Generic On-line/Off-line Threshold Signatures

Chris Crutchfield, David Molnar, David Turner, and David Wagner

University of California, Berkeley
{cyc, dmolnar, dbturner, daw}@cs.berkeley.edu

Abstract. We present *generic on-line/off-line threshold signatures*, in which the bulk of signature computation can take place "off-line" during lulls in service requests [6]. Such precomputation can help systems using threshold signatures quickly respond to requests. For example, tests of the Pond distributed file system showed that computation of a threshold RSA signature consumes roughly 86% of the time required to service writes to small files [12]. We apply the "hash-sign-switch" paradigm of Shamir and Tauman [16] and the distributed key generation protocol of Gennaro et al. [8] to convert any existing secure threshold digital signature scheme into a threshold on-line/off-line signature scheme. We show that the straightforward attempt at proving security of the resulting construction runs into a subtlety that does not arise for Shamir and Tauman's construction. We resolve the subtlety and prove our signature scheme secure against a static adversary in the partially synchronous communication model under the one-more-discrete-logarithm assumption [2]. The on-line phase of our scheme is efficient: computing a signature takes one round of communication and a few modular multiplications in the common case.

Keywords: On-line/Off-line, Signature Schemes, Threshold Cryptography, Chameleon Hash Functions, Bursty Traffic.

1 Introduction

We present *generic on-line/off-line threshold signatures* to improve the performance of threshold signature schemes, and we show how to construct such signatures from existing threshold signature schemes. In a threshold signature scheme, given a group of n players, and a threshold $t < n$, no subset of the players of size at most t can generate a signature. In other words, unlike standard signature schemes — in which a single player must protect his or her secret key — at most t of the n players in a threshold signature scheme may be compromised without endangering the security of the signature scheme.

Threshold signatures have been applied in several areas to avoid concentrating trust in any single entity. For example, OceanStore [10, 12] is a large-scale distributed data storage system that requires the computation of threshold signatures by an "inner ring" of servers for performing a Byzantine agreement when writing a file. Latency tests in Pond [12], the OceanStore prototype, show that for a 4 KB write, 77.8 ms out of 90.2 ms total time to service the write operation is spent on computing Shoup's RSA threshold signature scheme [17]. Therefore, computation is the dominant factor; although network communication and local file system access contribute to the time, the bulk of the contribution to service time comes from computing the threshold signatures [12].

Optimizing threshold signature computation is particularly important for distributed file systems because small file writes are common [14]. For example, Baker et al. found that for a file trace from the Sprite file system, 80% of all sequential transfers were less than 2300 bytes in length [1]. For larger files in Pond (2 MB), there is little change in the time spent computing the threshold signature; instead, the time spent on writing the file dominates the threshold signature time. Even so, because threshold signature computation takes up 86% of the time to service a small write in Pond, optimizing this computation improves the common case. Threshold signatures have also been applied as part of other applications, such as distributed certificate authorities, so increasing their performance can help these applications as well [19].

Our Approach. In an on-line/off-line scheme [6], servers can perform the bulk of the computation in an *off-line phase* before even seeing the message to be signed. The results of this precomputation are saved and then used in the *on-line phase* when a message must be signed. Because distributed systems often have "bursty" traffic, resources are available for such precomputation. For example, during the day and evening, traffic is high, but during the night and morning, traffic is low. Enabling threshold signatures to be computed off-line allows systems such as OceanStore to build up a stockpile of precomputed values while traffic is low. These values can be used to quickly sign messages later when traffic is high. Furthermore, other distributed file systems have been observed to have bursty traffic [15, 18], and so they can enjoy the benefits of our on-line/off-line threshold signature scheme. In addition, our scheme offers several advantages over existing schemes. For instance, implementing an on-line/off-line scheme such as Threshold DSS [7] may not be a desirable solution if much effort has already been spent on implementing a threshold signature scheme. Another advantage is that any threshold signature scheme secure against random message attack can be used with our on-line/off-line scheme to create a threshold scheme secure against an adaptive chosen message attack.

The main idea of our scheme is to apply the "hash-sign-switch" paradigm of Shamir and Tauman [16] to a threshold signature scheme. In this paradigm, we make use of a *chameleon hash function,* which is a special type of two-argument hash function $CH_{HK}(m, r)$ endowed with a public and secret key [9]. Knowledge of the public key HK allows one to evaluate the hash function, while knowledge of the secret key allows one to find collisions. Shamir and Tauman show that any standard signature scheme can be converted to an on-line/off-line scheme as follows: for the off-line phase, compute a standard signature on $CH_{HK}(a, r)$, where a and r are chosen randomly. Then, at the on-line phase, given the message m, use the secret key to find an r' such that $CH_{HK}(m, r') = CH_{HK}(a, r)$. The signature on $CH_{HK}(a, r)$ together with r' then forms a signature on the message

m; in a sense, we "switch" m for the random value a. We refer to the signed value of $CH_{HK}(a, r)$ as the *signature stamp.* If finding a collision in the chameleon hash is more efficient than signing the message directly (as is the case for several chameleon hash functions), this is a net performance win.

Overview of Our Construction. For our work, we focus on the specific chameleon hash function $CH_{HK}(m, r) = g^r h^m \bmod p$ with public key $HK = (p, g, h)$ and the secret key y is the discrete logarithm of h to the base g. We show how to use the discrete logarithm distributed key generation algorithm of Gennaro et al. [8] to perform chameleon hash key generation and computation of the signature stamp. We then show an efficient distributed algorithm for finding collisions with low overhead per player. We stress that *no trusted dealer* is required by our scheme; given an underlying threshold signature scheme with distributed key generation and distributed signing algorithms, we obtain a fully distributed signature scheme.

We also show methods for guaranteeing the robustness of our scheme using zero-knowledge proofs for verification. We provide two variants. The first is non-interactive and secure in the Random Oracle Model. The second uses an observation of Damgård and Dupont to obtain robustness at the cost of limited interaction but is secure without random oracles [4]. In both cases, instead of running verification each time a signature must be generated, we decide to forego this step and be *optimistic* because, as observed in [4], the signature shares will be correct almost always. If the signature created is not valid, then we can run the verification procedure in order to expose the corrupted players. The full details for our signature scheme appear in Sect. 3.

