret sharing problem in the case where all agents are selfish but rational. In recent work, Kol and Naor show that in the non-simultaneous communciation model (i.e. when rushing is possible), there is no Nash equilibrium that ensures all agents learn the secret. However, they describe a mechanism for this problem that is an ϵ-*Nash equilibrium*, i.e. it is close to an equilibrium in the sense that no player can gain more than ϵ utility by deviating from it.

Unfortunately, the Kol and Naor mechanism, and, to the best of our knowledge, all previous mechanisms for this problem require each agent to send $O(n)$ messages in expectation, where n is the number of agents. This may be problematic for some applications of rational secret sharing such as secure multiparty computation and simulation of a mediator.

We address this issue by describing a mechanism for rational n-out-of-n secret sharing that is an ϵ-Nash equilibrium, and is *scalable* in the sense that it requires each agent to send only an expected $O(1)$ bits. Moreover, the latency of our mechanism is $O(\log n)$ in expectation, compared to $O(n)$ expected latency for the Kol and Naor result. We also design mechanisms for a relaxed variant of rational m-out-of-n secret sharing where $m = \Theta(n)$ that require each processor to send $O(\log n)$ bits and have $O(\log n)$ latency. Our mechanisms are non-cryptographic, and are not susceptible to backwards induction.

> *"Three can keep a secret if two of them are dead."*
> - Benjamin Franklin

Categories and Subject Descriptors

F.m [**Theory of Computation**]: Miscellaneous

General Terms

Theory

Keywords

Game Theory, Cryptography, Secret Sharing, Multiparty Computation

1. INTRODUCTION

Secret sharing is one of the most fundamental problems in security, and is an important primitive in many cryptographic protocols, including secure multiparty computation. Recently, there has been interest in solving *rational secret sharing* [8, 4, 5, 1, 11]. In this setting, there are n selfish but rational agents, and we want to distribute shares of a secret to each agent, and design a protocol for the agents ensures that: (1) if any group of m agents follow the protocol they will all learn the secret; and (2) knowledge of less than m of the shares reveals nothing about the secret. Moreover, we want our protocol to be a *Nash equilibrium* in the sense that no player can improve their utility by deviating from the protocol, given that all other players are following the protocol.

Unfortunately, all previous solutions to this problem require each agent to send $O(n)$ messages in expectation, and so do not scale to large networks. Rational secret sharing is a primitive for rational multiparty computation, which can be used to compute an arbitrary function in a completely decentralized manner, without a trusted external party. A typical application of rational multiparty computation might be to either run an auction, or to hold a lottery to assign resources in a network. It is easy to imagine such applications where the number of players is large, and where it is important to have algorithms whose bandwidth and latency costs scale well with the number of players. Moreover, in a game theoretic setting, standard tricks to circumvent scalability issues, like running the protocol only on a small subset of the players, may be undesirable since they could lead to increased likelihood of bribery attacks.

In this paper, we address this issue by designing scalable

mechanisms for rational secret sharing. Our main result is a protocol for rational n-out-of-n secret sharing that (1) requires each agent to send only an expected $O(1)$ bits; and (2) has $O(\log n)$ expected latency. We also design scalable mechanisms for a relaxed variant of m-out-of-n rational secret sharing in the case where m is $\Theta(n)$.

1.1 The Problem

We assume there are n rational but selfish players. The players' utility functions are such that they prefer to learn the secret, but also prefer that other players not learn the secret. Following previous work [8, 4, 5, 11], we assume all the players have the same utility function, which is specified by three constants U_+, U, and U_-. Here U_+ is the utility to a player if he alone learns the secret, U is the utility if he learns the secret but at least one other player learns it as well, and finally U_- is the utility when the player does not learn the secret. We further assume that $U_+ > U > U_-$, so that the players' preferences are strict. We note that, as in previous work, the utility function does not distinguish the cases where different numbers of other players may learn the secret.

We will assume that the secret is an arbitrary element of a large (fixed) finite set $\mathcal{S}$. The shares are provided to the players by a dealer, who is active only once, at the beginning of the game. The players must then communicate with each other in order to reconstruct the secret.

We assume that the communication between the players is *point-to-point* and through private channels. In other words, if player A sends a message to player B, then a third player C is not privy to the message that was sent, or indeed even to the fact of a message having been sent. We assume communication is synchronous in that there is an upperbound known on the maximum amount of time required to send a message from one player to another. However, we assume *non-simultaneous* communication, and thus allow for the possibility of *rushing*, where a player may receive messages from other players in a round before sending out its own messages.

Our goal is to provide protocols for the dealer and rational players such that the players following the protocol can reconstruct the secret. Moreover, we want a protocol that is *scalable* in the sense that the amount of communication and the latency of the protocol should be a slow growing function of the number of players.

We are also concerned with the problem of rational multiparty computation, which is an important application of n-out-of-n rational secret sharing. In the problem of rational multiparty computation, there are n rational agents, and each agent i has an input x_i to a function f of n variables. Each player prefers to learn the output of f, but also prefers that other players do not. In particular, the utility functions are the same as for the secret sharing problem. The goal is to design a protocol for rational players that will ensure that all players learn the output of f. As described in [5], the algorithm from [3] can be used to solve the problem of rational multiparty computation provided one has a solution to n-out-of-n rational secret sharing.

1.2 Our Results

In their recent work, Kol and Naor [8] show that in the non-simultaneous broadcast model (i.e. when rushing is possible), there is no Nash equilibrium that ensures all agents

learn the secret. They thus consider the case of designing an ϵ-*Nash equilibrium* for the problem in this communication model, where an ϵ-Nash equilibrium is close to an equilibrium in the sense that no player can gain more than ϵ utility by deviating from it. In addition, their protocol is Monte Carlo, succeeding with probability $1-\delta$ for any fixed positive δ.

Our main result is a scalable, Monte Carlo protocol for n-out-of-n rational secret sharing that is also an ϵ-Nash equilibrium. This result is summarized in the following theorem.

THEOREM 1.1. *For any fixed positive ϵ, δ there exists a protocol for rational n-out-of-n secret sharing that with probability $1 - \delta$ has the following properties:*

- *The protocol is an ϵ-Nash equilibrium*

- *All players learn the secret*

- *The protocol, in expectation, requires each player to send $O(1)$ bits, and has latency $O(\log n)$*

As discussed above, the n-out-of-n case for rational secret sharing is a critical component of rational multiparty computation. Our result for n-out-of-n rational secret sharing enables a protocol for rational multiparty computation that is an ϵ-Nash equilibrium in the point-to-point, synchronous, non-simultaneous communication model. Moreover, it reduces worst case bandwidth by a multiplicative factor of n, and latency by a multiplicative factor of $\Theta(n/\log n)$ over the rational multiparty protocol in this communication model from [8].

