Brief Announcement: Proactive Secret Sharing with a Dishonest Majority

Shlomi Dolev
Ben-Gurion University
Beer-Sheva, Israel
dolev@cs.bgu.ac.il

Karim ElDefrawy
HRL Laboratories
Malibu, California, USA
kmeldefrawy@hrl.com

Joshua Lampkins
HRL Laboratories
Malibu, California, USA
jdlampkins@hrl.com

Rafail Ostrovsky
University of California Los Angeles
Los Angeles, California, USA
rafail@cs.ucla.edu

Moti Yung
Snapchat & Columbia University
New York, New York, USA
moti@cs.columbia.edu

ABSTRACT

In a secret sharing scheme a dealer shares a secret s among n parties such that an adversary corrupting up to t parties does not learn s, while any $t + 1$ parties can efficiently recover s. Over a long period of time all parties may be corrupted thus violating the threshold, which is accounted for in *Proactive Secret Sharing (PSS)*. PSS schemes periodically rerandomize (refresh) the shares of the secret and invalidate old ones. PSS retains confidentiality even when *all parties* are corrupted over the lifetime of the secret, but no more than t during a certain window of time, called the refresh period. Existing PSS schemes only guarantee secrecy in the presence of an honest majority with less than $n/2$ total corruptions during a refresh period; an adversary corrupting a single additional party, even if only passively, obtains the secret. This work is *the first feasibility result demonstrating PSS tolerating a dishonest majority*, it introduces the first PSS scheme secure against $t < n$ passive adversaries without recovery of lost shares, it can also recover from honest faulty parties losing their shares, and when tolerating e faults the scheme tolerates $t < n - e$ passive corruptions. A non-robust version of the scheme can tolerate $t < n/2 - e$ active adversaries, and mixed adversaries that control a combination of passively and actively corrupted parties that are a majority, but where less than $n/2 - e$ of such corruptions are active. We achieve these high thresholds with $O(n^4)$ communication when sharing a single secret, and $O(n^3)$ communication when sharing multiple secrets in batches.

Keywords

secret sharing, dishonest majority, proactive security, proactive secret sharing, non-robust secret sharing

1. BACKGROUND AND RELATED WORK

Over a long period of time all parties in a secret sharing scheme [9, 3] may be compromised and the threshold may be temporarily violated. The *proactive security model* [7] deals with an adversary's ability to eventually compromise all the parties, it protects against a *mobile adversary* capable of corrupting all parties in a distributed system or protocol during the execution but with the following limitations: (1) only a constant fraction of parties can be corrupted during any round of the protocol; (2) parties are periodically rebooted (or reset) to a predictable and pristine initial state, guaranteeing a small fraction of corrupted parties, assuming that the corruption rate is not more than the reboot rate.

Existing Proactive Secret Sharing (PSS) schemes, e.g., [7, 5, 8, 1, 2], are insecure when a majority of the parties are compromised, even if the compromise is only passive. Such schemes typically store the secret as the free term in a polynomial of degree $t < n/2$, thus once an adversary compromises $t + 1$ parties (even if only passively), it can reconstruct the polynomial and recover the secret. New and different techniques other than the above are required to construct PSS secure against dishonest majorities. Developing such techniques and the first PSS scheme secure against dishonest majorities is the main contribution of our work. Our PSS scheme provides security against $t < n$ passive adversaries without recovery of lost shares, it can also recover from honest faulty parties losing their shares, and when tolerating e faults it is secure against $t < n - e$ passive corruptions. A non-robust version of the scheme can tolerate $t < n/2 - e$ active adversaries, and mixed adversaries that control a combination of passively and actively corrupted parties that are a majority, but where less than $n/2 - e$ of such corruptions are active. Existing PSS schemes *cannot handle* such high combined corruption thresholds.

2. PRELIMINARIES

Below are preliminaries for the rest of the paper.

2.1 System and Network Model

We consider a set of n parties $\mathcal{P} = \{P_i\}_{i=1}^n$ connected via a synchronous network, and an authenticated broadcast channel. Each pair of parties also establish private secure authen-

ticated communication channels which can be instantiated via appropriate encryption and digital signature schemes.