A Subtlety In The Proof. Surprisingly, the straightforward adaptation of the proof of Shamir and Tauman for non-threshold on-line/off-line signature schemes *fails* to establish security for our new on-line/off-line threshold scheme. The subtlety is that in our scheme, the "signature stamp" value $CH_{HK}(m, r)$ is disclosed to all players at the close of our off-line threshold phase, including the adversary. While m and r are not disclosed, the output of the chameleon hash must be broadcast to allow for "black-box" use of the underlying threshold signature scheme in creating the stamp. As a result, any attempt at simulating the adversary's view of a signature query is "pinned down" by the value of the chameleon hash encoded in the stamp. In contrast, Shamir and Tauman do not reveal any chameleon hash values associated with a message to the adversary until *after* a signing query for that message is made. Therefore, their reduction is not "pinned down" in the same way and can easily answer adversary signing queries by simply evaluating the chameleon hash function on the queried message. While this is not an attack on the threshold on-line/off-line scheme, it shows that a new idea appears necessary to prove the scheme secure.

We resolve this subtlety by first introducing a new assumption for chameleon hash functions, which we call the *one-more-r assumption.* Informally, the new assumption says that given a sequence of random "challenge" outputs $v_1, \dots, v_n$

of the chameleon hash function, the adversary may adaptively pick values v_i, provide messages m_i, and then learn r_i such that $CH_{HK}(m_i, r_i) = v_i$. Then, even given this extra information, the adversary has negligible advantage at inverting the chameleon hash on any challenge value not picked. We show that this new assumption is sufficient to prove security of our scheme. Then we justify the assumption in the case of the $g^r h^m \bmod p$ chameleon hash by showing it is implied by the *one-more-discrete-logarithm* assumption of Bellare et al [2]. This establishes the security of our scheme based on a standard assumption. The details for showing our scheme is existentially unforgeable and robust against a static adversary are in Sect. 5.

Performance Results. We analyze the performance of our scheme in Sect. 6. We show the cost of our off-line phase is dominated by the cost of the distributed discrete logarithm key generation protocol. While our off-line phase in consequence requires several rounds of communication and computation, we argue that this overhead uses resources that would otherwise sit idle. If a new request arrives at a server during a busy time, the servers can simply fall back to directly computing a threshold signature.

Finally, we show that our optimistic on-line phase obtains a factor of $\mathcal{O}\left(\frac{k}{t}\right)$ improvement in computation compared to Shoup's RSA threshold signature scheme, where k is a security parameter, while also requiring only one round of communication [17]. For example, with the parameters suggested for Pond, this is a factor of 1024 improvement. Therefore, we believe our scheme will yield in many cases much improved performance for any distributed system that currently uses threshold signature schemes.

1.1 Previous Work

The first on-line/off-line signature scheme was developed by Even, Goldreich, and Micali [6]. This scheme allowed for the conversion of any standard signature scheme into a one-time on-line/off-line signature scheme. Their result, however, increased the size of the signature by a quadratic factor. In order to mitigate this, Shamir and Tauman [16] applied the results of Krawczyk and Rabin [9], using chameleon hash functions to construct a one-time on-line/off-line signature scheme that only increases the size of the signature by a factor of two. Although smart cards appear to be an important application of on-line/off-line signatures as noted in [6, 16], the application to bursty traffic has received little attention.

The origins of threshold signatures and threshold cryptography can be traced back to Desmedt and Frankel [5]. Some examples of threshold signatures include a robust threshold DSS signature scheme, which is an on-line/off-line scheme, by Gennaro et al. [7], and a robust, non-interactive threshold RSA signature scheme by Shoup [17]. The latter construction is the signature scheme implemented in Pond [12], a prototype version of the OceanStore [10] design, and partly our motivation for this paper.

Table 1. Comparison between Shoup's Threshold RSA and our On-line/Off-line Threshold Scheme

Threshold Sig. Schemes:	Shoup's RSA Scheme	Our On-line/Off-line Scheme
Key Generation	$\mathcal{O}(k^2 nt \log t + k^3) + K_{\mathrm{RSA}}$	$K_{\mathrm{On/Off}} + K_{\mathrm{DKG}}$
Off-line Phase	None	$3K_{\mathrm{DKG}} + \mathcal{O}(k^2) + \tau$
On-line Player	$\mathcal{O}(k^3)$	$\mathcal{O}(k^2)$
On-line Reconstruction	$\mathcal{O}(tk^3)$	$\mathcal{O}(t^2k^2)$
On-line Rounds of Comm.	1	1

1.2 Our Results

We compare our optimistic on-line/off-line threshold signature scheme with that of Shoup's signature scheme [17]. Shoup describes two variants of an RSA threshold signature scheme, and it is the first variant that we compare our scheme against. In both schemes, let n be the number of players, $t < \frac{n}{3}$ be the threshold[1], and $k \in \mathbb{N}$ be a security parameter. Our construction requires $2t+1$ players to construct a signature and tolerates the participation of at most t corrupted players. We analyze the bit complexity of both schemes using the following metrics and show the results in Table 1:

- Key Generation Complexity — Work done to perform key generation and distributing private key shares among the players. Let K_{RSA} denote the bit complexity for generating the RSA public and private keys, let $K_{\mathrm{On/Off}}$ denote the bit complexity for generating public and private keys in our scheme, and let K_{DKG} denote the bit complexity for distributed key generation.
- Off-line Phase Complexity — Work done to perform precomputation, meaning the computation performed for a signature before a message arrives. Furthermore, let τ be the bit complexity for generating a standard threshold signature.
- On-line Player Complexity — Work done by a player in computing its signature share when a message arrives. Note that all players compute their signature share in parallel.
- On-line Reconstruction Complexity — Work done by the players in combining all of the signature shares and creating a signature.
- On-line Rounds of Communication — Number of rounds the players need to generate a signature.

Note that Shoup's RSA signature scheme is not considered to be an on-line/off-line scheme because no precomputation is performed. Furthermore, an optimistic version of Shoup's scheme does not reduce its asymptotic complexity in the on-line phase. Finally, referring to Table 1, we see that both schemes only require one round of communication because all of the members of the group

[1] Shoup's RSA threshold signature scheme can actually tolerate a threshold of $t < \frac{n}{2}$ and only needs $t+1$ players to generate a signature.

do not have to wait for each other when a message m arrives; instead, they can immediately compute their signature shares for m. Because we can set the modulus in both schemes to be of the same size, we can compare fairly based on the bit complexity. A more complete analysis that includes robustness can be found in Sect. 6.