In this paper, we also consider the problem of m-out-of-n rational secret sharing for the case where $m < n$. Designing scalable algorithms for this problem is challenging because of the tension between reduced communication, and the need to ensure that *any* active set of m players can reconstruct the secret. For example, consider the case where each player sends $O(\log n)$ messages. If $m = o(n/\log n)$, even if the set of active players is chosen *randomly*, it is likely that there will be some active player that will never receive a message from any other active player. Moreover, even if $m = \Theta(n)$, if the set of active players is chosen in a worst case manner, it is easy to see that a small subset of the active players can easily be isolated so that they never receive messages from the other active players, and are thus unable to reconstruct the secret.

Despite the difficulty of the problem, scalable rational secret sharing for the m-out-of-n case may still be of interest for applications like the Vanish peer-to-peer system [2]. To determine what might at least be possible, we consider a significantly relaxed variant of the problem. In particular, we require $m = \Theta(n)$ and that the set of m active players be chosen independently of the random bits of the dealer. Our result is given in the following theorem, whose sketch is given in Section 4.

THEOREM 1.2. *For any fixed positive $\epsilon, \delta, , \lambda$, and threshold τ, there exists a protocol for rational secret sharing with absent players, which with probability $1-\delta$ has the following properties, provided that the subset of m active players is chosen independently of the random bits of the dealer:*

- *The protocol is an ϵ-Nash equilibrium.*

- *The protocol ensures that if at least a $(\tau + \lambda)$ fraction of the players are active, (i.e. $m/n \geq \tau + \lambda$) then all active players will learn the secret; and if less than a $(\tau - \lambda)$ fraction of the players are active, (i.e. $m/n \leq \tau - \lambda$) then the secret can not be recovered*

- *The protocol requires each player to send $O(\log n)$ bits, and has latency $O(\log n)$*

1.3 Our Approach

Our protocol shares many similarities with the approach of Kol and Naor in [8]. Following are the techniques we use from their protocol. First, each player i receives a set of lists from the dealer, and the j-th element of each list is used in round j. Second, we make use of two lists from [8]: (1) the list L_i that contains potential secrets, for some value t^*, the t^*-th element in each of these lists is the true secret; and (2) the list S_i that contains shares of an indicator sequence, for each round j, that reveals whether or not the j-th round is the t^*-th round. Finally, we make use of a clever technique from [8] to avoid backwards induction. In particular, one player, chosen uniformly at random by the dealer, is designated the *short* player and the remaining players are designated the *long* players. The length of the short player's lists is determined by geometric distributions, and the length of the long players' lists exceeds the length of the short player's list by an amount that is geometrically distributed. This scheme protects against backwards induction by making it difficult for any player to know exactly how many rounds the protocol will last.[1]

Following are the novel techniques in our approach. First, we restrict communication in our protocol to a certain type of complete binary tree; this is critical for ensuring that the total number of messages sent by each player is scalable. Second, we make use of an iterated secret sharing scheme over this tree to divide up shares of secrets among the players. This scheme is similar to that used in recent work by King and Saia [7] on the problem of scalable Byzantine agreement, and suggests deeper connections between the two problems. Finally, we make use of an iterated tag and hash scheme [13, 12] to enable checking of secrets and iterated shares of secrets. Tag and hash schemes were used previously in [8], but applying them in our setting, where communication is severely curtailed, is technically challenging.

1.4 Related Work

Since its introduction by Halpern and Teague in [5], there has been significant work on the problem of rational secret sharing, including results of Halpern and Teague [5], Gordon and Katz [4], Abraham et al. [1], Lysyanskaya and Triandopoulos [11] and Kol and Naor [8]. All of this related work except for [8], assumes the existence of simultaneous communication, either by broadcast or private channels. Several of the protocols proposed [4, 1, 11] make use of cryptographic assumptions and achieve equilibria under the assumption that the players are computationally bounded. The protocol from [1] is robust to coalitions; and the protocol from [11]

works in the situation where players may be either rational or adversarial.

The work of Kol and Naor [8] is closest to our own work in that they do not assume simultaneous communication, and do not make cryptographic assumptions. As we have already discussed, our protocols make use of several clever ideas from their result.

Additional work by Lepinski et al. [10, 9] and Izmalkov et al. [6] describe protocols for rational secure multiparty computation that are fair and robust to coalitions. However, their results rely on the physical assumption of "secure envelopes" [10], which seem difficult to implement for participants that are physically distant.

2. RATIONAL SECRET SHARING WHEN ALL PLAYERS ARE PRESENT

We now describe our mechanism for n-out-of-n secret sharing. We give an informal description of the dealer's and players' protocols in Sections 2.1 and 2.2. The formal descriptions of these protocols appears in Algorithm 1 and Algorithm 4.

2.1 Dealer's Protocol

The dealer is active only once at the beginning of the game, and during this phase of the game the players' inputs are prepared.

Communication between the players in our protocol will be restricted to sending messages to their neighbors in a complete binary tree.[2] First the players are assigned to nodes of a complete binary tree with n leaves. Each player is assigned to one leaf. Next the layer of internal nodes just above the leaves is assigned to players. This is done in such a way as to ensure that the player assigned to the parent node of two leaves is one of the two players assigned to those leaves. Then the remaining internal nodes are assigned arbitrarily to players who have not yet been assigned an internal node. Since there are $n - 1$ internal nodes in a tree with n leaf nodes, we can ensure that no player is assigned to more than one internal node. One of the players will not be assigned to an internal node, and this will not matter.

Next, the dealer will select a player uniformly at random to be the "short player". The rest of the players will be "long players". It is important that the players *not be told* which player is the short player. This is necessary to avoid backwards induction as is discussed in Section 3.2.

Now the dealer independently samples three random variables X, Y and Z from a geometric distribution with parameter β. X will be the definitive iteration, or the round of the game in which the true secret is revealed. Y will be the amount of padding on the short player's input. Z will be the amount of additional padding on the long players' input. Thus, the short player will receive enough input to last for $X + Y$ rounds of the game, while the long players will receive enough input to last for $X + Y + Z$ rounds of the game.

The players' input will consist of several lists, with one list element for each round of the game. Thus the lists sent

---

[1]It is also what makes their protocol and ours Monte Carlo, since with some small but constant probability, all the geometric random variables may be so small that it is better for each player to randomly try to guess the secret from the elements in their L_i list than to follow the protocol.