Time Periods and Refresh Phases: We assume that all parties are synchronized via a global clock. Time is divided into *time periods or epochs*; at the beginning of each period all parties engage in an interactive refresh protocol (also called refresh phase). At the end of the refresh phase all parties hold new shares for the same secret, and delete their old shares. We note that honest parties *must* delete their old shares so that if they get compromised in future periods, the adversary *cannot* recover their shares from old periods.

2.2 Variations of Secret Sharing

A Secret Sharing (SS) scheme consists of two protocols, `Share` and `Reconstruct`. `Share` allows a dealer to share a secret, s, among n parties such that it remains secure against an adversary that controls up to t parties, while allowing any group of $t+1$ or more uncorrupted parties to reconstruct the secrets via `Reconstruct`. A problem with standard secret sharing, e.g., Shamir's scheme [9], is that a dishonest dealer may deal inconsistent shares from which $t+1$ or more parties may not be able to reconstruct the secret. This malicious behavior can be prevented by augmenting the secret sharing scheme with homomorphic commitments; this is essentially what a Verifiable Secret Sharing (VSS) scheme achieves. A VSS scheme allows parties to verify that a dealer has correctly shared a secret. We utilize the VSS scheme of Feldman [4], where security is based on the hardness of computing discrete logarithms over $\mathbb{Z}_p$ for a large prime p. The definition of a Proactive Secret Sharing (PSS) scheme is similar to that of a standard SS scheme, with the addition of two new protocols to perform `Refresh` and `Recovery` for securing the secrets against a mobile adversary that can corrupt all n parties over a long period of time. The `Refresh` protocol refreshes data to prevent a mobile adversary from collecting (over a long period) a large number of shares that could exceed the reconstruction threshold and thus reveal the secret. The `Recovery` protocol allows rebooted parties, or faulty honest ones, to recover their shares and thus prevents the adversary from destroying the shared secret(s).

3. PROACTIVE SECRET SHARING SECURE AGAINST A DISHONEST MAJORITY

This section first overviews our PSS scheme and the intuition behind it, it then focuses on the two new protocols for refreshing and recovering shares; protocols for sharing and reconstructing a secret in our PSS scheme are similar to [6] and are briefly described due to space constraints. All protocols are secure against a dishonest majority.

3.1 Intuition and Overview of Operation

Field operations occur over a finite field $\mathbb{Z}_p$ for some prime p. Let α be a generator of $\mathbb{Z}_p^*$ and let $\beta = \alpha^{-1}$. In the case of multiple secrets, secrets will be stored at locations that are multiples of β, i.e., if $f(x)$ is a sharing polynomials then $f(\beta^1)$ and $f(\beta^2)$ will evaluate to secret 1 and secret 2 respectively, while shares will be computed as the evaluation of $f(x)$ at different values of α, i.e., $f(\alpha^1)$ and $f(\alpha^2)$ are the shares of party 1 and 2 respectively. We note that in the case of sharing a single secret, only one β is needed, and in that case it will not be the inverse of α, traditionally it has been the case that for single secrets $\beta = 0$, thus the secret

s is stored at the free term of the sharing polynomial, i.e., $f(0) = s$. The shares for a single secret can be evaluations of $f(x)$ at indices of the parties, i.e. $f(1)$, $f(2)$ ... $f(n)$, or at pre-agreed upon points corresponding to each party such as $f(\alpha^1)$, $f(\alpha^2)$, ..., $f(\alpha^i)$.

To simplify the illustration we assume here and in description of our share, reconstruct and refresh protocols (`DM-Share`, `DM-Reconstruct`, and `DM-Refresh`), that the adversary only compromises nodes temporarily, so only refreshing of shares is needed. If parallel share recovery (`DM-Recover`) for rebooted nodes is required, the tolerated threshold is decreased by the maximum number of nodes to be rebooted in parallel or that can loose their shares (this can also be due to non-malicious faults). If nodes are serially rebooted such that only a single share is to be recovered at any instant, then the tolerated thresholds are decreased by 1.