2 Preliminaries

Definition 1 (Negligible Function). A function $\eta : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for all $c > 0$, $\eta(n) < \frac{1}{n^c}$ for all sufficiently large n.

Definition 2 (Discrete Logarithm Assumption). Let $p = 2q+1$ be a prime where q is a random k-bit prime, and let g be a generator for a subgroup of $\mathbb{Z}_p^*$ with order q. For all probabilistic polynomial time algorithms A, if x is chosen uniformly at random from $\mathbb{Z}_q$ and $h = g^x \pmod{p}$, then $\Pr[\mathsf{A}(p, q, g, h) = x] \leq \eta(k)$, where η is a negligible function.

Definition 3 (Chameleon Hash Function). Given a public key HK and a private key or *trapdoor* TK, which are generated with respect to a security parameter k, a message $m \in \mathcal{M}$, and a random $r \in \mathcal{R}$ where $\mathcal{M}$ is the message space, and $\mathcal{R}$ is some finite space, we denote a chameleon hash function [9] by $CH_{HK}(m, r)$, which is a hash function with the following properties:

– **Collision Resistance**. Given any probabilistic polynomial time malicious entity $\mathcal{A}$ that does not know the private key TK, but only the public key HK, define its *advantage* to be the probability of finding (m_1, r_1) and (m_2, r_2) such that $CH_{HK}(m_1, r_1) = CH_{HK}(m_2, r_2)$. We require the advantage of $\mathcal{A}$ to be negligible.
– **Trapdoor Collisions**. There exists a polynomial time algorithm A such that on inputs the pair (HK, TK), a pair $(m_1, r_1) \in \mathcal{M} \times \mathcal{R}$, and a message $m_2 \in \mathcal{M}$, then A outputs r_2 such that $CH_{HK}(m_1, r_1) = CH_{HK}(m_2, r_2)$.
– **Uniform Probability Distribution**. If $r_1 \in \mathcal{R}$ is distributed uniformly, $m_1 \in \mathcal{M}$, and $(m_2, r_2) \in \mathcal{M} \times \mathcal{R}$ such that $CH_{HK}(m_1, r_1) = CH_{HK}(m_2, r_2)$, then r_2 is computationally indistinguishable from uniform over $\mathcal{R}$.

Throughout the rest of this paper, we will work with a particular family of chameleon hash functions based on discrete logarithms. We do so because the discrete logarithm-based hash function is best suited for using Lagrange interpolation. There are also other chameleon hash functions, such as those based on factoring, for example, but the mathematics involved in the interpolation would not be as convenient.

Let $k \in \mathbb{N}$ be a security parameter. We begin by picking a k-bit Germain prime $p' \in \mathbb{N}$, which has the property that $p = 2p' + 1$ and p' are both primes. Although it is not known if there are infinitely many Germain primes, we will assume that we can find one of the appropriate size. Let g' be a generator for $\mathbb{Z}_p^*$. Now let $Q_p \subset \mathbb{Z}_p^*$ denote the subgroup of quadratic residues generated by

$g \equiv (g')^2 \pmod{p}$, so that $|Q_p| = \frac{p-1}{2} = p'$. Finally, pick the private key $y \in \mathbb{Z}_{p'}^*$. Then we define our chameleon hash function $CH_{HK} : \mathbb{Z}_{p'} \times \mathbb{Z}_{p'} \to Q_p$ to be

$$CH_{HK}(m, r) = g^{r+ym} \equiv g^r h^m \pmod{p}$$

where $h \equiv g^y \pmod{p}$ and the public key is $HK = (p, g, h)$. Although we choose to work over the group $\mathbb{Z}_p^*$, one could also work with ECC groups.

Definition 4 (Signature Scheme). A signature scheme $\mathcal{S}$ is a triple of randomized algorithms (Key-Gen, Sig, Ver) where:

- Key-Gen: $1^* \to \mathcal{PK} \times \mathcal{SK}$ is a key generation algorithm such that on input 1^k, where $k \in \mathbb{N}$ is a security parameter, it outputs (PK, SK), such that $PK \in \mathcal{PK}$, the set of all public verification keys, and $SK \in \mathcal{SK}$, the set of all secret keys.
- Sig : $\mathcal{SK} \times \mathcal{M} \to \mathcal{SIGS}$ is a signing algorithm such that $\mathcal{M}$ is the message space and $\mathcal{SIGS}$ is the signature space. For shorthand, let $S_{SK}(m) = \mathsf{Sig}(SK, m)$ for all $m \in \mathcal{M}$.
- Ver : $\mathcal{PK} \times \mathcal{M} \times \mathcal{SIGS} \to \{\mathsf{Reject}, \mathsf{Accept}\}$ is a verification algorithm such that $\mathsf{Ver}(PK, m, \sigma) = \mathsf{Accept}$ if and only if σ is a possible output of $\mathsf{Sig}(SK, m)$. Again, for shorthand, let $V_{PK}(m, \sigma) = \mathsf{Ver}(PK, m, \sigma)$ for all $m \in \mathcal{M}$ and $\sigma \in \mathcal{SIGS}$.

Definition 5 (Threshold Signature Scheme). Given a signature scheme $\mathcal{S} = $ (Key-Gen, Sig, Ver), a threshold signature scheme $\mathcal{TS}$ for $\mathcal{S}$ is a triple of randomized algorithms (Thresh-Key-Gen, Thresh-Sig, Ver) for a set of n players $\mathcal{P} = \{P_1, P_2, \ldots, P_n\}$ with threshold value t where:

- Thres-Key-Gen is a distributed key generation algorithm used by the players to create $(PK, SK) \in \mathcal{PK} \times \mathcal{SK}$ such that each $P_i \in \mathcal{P}$ receives a share SK_i of the secret key SK.
- Thresh-Sig is a distributed signing algorithm used by the players to create a signature for a message $m \in \mathcal{M}$ such that the output of the algorithm is $S_{SK}(m)$. This algorithm can be decomposed into two algorithms: signature share generation and signature reconstruction.

In this paper, we assume that $\mathcal{TS}$ is *simulatable*, as defined in Gennaro et al. [7]. This means that there exists a simulator $\mathsf{SIM}_1^{\mathcal{TS}}$ which, on input PK, simulates the view of the adversary for a run of Thresh-Key-Gen that fixes the public key to be PK. In addition, there exists a simulator $\mathsf{SIM}_2^{\mathcal{TS}}$ for Thresh-Sig, such that on input the public key PK, the message v, the signature σ of v, and the key shares $x_{i_1}, x_{i_2}, \ldots, x_{i_t}$ of the servers controlled by the adversary, simulates the view of the adversary for a run of Thresh-Sig on v that produces σ.