[2]Recall that a complete binary tree is a binary tree in which all the internal nodes have two descendants; all the leaves are at the two deepest levels; and the leaves on the deepest level are as far left as possible.

Algorithm 1 *Dealer's Protocol*

n players, $\mathcal{S}$: set of potential secrets, $\mathbb{F}_q$: field of size q. $\beta \in (0,1)$: geometric distribution parameter.

1. Create a complete binary tree with n leaf nodes. Assign players to leaf nodes from left to right. Assign players to parents of leaf nodes by choosing the player assigned to one of the children of the node. Assign all other internal nodes to players who have not yet been assigned to an internal node, top down and left to right. Thus each player is assigned to a distinct leaf node, and each player except one, is assigned to a distinct internal node.

2. Choose X, Y, and Z independently from a geometric distribution with parameter β.

3. Create the following lists for each player i:

   - List L_i for potential secrets
   - List S_i for shares of the indicator sequence.
   - For each node w with which i is associated, list T_i^w for tags.
   - For each node w with which i is associated and each neighbor w' of w, list $H_i^{(w,w')}$ for verification functions.

4. (Populate the lists) For each round $t \le X + Y + Z$:

   (a) if $t = X$, $\sigma \leftarrow$ true secret.
   else $\sigma \leftarrow$ random element from $\mathcal{S}$
   Add σ to all the lists L_i at position t.

   (b) if $t = X$, $h_t \leftarrow 0$
   else $h_t \leftarrow$ random non-zero element from $\mathbb{F}_q$.

   (c) For each node w of the tree, create a 'share' v_t^w (to be set by *RecursiveShares*, subroutine 2)

   (d) Call *RecursiveShares*(root, h_t)

   (e) For each leaf w add v_t^w to the shares list S_i of corresponding player i at position t.

   (f) For each node w in the binary tree choose $(g_t^w, \bar{g}_t^w)$ uniformly at random from $\mathbb{F}_q^2$ and add it to T^w at position t

   (g) For node each w in the binary tree:

   i. if w has a parent $p(w)$:
   $V_p \leftarrow VerificationFunction(h_t, g_t^{p(w)})$
   Add V_p to list $H^{w,p(w)}$ at position t.

   ii. if w has a left child $\ell(w)$:
   $V_\ell \leftarrow VerificationFunction(v_t^{\ell(w)}, \bar{g}_t^{\ell(w)})$
   Add V_ℓ to list $H^{w,\ell(w)}$ at position t.

   iii. if w has a right child $r(w)$:
   $V_r \leftarrow VerificationFunction(v_t^{r(w)}, \bar{g}_t^{r(w)})$
   Add V_r to list $H^{w,r(w)}$ at position t.

5. Choose a player i uniformly at random to be the short player. Truncate each list (L_i S_i, and multiple tag and hash lists) associated with player i to be length $X + Y$ *i.e.* delete the last Z elements of each such list. (All other players' lists remain of length $X + Y + Z$.)

6. Send the data to the appropriate players

Algorithm 2 *RecursiveShares* (node w , integer y):

1. $v_t^w \leftarrow y$.

2. If w has children $\ell(w)$ and $r(w)$:

   (a) Choose random slope μ.

   (b) Let f be the line with slope μ and y-intercept y.

   (c) *RecursiveShares*($\ell(w), f(-1)$).

   (d) *RecursiveShares*($r(w), f(1)$).

Algorithm 3 *VerificationFunction* ($\mathbb{F}_q$-element v, $\mathbb{F}_q$-element g): // v is the message, g is the tag

1. Choose a and b uniformly at random $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$

2. $c = a * v + b * g$.

3. return (a, b, c).

Algorithm 4 *Player's Protocol* for Player i

On round t:

Up-Stage:

1. Send $(S_i[t], \bar{g}_t^i)$ to parent in the tree.

2. If player i is also internal node w, then:

   (a) Receive $(v_t^{\ell(w)}, \bar{g}_t^{\ell(w)})$ from left child and check using verification function $H_t^{w,\ell(w)}$. If no message is received or message does not satisfy verification function, output $L_i[t]$ and QUIT. Otherwise set $f(-1) \leftarrow v_t^{\ell(w)}$.

   (b) Receive $(v_t^{r(w)}, \bar{g}_t^{r(w)})$ from left child and check using verification function $H_t^{w,r(w)}$. If no message is received or message does not satisfy verification function, output $L_i[t]$ and QUIT. Otherwise set $f(1) \leftarrow v_t^{r(w)}$.

   (c) Reconstruct a degree 1 polynomial f from $(-1, f(-1))$ and $(1, f(1))$.

   (d) if w is not the root, $v_t^w \leftarrow f(0)$
   Send $(v_t^w, \bar{g}_t^w)$ to parent in the tree.

   (e) Else if w is the root $h_t \leftarrow f(0)$.

Down-Stage:

1. If player i is the root: send (h_t, g^{root_t}) to left and right children.

2. Else if i is node w (including leaves):

   (a) Receive $(h_t, g_t^{p(w)})$ from parent $p(w)$. If no message received output $L_i[t]$ and QUIT.

   (b) Use verification function $H^{(w,p(w))}[t]$ to check validity of message. If message does not satisfy verification function, output $L_i[t]$ and QUIT.

   (c) If w is an internal node, then send indicator (h_t, g_t^w) to left and right children.

   (d) If $h_t = 0$, output $L_i[t]$ and QUIT.

3. $t \leftarrow t + 1$

to the short player will be of length $X + Y$ and those sent to long players will be of length $X + Y + Z$. Each player i will receive the following four lists:

- A list L_i consisting of elements of $\mathcal{S}$, or potential secrets. The true secret will be at index X of the list.

- A list S_i containing the shares of the indicator sequence for player i (in his role as a leaf of the tree.)

- For each node w (including leaves) with which i is associated, player i will receive a list T_i^w of pairs of tags (one for the *Up-Stage*, one for the *Down-Stage*) which will be used in the tag and hash scheme described below.

- For each node w with which i is associated, for each neighbor w' of w, player i will receive a list $H_i^{w,w'}$ with verification functions to check the authenticity of messages from w'.

The dealer's next task is to populate these lists. For simplicity of exposition, we assume that all the lists are created to be of length $X + Y + Z$, and that the short player's lists are truncated to length $X + Y$ later, before being sent to him. The dealer will fill the potential secret-lists L_i with elements chosen independently and uniformly at random from the set $\mathcal{S}$. These lists have the property that the true secret is in position X of the list.

Let q be a (fixed) large prime. The indicator sequence and tags will all be elements of the field $\mathbb{F}_q$ and all of the algebraic operations described hereafter will be over $\mathbb{F}_q$.