Per the discussion in Section 1, to tolerate a dishonest majority it is not enough to store secrets in the free term, or as other points on a polynomial. What is needed is to encode secrets in a form resistant to a dishonest majority of up to $n-1$ parties. This can be achieved by first additively sharing the secret into n random summands (this provides security against $t < n$ passive adversaries), then those random additive summands may be shared and proactively refreshed using methods that can tolerate $t < n/2$ active adversaries with aborts, i.e., if less than $n/2$ of the parties are actively corrupted their misbehavior will be detected and flagged by the rest of the honest parties (constituting a majority) while ensuring confidentiality of the shared secret, even if up to $n/2$ passively corrupted parties exist among the remaining parties. This is the blueprint that we follow. Specifically, we start from the gradual secret sharing schemes from [6], develop two new protocols to verifiably generate random refreshing polynomials with the required properties, i.e., they have a random free term that encodes random additive shares that add up to zero. To recover shares with the above security guarantees, we observe that it is enough that the recovery protocol ensures security against $t < n/2$ ($t < n/2 - e$ with e faults) active adversaries, as passive adversaries only generate random polynomials and send them to the recovering party, i.e., if they respect the polynomial generation process, and as long as one party generates a random polynomial, the rest of the $n-1$ potentially passively corrupted parties will only see new random polynomials with the appropriate degrees.

3.2 Sharing and Reconstructing Secrets

`DM-Share`: to share a secret s the first step is to split it into n random summands, $s = rs_1 + rs_2 + \cdots + rs_n$. These n random summands are then each verifiably shared using a verifiable linear secret sharing scheme, e.g., using a verifiable version of Shamir's scheme (that relies on homomorphic commitments such as Feldman's scheme [4]) where each random summand rs_i is stored in the free term of a random polynomial of a specific degree. These n random sharing polynomials are of increasing degrees, where the degree of polynomial p_i is i where $i \in \{0, 1, 2, \ldots, n-2, n-2\}$; note that the first one is a constant and that two of the polynomials have degree $n-2$ if recovery of shares of one rebooted or faulty node is required. If only refreshing is required then degrees of the polynomials go up to $n-1$ instead.

`DM-Reconstruct`: to reconstruct a secret, each party broadcasts its shares, and each party then interpolates the n ran-

dom sharing polynomials $p_i(x)$ and recovers the n random summands from the free terms, i.e., $rs_i = p_i(0)$. The secret is reconstructed as the summation of all the recovered free terms, $s = \sum_{i=0}^{n-1} rs_i = \sum_{i=0}^{n-1} p_i(0)$ when degrees are from 0 to $n-1$.

3.3 Refreshing Shares for Dishonest Majority

`DM-Refresh`: In our refresh protocol, each party generates n random refreshing polynomials with the appropriate degrees (i.e., from a single constant term, corresponding to degree 0, to degree $n-1$); each party then verifiably shares these refreshing polynomials with the other parties by committing to their coefficients and distributing shares of these polynomials as their evaluation at the indices of the parties similar to Feldman's VSS [4]. These refreshing polynomials should satisfy the following condition: they have random constant coefficients that add up to 0 (to match the case when a single secret is shared in the free term). This can be ensured by homomorphically checking that the polynomials shared by each party have this property. This condition ensures that the shared secret remains unchanged. Once each party receives all the shares generated by other parties, they add them to their local shares and delete the shares that resulted from the previous execution of the refresh protocol.