Definition 6 (Distributed Key Generation). A Distributed Key Generation (DKG) protocol is often used in threshold signature schemes in order to construct the public key and private key. In a DKG protocol with n players, the public key is made known to all players, whereas the private key is known by none. Instead, each player receives a *key share*, from which they can — acting in concert — recover the private key. A DKG protocol is, of course, fully distributed, and requires no trusted dealer.

In this paper, we use a discrete logarithm-based DKG protocol (where the private key is y and the public key is $h = g^y$ for some g), namely the New-DKG protocol as defined by Gennaro et al. [8]. This protocol has the property that there exists a simulator $\mathsf{SIM}^{\mathrm{DKG}}$ that on input h can simulate the interactions of the DKG protocol with a set $\mathcal{P}_{\mathcal{A}} \subset \mathcal{P}$ of players controlled by the adversary $\mathcal{A}$, where $|\mathcal{P}_{\mathcal{A}}| \leq t$, such that the resulting public key produced is fixed to be h. In addition, as a result of this simulation, $\mathsf{SIM}^{\mathrm{DKG}}$ is able to recover the key shares held by the adversary's players $\mathcal{P}_{\mathcal{A}}$.

3 An On-line/Off-line Threshold Signature Scheme

We shall construct an optimistic, generic on-line/off-line threshold signature scheme $\mathcal{TS}^{\mathrm{On/Off}} =$ (On/Off-Thresh-Key-Gen, Thresh-Sig-Off-line, Thresh-Sig-On-line, Ver) that does not require the use of a trusted dealer, and we show how existing threshold signature schemes can be used in performing a threshold computation of the signature stamp off-line. Furthermore, we use the New-DKG protocol from Gennaro, *et al* [8].

3.1 Key Generation (Done once)

On/Off-Thresh-Key-Gen

Inputs: A threshold signature scheme $\mathcal{TS} =$ (Thresh-Key-Gen, Thresh-Sig, Ver), a set of n players $\mathcal{P} = \{P_1, P_2, \ldots, P_n\}$, a threshold $t < \frac{n}{3}$, and a security parameter $k \in \mathbb{N}$.
Public Output: A set of public keys.
Private Output: All players $P_i \in \mathcal{P}$ receive a set of private keys.

1. Run Thresh-Key-Gen on input 1^k to obtain $(PK, SK) \in \mathcal{PK} \times \mathcal{SK}$ and each $P_i \in \mathcal{P}$ receives the secret key share SK_i.
2. Create a random k bit Germain prime $p' \in \mathbb{N}$, where $p = 2p' + 1$ is also a prime, and let g be a generator for Q_p.
3. Use the DKG protocol to create $h = g^y$, where $y \in \mathbb{Z}_{p'}$ is the secret key and $P_i \in \mathcal{P}$ receives the share y_i for a degree t polynomial $p_y(x) \in \mathbb{Z}_{p'}[x]$ such that $p_y(0) = y$.
4. Check that $n < p'$ so that each player $P_i \in \mathcal{P}$ has index $i \in \mathbb{Z}^*_{p'}$. Otherwise abort.
5. Publish the public keys $(PK, HK = (p, g, h))$. All players $P_i \in \mathcal{P}$ retain (SK_i, y_i).

3.2 Off-line Phase (Done per message)

In the off-line phase, we will show how to construct the chameleon hash function and create the signature stamp in a distributed manner.

Thresh-Sig-Off-line

Inputs: The same set of n players $\mathcal{P}$ and a threshold $t < \frac{n}{3}$.
Private Output: A signature stamp.

1. Use the DKG protocol to create g^r, where $r \in \mathbb{Z}_{p'}$ so that P_i receives the share r_i for another degree t polynomial $p_r(x) \in \mathbb{Z}_{p'}[x]$ such that $p_r(0) = r$.
2. Use the DKG protocol to create h^m where $m \in \mathbb{Z}_{p'}$. Each player P_i receives a share m_i for a degree t polynomial $p_m(x) \in \mathbb{Z}_{p'}[x]$ such that $p_m(0) = m$.
3. Finally, the DKG protocol is used to generate shares z_i for each $P_i \in \mathcal{P}$ of a degree $2t$ polynomial $p_0(x) \in \mathbb{Z}_{p'}[x]$ such that $p_0(0) = 0$.
4. Now g^r and h^m are both known to the players, so $CH_{HK}(r, m) = g^r h^m \pmod{p}$.
5. Use Thresh-Sig to compute the signature stamp $S_{SK}(CH_{HK}(r, m))$.

3.3 On-line Phase (Done per message)

Thresh-Sig-On-line

Inputs: A subset $\mathcal{P}' \subset \mathcal{P}$ of size $2t + 1$ and a message $m' \in \mathbb{Z}_{p'}$.
Public Output: A signature for m'.

1. For each $P_i \in \mathcal{P}'$, define $\textsf{col-1}_i = r_i - y_i m'$ and $\textsf{col-2}_i = y_i m_i + z_i$, which are P_i's share of the trapdoor collision. Then, P_i broadcasts the pair $(\textsf{col-1}_i, \textsf{col-2}_i)$ to all of the other players in $\mathcal{P}'$.
2. Define $f_i(x)$ to be $f_i(x) = \prod_{P_j \in \mathcal{P}' \setminus \{P_i\}} \frac{j-x}{j-i}$, as in the definition of Lagrange interpolation. Now use Lagrange interpolation on the shares to compute the trapdoor collision

$$\begin{aligned} r' &= \sum_{P_i \in \mathcal{P}'} (\textsf{col-1}_i + \textsf{col-2}_i) f_i(0) \\ &= \sum_{P_i \in \mathcal{P}'} (r_i + y_i m_i + z_i - y_i m') f_i(0) \\ &\equiv r + ym - ym' \pmod{p'}. \end{aligned}$$

3. In this way, the signature for message m' is

$$(S_{SK}(CH_{HK}(m, r)), m', r').$$

Notice that the definition of $\mathsf{col\text{-}2}_i$ requires adding the share z_i. This is necessary because we have to multiply the secrets y and m, so each player computes $y_i m_i$ which becomes a share of a degree $2t$ polynomial that is not chosen uniformly at random; thus, adding the share z_i will make the polynomial random. Furthermore, this degree $2t$ polynomial is the reason for requiring $t < \frac{n}{3}$.