Next, to populate the shares-lists S_i, for each round $t \neq X$, the dealer will set h_t to be a uniformly random non-zero element of $\mathbb{F}_q$. h_X will be set to 0. Then for each t the dealer will create *iterated shares* of h_t as follows. First, the value h_t is assigned to the root of the tree. Next, for each node w of the tree for which a value v_t^w has been assigned, if w is a leaf, assign the value v_t^w to be the share of the corresponding player for round t. Otherwise, pick a random slope $\mu \neq 0$ and let f be the line with slope μ and y-intercept v_t^w (so that $f(0) = v_t^w$). Assign values $f(-1)$ and $f(1)$ to the left and right children of w, and recurse. The formal algorithm for this process is given in Algorithm 2, and an example run of this algorithm is illustrated in Figure 1.

The dealer's remaining task is to populate lists that implement the tag and hash verification scheme that will be used by the players. For each round t and each internal node w of the tree, the dealer generates tags g_t^w and $\bar{g}_t^w$ which are elements chosen independently and uniformly at random from $\mathbb{F}_q$. These elements are put into the appropriate lists T^w. The tag g_t^w will be used on the *Down-Stage* of the protocol for the players that is described in Section 2.2, and $\bar{g}_t^w$ will be used in the *Up-Stage* of this protocol.

Now, for each round t and each node w, let $p(w)$ denote the parent of w, and let $\ell(w)$ and $r(w)$ be the left and right children of w if any. The dealer samples uniformly random non-zero elements a_t^w and b_t^w from $\mathbb{F}_q$ and sets

$$c_t^w = a_t^w * h_t + b_t^w * g_t^{p(w)}$$

The triple (a_t^w, b_t^w, c_t^w) consists of the verification function for the *Down-Stage* round of t for the node w. Similarly, the dealer generates such verification triples (a, b, c) to check the validity of the messages $(v_t^{\ell(w)}, \bar{g}_t^{\ell(w)})$ and $(v^{r(w)}, \bar{g}_t^{r(w)})$ that
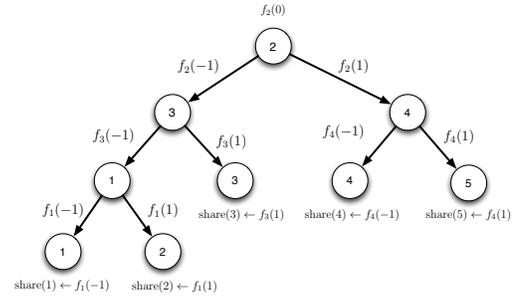

Figure 1: Example run of the dealer's protocol for a fixed round t. The function f_i is defined as $f_i(x) = a_i x + b_i$ where a_i and b_i are elements chosen independently and uniformly at random from the field $\mathbb{F}_q$. When $t = t^*$, b_2 is fixed at 0 and a_i is chosen uniformly at random; this ensures that $f_2(0)$ will be 0. When $t \neq t^*$, b_2 is chosen uniformly at random from $\mathbb{F}_q - 0$, which ensures that $f_2(0) \neq 0$.

are to be received by w in the *Up-Stage*. These verification function triples are put into the appropriate lists H^w. (See Algorithm 3.) Note that for any triple $(a, b, c) \in \mathbb{F}_q^3$ with $a, b \neq 0$ there are q pairs $(x, y) \in \mathbb{F}_q^2$ which satisfy $ax + by = c$, so the values of the message and tag cannot be guessed from the verification function, except with probability $1/q$.

The dealer's task is now almost done. All that remains is to truncate the short player's lists and communicate the input to the players. The formal description of the entire protocol for the dealer is given in Algorithm 1.

2.2 Player's Protocol

The players' protocol consists of two stages: the *Up-Stage* and the *Down-Stage*. In the *Up-Stage* for round t each player i will send his share $S_i[t]$ along with the tag $\bar{g}_t^i$ up the tree to the parent node of the leaf with which player i is associated. Each internal node w receives two shares and tags from its two children and after checking their validity with the provided verification function $(H_t^{w,\ell(w)})$, uses these to reconstruct a line f, where the non-tag part of the messages received from the left and right children respectively represent $f(-1)$ and $f(1)$. Once f is reconstructed, w will send his "share", $f(0)$ along with his tag $\bar{g}_t^w$ to his parent, up the tree, who will recursively do the same thing. This procedure continues recursively up the tree until the root is reached. At the root, the function f, reconstructed from the shares of the two children, when evaluated at 0 will yield the indicator for the current round. In particular, the root sets $h_t = f_{\text{root}}(0)$. We next move into the *Down-Stage*.

In the *Down-Stage* for round t the root will send h_t (determined by the end of the *Up-Stage*) and its tag g_t^{root} to both of its children. Now each child node w, receiving from its parent $p(w)$ the pair $(h_t, g_t^{p(w)})$, will first check this message for consistency using its verification function (a_t^w, b_t^w, c_t^w). Once the message is verified, w, if it is itself an internal node, will recursively append its own tag to h_t and send the pair (h_t, g_t^w) to its children. If w is a leaf, then if $h_t = 0$ then $t = X$ and the player associated with w will output the secret for the current round as the true secret, and quit the protocol. Otherwise the next round begins.
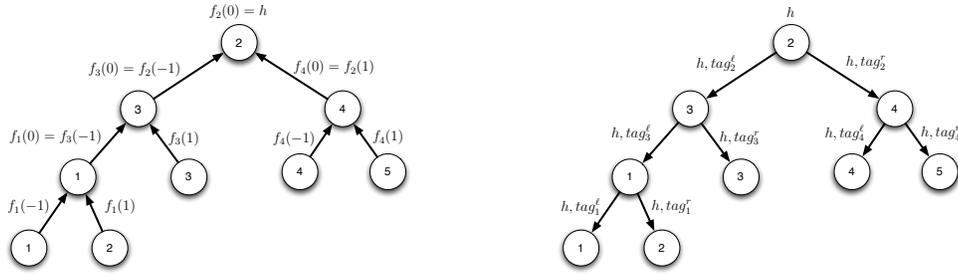
Figure 2: On the left is the up stage of the algorithm; on the right is the down stage.

If at any time in the protocol, a player detects deviation from the protocol by another player, then the player immediately outputs the potential secret for the current round as the true secret and quits the protocol. This detection of deviation could be in the form of no message being received when a message was supposed to be sent, or in the form of messages received that do not satisfy the verification function. Note that detection of cheating propagates to all the players, since once a player detecting cheating quits the protocol, everyone who was expecting messages from him will also fail to get these messages. The formal description of the protocol for the players is given in Algorithm 4, and an example run of this algorithm is illustrated in Figure 2.