3.4 Recovering Shares for Dishonest Majority

`DM-Recover`: To be able to recover shares of a single party, then instead of initially sharing the n additive summands with polynomials of degrees $n-1$ to 0 (i.e., a single constant), the degrees will be $n-2$ to 0, with two of the summands shared with two different polynomials of degree $n-2$. This allows $n-1$ parties to recover shares of a single party that is rebooted. (For e parallel recoveries polynomial degrees will be $n-e-1$ to 0.) In each refresh period there are n current sharing polynomials with degrees ranging from $n-2$ to 0, and each party has a share for each of these polynomials. When a single party P_{rc} is rebooted and needs to recover its shares, i.e., the evaluation of each of the current sharing polynomials at P_{rc}'s evaluation point α^{rc}, the other parties need to generate and verifiably share n random polynomials that evaluate to the same values as the current sharing polynomials at α^{rc}. To achieve this, parties generate and verifiably share n random recovery polynomials that evaluate to 0 at α^{rc}. All parties add their local shares of the current sharing polynomials to the shares of these random recovery polynomials; this results in n shared random recovery polynomials that have only the point at α^{rc} in common with the current sharing polynomials. All parties then send their shares of these n shared random recovery polynomials to P_{rc}, and P_{rc} can then interpolate these polynomials without learning anything about the secret or the actual sharing polynomials of the current period. We note that passively corrupted parties in the recovery will execute the protocol correctly, and actively corrupted parties are limited to $t < n/2-e$ with e faults; we only need a recovery protocol secure against active adversaries because only the recovering party receives information. Every other party generates random data and shares it with the rest of the parties, so there is no information related to the secret that is revealed to any party. As long as there is a single honest party, the random recovery polynomials that party generates ensures randomness of the overall recovery polynomials; this ensures that the only information P_{rc} learns are its n shares at α^{rc}.

4. CONCLUSION AND OPEN QUESTIONS

We present the *first feasibility result for Proactive Secret Sharing (PSS) for a dishonest majority and the first such PSS scheme.* Our main PSS scheme is secure against $t < n$ passive adversaries without recovery of lost shares, and when tolerating e faults that result in lost shares the scheme resists $t < n-e$ passive corruptions. A non-robust version of the scheme can tolerate $t < n/2-e$ active adversaries, and mixed adversaries that control a combination of passively and actively corrupted parties that are a majority, but where less than $n/2-e$ of such corruptions are active. The following issues remain open: (i) It is unclear what is the lowest possible communication required for a PSS scheme secure against a dishonest majority. We achieve $O(n^3)$ for batches of secrets, it remains open if this can be reduced to $O(n)$ or $O(1)$. We conjecture that $O(n)$ is the lower bound for our PSS blueprint which first shares the secret via an additive sharing scheme. Such an additive scheme does not seem to be amenable to batching in a straightforward manner; it is currently not obvious to us how to batch it without destroying the secret. (ii) There are currently no PSS schemes secure against dishonest majorities of up to $n-1$ and that operate over asynchronous networks.

5. REFERENCES

[1] J. Baron, K. ElDefrawy, J. Lampkins, and R. Ostrovsky. How to withstand mobile virus attacks, revisited. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*, PODC '14, pages 293–302, New York, NY, USA, 2014. ACM.

[2] J. Baron, K. ElDefrawy, J. Lampkins, and R. Ostrovsky. Communication-optimal proactive secret sharing for dynamic groups. In *Proceedings of the 2015 International Conference on Applied Cryptography and Network Security*, ACNS '15, 2015.

[3] G. R. Blakley. Safeguarding cryptographic keys. *Proc. of AFIPS National Computer Conference*, 48:313–317, 1979.

[4] P. Feldman. A practical scheme for non-interactive verifiable secret sharing. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, SFCS '87, pages 427–438, Washington, DC, USA, 1987. IEEE Computer Society.

[5] A. Herzberg, S. Jarecki, H. Krawczyk, and M. Yung. Proactive secret sharing or: How to cope with perpetual leakage. In *CRYPTO*, pages 339–352, 1995.

[6] M. Hirt, U. Maurer, and C. Lucas. A Dynamic Tradeoff between Active and Passive Corruptions in Secure Multi-Party Computation. In R. Canetti and J. A. Garay, editors, *Advances in cryptology - CRYPTO 2013 : 33rd Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013 : proceedings*, volume 8043 of *Lecture notes in computer science*, pages 203–219, Heidelberg, 2013. Springer.

[7] R. Ostrovsky and M. Yung. How to withstand mobile virus attacks (extended abstract). In *PODC*, pages 51–59, 1991.

[8] D. Schultz. *Mobile Proactive Secret Sharing*. PhD thesis, Massachusetts Institute of Technology, 2007.

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