3.4 Verification (Done per message)

Given the signature (σ, m', r'), where $\sigma \in \mathcal{SIGS}$, simply check that

$$V_{PK}(CH_{HK}(m', r'), \sigma) = \mathsf{Accept}$$

holds true, as in the standard signature scheme.

3.5 Signature Share Verification (Performed if necessary)

If $V_{PK}(CH_{HK}(m', r'), \sigma) = \mathsf{Reject}$, then some players are sending incorrect shares. In order to ensure robustness, we must be able to construct a valid signature. The naïve solution of trying all possible subsets of size $2t+1$ to construct a valid signature is unacceptable because there are an exponential number of such subsets. Instead, we will identify and remove the corrupted players. To do so, we have each player in $\mathcal{P}$ check the validity of the pair $(\mathsf{col\text{-}1}_i, \mathsf{col\text{-}2}_i)$ for each player $P_i \in \mathcal{P}'$:

1. **Verifying** $\mathsf{col\text{-}1}_i$. Because g^{r_i} and g^{y_i} are known values from the DKG protocol, we can compute for each $P_i \in \mathcal{P}'$, $g^{r_i} \cdot (g^{y_i})^{-m'} = g^{r_i - y_i m'} \pmod p$ and confirm that $g^{\mathsf{col\text{-}1}_i} = g^{r_i - y_i m'}$ as desired.
2. **Verifying** $\mathsf{col\text{-}2}_i$. Although we have access to g^{z_i} from the DKG protocol, we do not have $g^{y_i m_i}$. Instead, what we will do is confirm that the discrete logarithm of $g^{\mathsf{col\text{-}2}_i} g^{-z_i} = g^{\mathsf{col\text{-}2}_i - z_i}$ to the base g^{m_i} is equal to the discrete logarithm of g^{y_i} to the base g. Now we can apply Chaum and Pedersen's ZKP for equality of discrete logarithms [3] with the Fiat-Shamir heuristic: Let $d = g^{y_i}, e = g^{m_i}$, and $f = g^{\mathsf{col\text{-}2}_i - z_i}$. Player $P_i \in \mathcal{P}'$ chooses $r \in \mathbb{Z}_{p'}$ uniformly at random and computes $H(g, d, e, f, g^r, e^r) = c$, where H is a random oracle and c is the challenge. P_i computes $v = y_i c + r$ and broadcasts the pair (c, v). Finally, all players compute and confirm that $H(g, d, e, f, g^v d^{-c}, e^v f^{-c}) = c$.

If any of the shares are deemed incorrect, then broadcast a *complaint* against P_i. If there are at least $t+1$ complaints, then clearly P_i must be corrupt since with at most t malicious players, there can be at most t false complaints. Also, if P_i is corrupt, there will always be enough honest players to generate at least $t+1$ complaints and P_i will surely be disqualified in this case. Once eliminated, P_i is removed from $\mathcal{P}'$ and is replaced with a new player, thus resulting in a new signature. As long as at most t players are corrupted, there will always be enough honest players to create a valid signature.

4 Security Model

4.1 Security Definitions

We define two assumptions that we will use in our proof. The first is the one-more-discrete-logarithm assumption introduced by Bellare et al. [2]

Definition 7 (One-More-Discrete-Logarithm Assumption). We let $p = 2q+1$ be a prime where q is a random k-bit prime, and let g be a generator for a subgroup of $\mathbb{Z}_p^*$ with order q. We let $n : \mathbb{N} \rightarrow \mathbb{N}$ be a function of k. Now let $\{x_1, x_2, \ldots, x_{n(k)}, x_{n(k)+1}\}$ be elements of $\mathbb{Z}_q$ chosen uniformly at random, and for each $i \in \{1, 2, \ldots, n(k)+1\}$, define $z_i = g^{x_i} \pmod p$. Now let the adversary $\mathcal{A}$ have access to a discrete log oracle DLog such that if $x \in \mathbb{Z}_q$, $z = g^x \pmod p$, then $\mathsf{DLog}(g, z) = x$. In the one-more discrete-logarithm problem [2], $\mathcal{A}^{\mathsf{DLog}}$ is given $\{z_1, z_2, \ldots, z_{n(k)+1}\}$ and must output $\{x_1, x_2, \ldots, x_{n(k)+1}\}$ by querying DLog at most $n(k)$ times. The assumption is $\Pr[\mathcal{A}^{\mathsf{DLog}}(g, z_1, z_2, \ldots, z_{n(k)+1}) = (x_1, x_2, \ldots, x_{n(k)+1})] \leq \eta(k)$, where η is a negligible function.

We define a similar assumption that is related to finding collisions in a chameleon hash function. We will use this assumption to show our new scheme is secure. In Sect. 5.2, we show that this assumption is implied by the one-more-discrete-logarithm assumption for the chameleon hash function we use.

Definition 8 (One-More-R Assumption). As above, we let g be a generator for a subgroup of $\mathbb{Z}_p^*$ with order q, a k-bit prime. In addition, we let k' be randomly chosen from $\mathbb{Z}_q$ and let $h = g^{k'}$. We let $n : \mathbb{N} \rightarrow \mathbb{N}$ be a function of k. Now let $\{v_1, v_2, \ldots, v_{n(k)}, v_{n(k)+1}\}$ be randomly chosen elements in the range of $CH_{HK}(\cdot)$. Now we give the adversary $\mathcal{A}$ access to a $\mathsf{Get\text{-}An\text{-}R}(v, m)$ oracle, such that if $r = \mathsf{Get\text{-}An\text{-}R}(v, m)$, then $CH_{HK}(m, r) = v$. In the One-More-R problem, $\mathcal{A}^{\mathsf{Get\text{-}An\text{-}R}}$ is given $\{v_1, v_2, \ldots, v_{n(k)+1}\}$ and with at most $n(k)$ queries to $\mathsf{Get\text{-}An\text{-}R}$, must output $\{(m_1, r_1), (m_2, r_2), \ldots, (m_{n(k)+1}, r_{n(k)+1})\}$ such that $v_i = CH_{HK}(m_i, r_i)$. The assumption is that $\Pr[\mathcal{A}^{\mathsf{Get\text{-}An\text{-}R}}(g, h, v_1, v_2, \ldots, v_{n(k)+1}) = \{(m_1, r_1), (m_2, r_2), \ldots, (m_{n(k)+1}, r_{n(k)+1})\}] \leq \eta(k)$, where η is a negligible function.