3. ANALYSIS OF OUR PROTOCOL

There are two ways in which a player may try to deviate from the protocol. First, he may try to remain in the protocol and send messages not specified by the protocol, or "fake" messages, in order to confuse other players. Second, he may simply output a value for the secret and leave the protocol early. As the following proposition shows, the tag-and-hash verification scheme used in the protocol makes it hard for a player to successfully fool other players by sending fake messages.

PROPOSITION 3.1. *The probability that a faked message will satisfy the verification function is* $\frac{1}{q-1}$

PROOF. Suppose a player A is to send a message v with tag g to player B who has a verification triple (a, b, c) such that $av + bg = c$. Suppose further that A wants to send a fake message instead. Note that A has no information about the verification function except that the message he is supposed to send satisfies it. In fact, recall that the verification function was created by the dealer by randomly selecting a and b and setting c to the appropriate value. Consider the following two cases.
Case 1: A wants to send a particular fake message v'. In this case, there is a unique $g' = (c - av')/b$ such that the pair (v', g') satisfies the verification function. With no knowledge of a or b, A's chance of correctly guessing g' is $\frac{1}{q-1}$.
Case 2: A just wants to send a fake message, but does not care what it is. In this case, there are exactly q pairs (v', g') which satisfy the verification function; one of these is (v, g). Thus only $q - 1$ fake pairs work. On the other hand, there are $(q-1)^2$ pairs of elements in $\mathbb{F}_q^2$ which are not of the form $(v, *)$ or $(*, g)$. So again, with no information about a or b, A's chance of guessing a pair that works is $\frac{1}{q-1}$. $\square$

We will next consider the second type of deviation from the protocol, where a player deviates by leaving the protocol early. A player may consider deviating in this way if he either knows the secret with certainty, or with sufficiently high probability. The former only happens partway through the definitive iteration, when some of the players have learned this fact, but have not yet communicated it to the others. However, if a player decides not to transmit the fact that $h_t = 0$ to players lower in the tree, he does not gain anything by doing so. This is because the players who do not get the messages they are expecting will, in accordance with the protocol, output the secret for the current round, which is the correct secret. Thus, the only situation we still need to consider is whether a player can guess the position of the secret, with sufficiently high probability, before the definitive iteration. We deal with this situation in the next few lemmas.

LEMMA 3.2. *A player deviating from the protocol cannot increase his expected payoff by more than ϵ unless his probability of successfully learning the secret by deviating is at least* $\frac{U - U_- + \epsilon}{U_+ - U_-}$

PROOF. Let p be the threshold success probability required to make deviating worthwhile, ie to increase one's payoff by ϵ. Since the game ends when a player is caught deviating from the protocol (and certainly if the player stops sending messages, this *will* be detected by the other players), a failed attempt at cheating means that the player does not learn the secret. Thus a player's expected utility from cheating while everyone else follows the protocol is $pU_+ + (1-p)U_-$. On the other hand, following the protocol results in everyone learning the secret, with a utility of U. Thus, at the threshold for making cheating worthwhile, we have

$$pU_+ + (1-p)U_- = U + \epsilon$$

Solving this for p gives the desired result. $\square$

We will now compute the probability of a player successfully guessing the secret before it is actually revealed during the protocol.

LEMMA 3.3. *A player who initially received a list of length α has at most $\frac{3}{\alpha - t}$ chance of (correctly) guessing the position of the secret on round t if it has not already been revealed.*

PROOF. For $j \geq t$ let $P(\alpha, t, j)$ denote the probability that the secret is in the jth position in the list of potential secrets, given that the player initially received a list of length α, and $t - 1$ rounds of the protocol have ended without the secret being revealed.

Let $|L_i|$ be the length of the list of potential secrets received by player i. Recall that X, Y and Z are geometric random variables with parameter β, corresponding to the position of the secret, the amount of padding on the short player's list, and the amount of additional padding on the long players' lists respectively. Then $|L_i| = X + Y$ if i is the short player, and $|L_i| = X + Y + Z$ if i is a long player. Let ξ be the event that the player was chosen as the short player; its complement $\bar{\xi}$ is the event that the player was chosen as a long player. Now

$$P(\alpha, t, j) = \Pr(X = j | X \geq t, |L_i| = \alpha)$$

By the law of total probability, this equals

$$\Pr(X = j | X \geq t, |L_i| = \alpha, \xi) \Pr(\xi | X \geq t, |L_i| = \alpha)$$
$$+ \Pr(X = j | X \geq t, |L_i| = \alpha, \bar{\xi}) \Pr(\bar{\xi} | X \geq t, |L_i| = \alpha)$$

Bounding both conditional probabilities by 1 and observing that the event $|L_i| = \alpha \wedge \xi$ is the same as $X + Y = \alpha$ while the event $|L_i| = \alpha \wedge \bar{\xi}$ is the same as $X + Y + Z$, we see that

$$P(\alpha, t, j) \leq \Pr(X = j | X \geq t, X + Y = \alpha)$$
$$+ \Pr(X = j | X \geq t, X + Y + Z = \alpha)$$

so it suffices to bound these two conditional probabilities. Let $P_{short}(\alpha, t, j)$ and $P_{long}(\alpha, t, j)$ respectively be the two summands on the right side above; we will now compute these.

$$\begin{aligned}
P_{short}(\alpha, t, j) &= \Pr(X = j | X \geq t, X + Y = \alpha) \\
&= \frac{\Pr(X = j, X \geq t, X + Y = \alpha)}{\Pr(X \geq t, X + Y = \alpha)} \\
&= \frac{\Pr(X = j, Y = \alpha - j)}{\sum_{x=t}^{\alpha-1} \Pr(X = x, Y = \alpha - x)} \\
&= \frac{\Pr(X = j) \Pr(Y = \alpha - j)}{\sum_{x=t}^{\alpha-1} \Pr(X = x) \Pr(Y = \alpha - x)} \\
&= \frac{(1-\beta)^{j-1}\beta(1-\beta)^{\alpha-j-1}\beta}{\sum_{x=t}^{\alpha-1}(1-\beta)^{x-1}\beta(1-\beta)^{\alpha-x-1}\beta} \\
&= \frac{(1-\beta)^{\alpha-2}\beta^2}{(\alpha-t)(1-\beta)^{\alpha-2}\beta^2} \\
&= \frac{1}{\alpha-t}
\end{aligned}$$