4.2 Adversarial Model

We assume that there is a static adversary $\mathcal{A}$ that corrupts some subset of the players in $\mathcal{P}$ before beginning the threshold signature scheme. Furthermore, we can analyze two different types of static adversaries: one that compromises before the off-line phase and the other compromises after the off-line phase terminates. We assume the former case in our proof of existential unforgeability. As for the communication model, we assume that all players are connected by secure point-to-point channels. Furthermore, we will assume a partially synchronous communication model during the key generation and off-line phases for the purpose of using the DKG protocol of Gennaro et al. [8].

5 Proof of Security

5.1 Robustness

Theorem 1. *Suppose that an adversary corrupts at most $t < \frac{n}{3}$ players. Then, our on-line/off-line threshold signature scheme $\mathcal{TS}^{\mathrm{On/Off}}$ is robust.*

Proof. We need to show completeness, soundness, and zero knowledge simulatability of the signature share verification protocol when verifying $\textsf{col-2}_i$ from player $P_i \in \mathcal{P}'$.

- **Completeness**: An honest player $P_i \in \mathcal{P}'$ should convince any verifier that the protocol was followed with high probability. In fact, if the signature share verification protocol is correctly followed, then the verifier will accept with probability 1.
- **Soundness**: No corrupted player $P_i \in \mathcal{P}'$ should be able to fool any verifier into accepting incorrect shares with high probability. Using the definitions for e, d, and f from Sect. 3.5, we require that both

$$\begin{array}{l} g^v d^{-c} \equiv g^r \pmod{p} \\ e^v f^{-c} \equiv e^r \pmod{p} \end{array} .$$

  Therefore, $g^v d^{-c} \equiv g^{v - y_i c} \equiv g^r \pmod{p}$ if and only if $v \equiv y_i c + r \pmod{p'}$. In addition, $e^v f^{-c} \equiv g^{m_i v} g^{(\textsf{col-2}_i - z_i)(-c)} \equiv (g^{m_i})^r \pmod{p}$, which implies that $m_i v - c(\textsf{col-2}_i - z_i) \equiv m_i r \pmod{p'}$. By using $e^v f^{-c} \equiv e^r \pmod{p}$ from above, we see that $m_i y_i c \equiv c(\textsf{col-2}_i - z_i) \pmod{p'}$. If $c \not\equiv 0 \pmod{p'}$, then clearly $\textsf{col-2}_i$ is the correct share. If $c \equiv 0 \pmod{p'}$, then $\textsf{col-2}_i$ may be incorrect. By the Discrete Logarithm Assumption, no probabilistic polynomial time adversary can produce such a v with non-negligible probability.
- **Zero Knowledge Simulatability**: No cheating verifier should learn anything useful after running the protocol. We can easily construct a simulator S which simulates the view of the verifier when verifying P_i's $\textsf{col-2}_i$. To do so, S selects c and v uniformly at random and fixes $H(g, d, e, f, g^v d^{-c}, e^v f^{-c})$ to be c, since we are working in the Random Oracle model. Thus, S has recreated the view of the verifier without knowing P_i's secret key share y_i, so the signature share verification protocol has zero knowledge.

As a result, our on-line/off-line threshold signature scheme is robust. We sketch an alternative approach without random oracles in Sect. 7. $\square$

5.2 Existential Unforgeability

The proof of existential unforgeability will be in a similar style to the proof in Shamir and Tauman [16]. First we make use of the following Lemma to show that our One-More-R assumption is implied by a standard assumption:

Lemma 1. *Suppose that there exists an adversary $\mathcal{B}$ that breaks the One-More-R assumption for the discrete logarithm chameleon hash with advantage ε. We show how to construct an algorithm $\mathcal{A}$ that breaks the One-More-Discrete-Log assumption with advantage greater than ε.*

Proof. We let $\mathcal{A}$ respond to $\mathcal{B}$'s queries in the One-More-R problem. $\mathcal{A}$ is given as input g and $\{z_1, z_2, \ldots, z_{n(k)+1}\}$. Let $\mathcal{A}$ be described as follows:

1. Pick y uniformly at random in $\mathbb{Z}_{p'}$.
2. Let $h = g^y$, and initialize $\mathcal{B}$ with g and h.
3. For $1 \leq i \leq n(k)+1$, pick m_i uniformly in $\mathbb{Z}_{p'}$ and let $v_i = z_i h^{m_i}$.
4. Send $\mathcal{B}$ the set $\{v_1, v_2, \ldots, v_{n(k)+1}\}$.
5. Whenever $\mathcal{B}$ makes a $\mathsf{Get\text{-}An\text{-}R}(v, m)$ query, receive $t = \mathsf{DLog}(g, v)$. Return the value $t - ym$ to $\mathcal{B}$.
6. If $\mathcal{B}$ successfully outputs $\{(m'_1, r'_1), (m'_2, r'_2), \ldots, (m'_{n(k)+1}, r'_{n(k)+1})\}$ where $CH_{HK}(m'_i, r'_i) = v_i$ for all i, $\mathcal{A}$ returns $\{x_1, x_2, \ldots, x_{n(k)+1}\}$ where $x_i = r'_i + y(m'_i - m_i)$. Otherwise, abort.

Clearly, we have $\varepsilon < \mathsf{Adv}\ \mathcal{B} \leq \mathsf{Adv}\ \mathcal{A}$. □

Using the One-More-R assumption, we can prove that our on-line/off-line threshold signature scheme is secure against adaptive chosen message attack.

Theorem 2. *Let $\mathcal{TS}$ = ($\mathsf{Thresh\text{-}Key\text{-}Gen}$, $\mathsf{Thresh\text{-}Sig}$, Ver) be a given simulatable threshold signature scheme. Then we let $\mathcal{TS}^{\mathrm{On/Off}}$ = ($\mathsf{On/Off\text{-}Thresh\text{-}Key\text{-}Gen}$, $\mathsf{Thresh\text{-}Sig\text{-}Off\text{-}line}$, $\mathsf{Thresh\text{-}Sig\text{-}On\text{-}line}$, Ver) be the resulting On-line/Off-line Threshold Signature scheme. If $\mathcal{TS}^{\mathrm{On/Off}}$ is existentially forgeable by an q-adaptive chosen message attack with success probability ε, then one of the following must hold:*

1. *There exists a probabilistic algorithm that breaks either the One-More-R assumption or the collision resistance of CH_{HK} with probability at least $\frac{\varepsilon}{2}$.*
2. *The underlying threshold signature scheme $\mathcal{TS}$ is existentially forgeable by a q-random message attack with probability at least $\frac{\varepsilon}{2}$.*

Proof. Suppose that an adversary $\mathcal{A}$ forges a signature in the $\mathcal{TS}^{\mathrm{On/Off}}$ scheme with a q-chosen message attack with probability ε. Now let $\{m_1, m_2, \ldots, m_q\}$ be the q messages chosen by $\mathcal{A}$ to be signed by the $\mathcal{TS}^{\mathrm{On/Off}}$ scheme. Let $\{(\sigma_1, m_1, r_1), \ldots, (\sigma_q, m_q, r_q)\}$ be the signatures produced in this fashion by the $\mathcal{TS}^{\mathrm{On/Off}}$ scheme. Then $\mathcal{A}$ outputs a signature forgery (σ, m, r) such that $V_{PK}(CH_{HK}(m, r), \sigma) = \mathsf{Accept}$ and $m \neq m_i$ for all i, with probability ε. Moreover, either there exists an i such that $CH_{HK}(m_i, r_i) = CH_{HK}(m, r)$ or there does not exist such an i. One of these cases occurs with probability at least $\frac{\varepsilon}{2}$.

If the first case holds with probability at least $\frac{\varepsilon}{2}$, then we define a simulator $\mathcal{S}$ that breaks the One-More-R assumption. $\mathcal{S}$ is given as input the public bases g and h, as well as the set of challenges $\{v_1, v_2, \ldots, v_{n(k)+1}\}$.

$\mathcal{S}$ simulates the $\mathsf{On/Off\text{-}Thresh\text{-}Key\text{-}Gen}$ phase with $\mathcal{A}$. When the simulation gets to the point where h is to be generated by using the DKG protocol, $\mathcal{S}$ uses $\mathsf{SIM}^{\mathrm{DKG}}(h)$, the DKG simulator, to "fix" the result of the DKG run to be h.

On the i^{th} run of the $\mathsf{Thresh\text{-}Sig\text{-}Off\text{-}line}$ phase, $\mathcal{S}$ simulates the phase as normal. However, when it reaches the point where h^m is to be generated using the DKG protocol, it uses $\mathsf{SIM}^{\mathrm{DKG}}(v_i g^{-r})$ to fix the value of h^m so that the resulting chameleon hash $g^r h^m$ equals the given v_i value. $\mathcal{S}$ then simulates the rest of the phase as normal.

On the j^{th} run of the $\mathsf{Thresh\text{-}Sig\text{-}On\text{-}line}$ phase, with input m'_j specified by $\mathcal{A}$, $\mathcal{S}$ simulates the phase as normal. Suppose that the players involved are $\mathcal{P}' \subset \mathcal{P}$. Of the players in $\mathcal{P}'$, without loss of generality let $\mathcal{P}_{\mathcal{A}} = \{P_1, P_2, \ldots, P_t\} \subset \mathcal{P}'$ be the players controlled by the adversary $\mathcal{A}$. Since $\mathcal{S}$ "controls" more than t players, it is able to reconstruct the values of r_i, y_i, m_i, and z_i for all $P_i \in \mathcal{P}_{\mathcal{A}}$ from its own shares, since all were generated by the DKG protocol. Hence $\mathcal{S}$ is able to recover $\mathsf{col\text{-}1}_i$ and $\mathsf{col\text{-}2}_i$ for all $P_i \in \mathcal{P}_{\mathcal{A}}$. Now $\mathcal{S}$ fixes $P_l \in \mathcal{P}' \setminus \mathcal{P}_{\mathcal{A}}$. For each $P_i \in \mathcal{P}' \setminus (\mathcal{P}_{\mathcal{A}} \cup \{P_l\})$, $\mathcal{S}$ picks $\mathsf{col\text{-}1}_i$ and $\mathsf{col\text{-}2}_i$ uniformly at random and broadcasts them. In addition, $\mathcal{S}$ queries the $\mathsf{Get\text{-}An\text{-}R}$ oracle on m'_j and v_j to receive r'_j. With this information $\mathcal{S}$ can simply fix the value of $(\mathsf{col\text{-}1}_l, \mathsf{col\text{-}2}_l)$ such that the interpolation of all the $\mathsf{col\text{-}1}_i + \mathsf{col\text{-}2}_i$ values comes out to be r'_j.

At the end, $\mathcal{A}$ produces (σ, m, r) such that $V_{PK}(CH_{HK}(m,r), \sigma) = \mathsf{Accept}$ and there exists an i such that $CH_{HK}(m,r) = v_i$. If v_i was not used by $\mathcal{S}$ in a run of $\mathsf{Thresh\text{-}Sig\text{-}On\text{-}line}$, then $\mathcal{S}$ has produced One-More-R value, namely r. On the other hand, if v_i was used by $\mathcal{S}$, then we have a collision with CH_{HK}.

If the second case holds with probability at least $\frac{\varepsilon}{2}$, then we define a simulator $\mathcal{S}$ that existentially forges a signature under a random message attack on the underlying threshold signature $\mathcal{TS}$. In addition, we let $\mathsf{SIM}_1^{\mathcal{TS}}$ and $\mathsf{SIM}_2^{\mathcal{TS}}$ be defined as in Definition 5.

$\mathcal{S}$ simulates the $\mathsf{On/Off\text{-}Thresh\text{-}Key\text{-}Gen}$ phase as normal, except during the execution of $\mathsf{Thresh\text{-}Key\text{-}Gen}$. In this case, $\mathcal{S}$ uses $\mathsf{SIM}_1^{\mathcal{TS}}$ to fix the public key for $\mathcal{TS}$ to be the public key for the signing oracle $\mathsf{Sig}_{\mathcal{TS}}$.

On the i^{th} run of the off-line phase, let $\mathcal{S}$ simulate it as normal, except for the computation of h^m and running $\mathsf{Thresh\text{-}Sig}$. Let $\mathcal{S}$ query $\mathsf{Sig}_{\mathcal{TS}}$, which outputs (v_i, σ_i), where v_i is chosen uniformly at random and $V_{PK}(v_i, \sigma_i) = \mathsf{Accept}$. Next, use $\mathsf{SIM}^{\mathrm{DKG}}(v_i g^{-r})$ to fix h^m. Finally, $\mathcal{S}$ then uses $\mathsf{SIM}_2^{\mathcal{TS}}$ to simulate a run of $\mathsf{Thresh\text{-}Sig}$ with $\mathcal{S}$ on input v_i, such that the output is fixed to σ_i. We can do this because our assumption is that $\mathsf{Thresh\text{-}Sig}$ is simulatable.