Also

$$\begin{aligned}
P_{long}(\alpha, t, j) &= \Pr(X = j | X \geq t, X + Y + Z = \alpha) \\
&= \frac{\Pr(X = j, X \geq t, X + Y + Z = \alpha)}{\Pr(X \geq t, X + Y + Z = \alpha)} \\
&= \frac{\Pr(X = j, Y + Z = \alpha - j)}{\sum_{x=t}^{\alpha-2} \Pr(X = x, Y + Z = \alpha - x)} \\
&= \frac{\Pr(X = j) \Pr(Y + Z = \alpha - j)}{\sum_{k=1}^{\alpha-t-1} \Pr(X = \alpha - k - 1, Y + Z = k + 1)}
\end{aligned}$$

where the last equality is via the change of variables $k = \alpha - x - 1$. Since $\Pr(Y + Z = a) = (a-1)(1-\beta)^{a-2}\beta^2$ we

have

$$\begin{aligned}
P_{long}(\alpha, t, j) &= \frac{(1-\beta)^{j-1}\beta(\alpha-j-1)(1-\beta)^{\alpha-j-2}\beta^2}{\sum_{k=1}^{\alpha-t-1}(1-\beta)^{\alpha-k-2}\beta k(1-\beta)^{k-1}\beta^2} \\
&= \frac{(\alpha-j-1)}{\sum_{k=1}^{\alpha-t-1} k} \\
&= \frac{2(\alpha-j-1)}{(\alpha-t)(\alpha-t-1)}
\end{aligned}$$

Since $j \geq t$, $(\alpha - j - 1)/(\alpha - t - 1) \leq 1$ and from above $P_{long}(\alpha, t, j) \leq \frac{2}{\alpha-t}$. It follows that

$$P(\alpha, t, j) \leq \frac{3}{\alpha - t}$$

as asserted. $\quad\square$

We make the observation that if player i received a list of length α at the beginning, and $t - 1$ rounds of the protocol did not reveal the secret, then if player i guesses that the true secret is in position j, $P(\alpha, t, j)$ represents the probability that this guess is correct. For $j \geq t + 1$ this also represents the probability that player can succesfully cheat by outputting the potential secret at position j and dropping out of the protocol. The situation is slightly different for $j = t$ because if the player is caught cheating on round t, all other players following the protocol will output the secret at position j so if that is correct, player i will gain no advantage from having guessed it. Thus, on round t the best probability of a successful cheat is obtained by guessing that the secret is in position $t + 1$. As we have just seen, the probability of success is at most $3/(\alpha - t)$.

3.1 Proof of Theorem 1.1

We are now ready to present the proof of Theorem 1.1

PROOF. By Proposition 3.1 and Lemma 3.2, if $q > 1 + \frac{U_+ - U_-}{(U - U_- + \epsilon)}$ then no player can gain ϵ or more by faking messages in the protocol.

By Lemma 3.3 and Lemma 3.2, to prevent cheating by guessing the secret and prematurely exiting the protocol, it suffices if

$$\frac{3}{\alpha - t} < \frac{U - U_- + \epsilon}{U_+ - U_-}$$

Now, we know that $t < X$ and $\alpha - t$ is at least the short player's padding, and thus $\alpha - t > Y$. Thus it suffices to have $Y > \frac{U_+ - U_-}{3(U - U_- + \epsilon)}$. But Y is a geometric random variable with parameter β, so we require that

$$\Pr\left(Y > \frac{U_+ - U_-}{3(U - U_- + \epsilon)}\right) = (1-\beta)^{\frac{U_+ - U_-}{3(U - U_- + \epsilon)}} \quad (1)$$

Now, if we set

$$\beta = 1 - (1 - \delta)^{\frac{3(U - U_- + \epsilon)}{U_+ - U_-}}$$

then $\Pr(Y > \frac{U_+ - U_-}{2(U - U_- + \epsilon)}) = 1 - \delta$ so that by Lemma 3.2, Lemma 3.3 and equation 1, with probability at least $1 - \delta$ no player can gain more than ϵ in expectation by guessing the position of the secret before the definitive iteration. It follows that with probability at least $1 - \delta$ the protocol is an ϵ-Nash equilibrium.

We now analyze the resource costs of our protocol. The communication tree has $2n-1$ nodes. Each player is mapped to one leaf and at most one internal node. Players only communicate with their neighbors in the tree. So on each round, in the *Up-Stage* player i acting as a leaf, sends two field elements (a share and a tag) up to his parent, and then two more acting as an internal node. In the *Down-Stage*, acting as an internal node he sends two field elements each to his left and right child. Thus each player sends at most eight field elements per round, and thus the expected number of messages sent is $O(1)$ and the expected number of bits sent is $O(\log q)$ which does not depend on n. Since the expected number of rounds is constant and the tree has $O(\log n)$ height, it follows that the expected latency is $O(\log n)$. $\square$

3.2 A Note on Backwards Induction

The *backwards induction* problem arises when a multi-round protocol has a last round number that is known to all players. In particular, if it is globally known that the last round of the protocol is ℓ, then on the ℓ-th round, there is no longer any fear or reprisal to persuade a player to follow the protocol. But then if no player follows the protocol in the ℓ-th round, then in the $(\ell-1)$-th round, there is no reason for any player to follow the protocol. This same logic continues backwards to the very first round.

The backwards induction problem can occur with protocols that make cryptographic assumptions, since there will always be some round number, ℓ, in which enough time has passed so that even a computationally bounded player can break the cryptography. Even though ℓ may be far off in the future, it is globally known that the protocol will end at round ℓ, and so by backwards induction, even in the first round, there is no incentive for a player to follow the protocol.

As in [8], we protect against backwards induction by having both long and short players. As the above analysis shows, if β is chosen sufficiently small, we can ensure that Y will be large enough so that the probability of making a correct guess as to when the protocol ends is too small to enable profitable cheating for any player. Thus, even when a player gets to the second to the last element in all its lists, it can not be very sure that the protocol will end in the next round. All players are aware of these probabilities at the beginning of the protocol, and thus each player knows that no other player will be able to accurately guess when the protocol ends.

4. SCALABLE RATIONAL SECRET SHARING WITH ABSENT PLAYERS

In this section we discuss m-out-of-n secret sharing where $m < n$. Here we want a subset of m or more of the players to be able to reconstruct the secret even when the remaining players are absent. However, fewer than m players should *not* be able to reconstruct the secret on their own. As discussed previously, it does not seem possible to design *scalable* algorithms for secret sharing in the case when either m is much smaller than n or the subset of size m of active players may be chosen in a completely arbitrary manner. We now describe the situation where (1) $m < n$, but $m = \Theta(n)$; and (2) the subset of active players does not depend on the random bits of the dealer. More precisely, we will present

an algorithm with parameters τ and λ (in addition to the parameters of our n-out-of-n algorithm) such that when the number of active players $m > (\tau + \lambda)n$ then with high probability the algorithm is an ϵ-Nash equilibrium for the active players and all the active players learn the secret. On the other hand, if the number of active players $m < (\tau - \lambda)n$, then with high probability the active players cannot reconstruct the secret.

We will assume, as in previous work, that the set of active players is known to all the active players, before the start of the algorithm. This set may be arbitrary or randomly chosen but it does not depend on the random communication tree created by the algorithm.

Here we present a sketch of the algorithm and analysis. Details will appear in the full version of the paper.

Since the algorithm is a variant of the n-out-of-n scheme, for conciseness, we describe here only the places where the two algorithms differ. Let c be a large constant (which we will specify later.) First, the dealer partitions the players into $\frac{n}{c \log n}$ mutually disjoint groups of $c \log n$ players each.[3] This is done using a random permutation of the players. (The first $c \log n$ players in the permutation are the first group, the next $c \log n$ are the second group and so on.) The next step is to create a communication tree as illustrated in Figure 3. The tree is a complete binary tree with $\frac{n}{c \log n}$ leaves (*i.e.* it has $\frac{2n}{c \log n} - 1$ nodes in all.) All of the nodes are *supernodes*, in that they correspond to groups of $c \log n$ players instead of single players. Each group is assigned to a distinct leaf node. Then all but one of the groups are additionally assigned to distinct internal nodes, as in the n-out-of-n algorithm.

Dealer's Protocol: In order to create the sharelist S_i for each player i, the dealer once again creates iterated shares starting with the value h_t (the indicator for round t) at the root of the tree, and creating 2-out-of-2 Shamir shares of the value at a node, as values for its children. Then at each leaf node, $\tau c \log n$-out-of-$c \log n$ Shamir shares are created for the value at the leaf node. These shares are added to the players' sharelists, one for each of the $c \log n$ players assigned to the leaf. Note that this means that $\tau c \log n$ or more of the $c \log n$ players at a leaf can reconstruct the value at the leaf, but fewer than $\tau c \log n$ of them cannot. The algorithm for populating the other lists (L_i, T_i^w, and $H_i^{w,w'}$) is similar to that used for the n-out-of-n algorithm.

Player's Protocol: The protocol for the players consists of an upward and downward phase. In the upward phase, the active players at each leaf node first send their iterated shares to all the other players at the same leaf, in a clique. Assuming that $\tau c \log n$ players are active at the leaf, they can now reconstruct the value at the leaf. This value is sent by all the active players to all the players at the parent supernode. The players in these supernodes use the shares received from the two children to reassemble the appropriate value for that supernode, check this value with the hash and tag scheme and then send the value up to the players in the supernode above them. This continues until the root node is reached, at which point, the value h_t for the current round is reconstructed.

---

[3]It is convenient to assume that $c \log n$ is an integer and divides n. In fact, for arbitrary n, there will be $\left\lfloor \frac{n}{\lceil c \log n \rceil} \right\rfloor$ groups, each with either $\lceil c \log n \rceil$ or $\lceil c \log n \rceil + 1$ players.
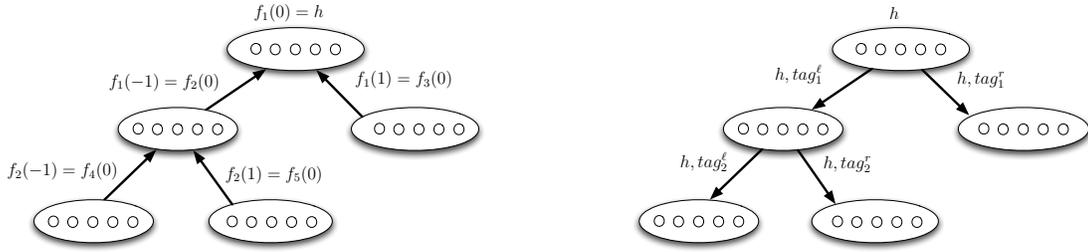
Figure 3: Illustration of the scalable m-out-of-n secret sharing algorithm. The supernodes each consist of $c \log n$ players, selected as described in Section 4. The players in each leaf supernode are connected in a clique. For any parent and child supernodes, the players in the child and the players in the parent are connected in a full bipartite graph. On the left is the up stage of the algorithm; on the right is the down stage.

In the downward phase, h_t is sent down the tree from each parent supernode to its children, using all-to-all communication among the players in the supernodes, until at the bottom, all players learn the value of h_t for that round. The players now act on the information they receive (or do not receive) exactly as they do in the n-out-of-n algorithm.

Analysis Sketch: In order to show correctness of the algorithm, we need to make two arguments. The first argument is that the active players are well distributed in the tree, so that at any stage of the players' protocol sufficiently many shares are available to do the desired reconstruction. The second argument is that is is an ϵ-Nash equilibrium for the active players to follow the protocol. The proof of this part is essentially identical to the proof that the n-out-of-n protocol was an ϵ-Nash equilibrium for all the players, and we omit it here. In the remainder of this section, we sketch why reconstruction is possible with high probability, despite absent players.

We note that for the value at an internal node to be reconstructed it is necessary that the values of both of its children be received. Therefore, it is necessary that there is at least one active player at each internal node. Since no other reconstruction is to be done at internal nodes, this is also sufficient. However, note that the group of $c \log n$ players assigned to an internal node is the same group as those assigned to some leaf. Hence if there are no active players at some internal node, then there is a leaf at which there are also no active players. Since at the leaves we will have a more stringent requirement for how many players need to be active, it is sufficient to consider the failure of the algorithm at the leaves.

Now, the value at a leaf is distributed as $\tau c \log n$-out-of-$c \log n$ Shamir shares to the $c \log n$ players associated with that leaf. Thus, this value can be reconstructed if and only if there are at least $\tau c \log n$ active players at the leaf. Moreover, in order for the protocol to succeed, the values at all the leaves must be reconstructable.

Recall that the players are assigned to leaves of the tree by the following random process: first choose a random permutation of the n players, then partition into $\frac{n}{c \log n}$ mutually disjoint groups by choosing successive contiguous blocks of length $c \log n$. Moreover, the choice of which players are active is made *independently* of the above process. It turns out that in this case, the number of active players in a fixed block is tightly concentrated around its mean. A more precise statement is in the following lemma, whose proof is an

application of the Azuma-Hoeffding inequality. We omit the details here.

LEMMA 4.1. *Let $b_1, b_2, \ldots, b_n$ be n bits such that exactly m of them are 1 and $n - m$ of them are 0. Let $\sigma \in S_n$ be a random permutation of n symbols and $b_{\sigma(1)}, b_{\sigma(2)}, \ldots, b_{\sigma(n)}$ be the induced permutation on bits. Then for any contiguous block $b_{\sigma(i+1)}, b_{\sigma(i+2)}, \ldots, b_{\sigma(i+c \log n)}$ of length $c \log n$, the number U of bits in the block which are 1 has expectation $\frac{mc \log n}{n}$ and satisfies*

$$\Pr\left(U - \frac{mc \log n}{n} > \lambda c \log n\right) \leq e^{-(\lambda^2 c \log n)/2}$$

and

$$\Pr\left(U - \frac{mc \log n}{n} < -\lambda c \log n\right) \leq e^{-(\lambda^2 c \log n)/2}$$

We can now prove

THEOREM 4.2. *Our m-out-of-n secret sharing algorithm is correct with high probability.*

PROOF. Imagine the permutation of players is a bit string such that every bit corresponds to one player. If a bit is 1 it means the corresponding player is active and if a bit is 0 the corresponding player is inactive.

Case (1) $m \geq (\tau + \lambda)n$. Consider a particular leaf node ℓ and let U_ℓ denote the number of active players at ℓ. By Lemma 4.1,

$$\Pr(U_\ell < \tau c \log n) = \Pr\left(U_\ell - \frac{mc \log n}{n} < \left(\tau - \frac{m}{n}\right)c \log n\right)$$

$$\leq \Pr\left(U_\ell - \frac{mc \log n}{n} < -\lambda c \log n\right)$$

$$\leq e^{-(\lambda^2 c \log n)/2}$$

Thus the reconstruction fails at ℓ with probability at most $e^{-(\lambda^2 c \log n)/2}$. Taking a union bound over all the $\frac{n}{c \log n}$ leaves we see that the probability that the algorithm fails (which happens only when the reconstruction fails at *some* leaf) is at most $ne^{-(\lambda^2 c \log n)/2} = n^{1-\lambda^2 c/2}$. Setting $c = \frac{2(k+1)}{\lambda^2}$ we see that the probability that the algorithm fails is at most $1/n^k$.

Case (2) $m \leq (\tau - \lambda)n$. Once again consider a particular leaf node ℓ, and let U_ℓ be the number of active players at ℓ.

By Lemma 4.1,

$$\Pr(U_\ell > \tau c \log n) = \Pr\left(U_\ell - \frac{mc \log n}{n} > \left(\tau - \frac{m}{n}\right) c \log n\right)$$

$$\leq \Pr\left(U_\ell - \frac{mc \log n}{n} > \lambda c \log n\right)$$

$$\leq e^{-(\lambda^2 c \log n)/2}$$

Therefore with probability at least $1 - e^{-(\lambda^2 c \log n)/2} = 1 - 1/n^{k+1}$ reconstruction fails at ℓ. Since the failure of the reconstruction at a single leaf is sufficient to break the protocol completely (since no reconstruction is possible at any node along the path from that leaf to the root) it follows that with probability at least $1 - 1/n^{k+1}$ the algorithm fails in this case. $\square$

5. CONCLUSIONS AND FUTURE WORK

We have presented *scalable* mechanisms for rational secret sharing problems. Our algorithms are scalable in the sense that the number of bits sent by each player is $O(1)$ and the latency is at most logarithmic in the number of players. For the n-out-of-n rational secret sharing, we give a scalable algorithm that is an ϵ-Nash equilibrium to solve this problem. For m-out-of-n rational secret sharing where (1) $m = \Theta(n)$; and (2) the set of active players is chosen independently of the random bits of the dealer, we give a scalable algorithm with threshold parameter τ that is an ϵ-Nash equilibrium and ensures that for any fixed, positive λ that if (1) at least a $m/n > \tau + \lambda$ fraction of the players are active, all players will learn the secret; and (2) if fewer than a $\tau - \lambda$ fraction of the players are active, then the secret can not be recovered.

Several open problems remain. First, while our algorithms lead to a $\Theta(n)$ multiplicative reduction in communication costs for rational secure multiparty computation (SMPC), the overall bandwidth for this problem is still very high. It is known that there are certain functions, for which rational SMPC can not be performed in a scalable manner (for example, the parity function). However, we ask: Can we find classes of well-motivated functions for which scalable SMPC is possible? This is related to our second open problem which is: Can we design scalable algorithms for simulating a class of well-motivated mediators? In some sense, this problem may be harder than the SMPC problem, since some types of mediators offer different advice to different players. In other ways, the problem is easier: a simple global coin toss is an effective mediator for many games. A final important problem is: Can we design coalition-resistant scalable algorithms for rational secret sharing?

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 53–62. ACM, 2006.

[2] R. Geambasu, T. Kohno, A. Levy, and H. Levy. Vanish: Increasing data privacy with self-destructing data. In *Proceedings of the 18th conference on USENIX security symposium*, pages 299–316. USENIX Association, 2009.

[3] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 218–229. ACM, 1987.

[4] S. Gordon and J. Katz. Rational secret sharing, revisited. *Security and Cryptography for Networks*, pages 229–241, 2006.

[5] J. Halpern and V. Teague. Rational secret sharing and multiparty computation: extended abstract. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, page 632. ACM, 2004.

[6] S. Izmalkov, S. Micali, and M. Lepinski. Rational secure computation and ideal mechanism design. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 585–594. IEEE, 2005.

[7] V. King and J. Saia. Breaking the O (n 2) bit barrier: scalable byzantine agreement with an adaptive adversary. In *Proceeding of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 420–429. ACM, 2010.

[8] G. Kol and M. Naor. Games for exchanging information. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 423–432. ACM, 2008.

[9] M. Lepinksi, S. Micali, and A. Shelat. Collusion-free protocols. In *ANNUAL ACM SYMPOSIUM ON THEORY OF COMPUTING*, volume 37, page 543. Citeseer, 2005.

[10] M. Lepinski, S. Micali, C. Peikert, and A. Shelat. Completely fair SFE and coalition-safe cheap talk. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 1–10. ACM, 2004.

[11] A. Lysyanskaya and N. Triandopoulos. Rationality and adversarial behavior in multi-party computation. *Advances in Cryptology-CRYPTO 2006*, pages 180–197, 2006.

[12] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 73–85. ACM, 1989.

[13] M. Wegman and J. Carter. New hash functions and their use in authentication and set equality. *Journal of computer and system sciences*, 22(3):265–279, 1981.