Each run of the on-line phase is simulated as normal by $\mathcal{S}$. At the end, $\mathcal{A}$ produces (σ, m, r) such that $V_{PK}(CH_{HK}(m,r), \sigma) = \mathsf{Accept}$ and for all i, $v_i \neq CH_{HK}(m,r)$. But in this case, $\mathcal{S}$ has forged a signature σ on a message $CH_{HK}(m,r)$ not queried to the signing oracle $\mathsf{Sig}_{\mathcal{TS}}$. □

From this, we can derive the following theorem:

Theorem 3. *Suppose that a static adversary corrupts at most t players before beginning the off-line phase. Then our on-line/off-line threshold signature scheme $\mathcal{TS}^{\mathrm{On/Off}}$ is existentially unforgeable against adaptive chosen message attacks* *assuming that the underlying threshold signature scheme $\mathcal{TS}$ is existentially unforgeable against random message attacks.*

6 Evaluation

We analyze the number of bit operations required by our scheme, as previously shown in Table 1. First, in our scheme, is the threshold key generation. The bit complexity of $\mathsf{Thresh\text{-}Key\text{-}Gen}$ for $\mathcal{TS}$, as well as generating a Germain prime is included in $K_{\text{On/Off}}$. Afterwards, we invoke the DKG protocol once, so this costs K_{DKG}. Thus, the key generation phase takes $K_{\text{On/Off}} + K_{\text{DKG}}$ bit operations.

Next, we analyze our off-line phase. First, we invoke the DKG protocol three times, so this gives $3K_{\text{DKG}}$. Next, we have g^r and h^m, so we multiply both terms to get $CH_{HK}(r, m)$. Moreover, a single multiplication requires $\mathcal{O}(k^2)$ bit operations over $\mathbb{Z}_p$. Finally, the signature stamp $S_{SK}(CH_{HK}(m, r))$ requires τ bit operations. Thus the off-line phases requires a total of $3K_{\text{DKG}} + \mathcal{O}(k^2) + \tau$ bit operations.

For our on-line complexity, we can separate a player's computational complexity for generating a signature share from the signature reconstruction complexity. Each player $P_i \in \mathcal{P}'$ performs two additions and two multiplications when computing $\mathsf{col\text{-}1}_i$ and $\mathsf{col\text{-}2}_i$. The on-line signature reconstruction requires computing $f_i(0)$, which is $2t$ multiplications, and this is done for all $P_i \in \mathcal{P}'$, so we have a total of $(2t+1)^2$ multiplications when we compute r'. Only addition of the $2(2t+1)$ shares as well as performing $2t$ subtractions when computing $f_i(0)$ is required giving a total of $2t(2t+1) + 2(2t+1) - 1 = 4t^2 + 6t + 1$ additions. Furthermore, each addition over $\mathbb{Z}_p$ requires $\mathcal{O}(k)$ bit operations. Already we see that the number of multiplications in the on-line phase is substantially fewer than k since the threshold t is quite small when compared to a k bit prime. If verification of the signature shares is required, then each share requires six modular exponentiations. A summary of the number of operations performed appears in Table 2.

We review the complexity of Shoup's RSA threshold signature scheme [17], which was also shown in Table 1. The key generation phase of Shoup's signature scheme requires a trusted party, but asymptotically the computation cost is the same as our distributed key generation. In Shoup's on-line phase, the reconstruction complexity, once again, can be separated from the share verification complexity. The reconstruction of the signature requires t modular exponentiations, $t - 1$ modular multiplications, and one invocation of the extended Eu-

Table 2. Our On-line Phase Computational Complexity

Our On-line Phase Complexity	Additions	Multiplications	Exponentiations
Player Signature Share	2	2	0
Signature Reconstruction	$4t^2 + 6t + 1$	$4t^2 + 4t + 1$	0
Signature Share Verification	0	3	6

clidean algorithm. Finally, verifying an individual signature share also requires six modular exponentiations and three modular multiplications. Although both threshold signature schemes have approximately the same signature share verification complexity, we have managed to avoid any modular exponentiations in the reconstruction complexity of our signature scheme.

7 Extensions

7.1 Using Merkle Trees for Batching

We explained earlier that computing a threshold signature when performing writes for small files in Pond [12] is expensive, while for large files, the time spent computing the threshold signature is negligible compared to the actual write. In the event that a threshold signature must be quickly computed on demand, our scheme immediately becomes attractive over other schemes. This is especially true for Pond when computing threshold signatures for small writes.

One way of improving performance is to batch messages using Merkle hash trees [11]. Instead of signing messages one by one, we wait for n messages to arrive and then build a Merkle tree over these messages. After computing the hash h_r on the root node, we sign h_r using our on-line/off-line threshold signature scheme. To verify the validity of any message m that was included in the batch, we hash m and using hashes from sibling nodes as well as other internal nodes, we can compute h_r. We then verify that h_r is in fact the hash that was signed. If there are a total of n messages and the batch size is B, then a total of $\lceil \frac{n}{B} \rceil$ signature stamps are needed. This approach does trade latency for throughput, and it depends on how much time can be spent waiting for messages to arrive on-line. In fact, Merkle trees for batching has been applied to Shoup's scheme in OceanStore in order to increase throughput for small updates [13].

7.2 Eliminating Random Oracles

By using the techniques in [4], which eliminates the random oracle from the verification step in Shoup's RSA threshold scheme, we can eliminate the random oracle H, but at the cost of including interaction.

8 Acknowledgments

We thank Lea Kissner, Emil Ong, Naveen Sastry, Umesh Shankar, and Hoeteck Wee for providing helpful feedback on an earlier draft of this work, as well as the anonymous referees for their helpful comments.

References

[1] Mary G. Baker, John H. Hartman, Michael D. Kupfer, Ken W. Shirriff, and John K. Ousterhout. Measurements of a Distributed File System. In *Proceedings of 13th ACM Symposium on Operating Systems Principles*, pages 198–212. Association for Computing Machinery SIGOPS, 1991.

[2] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The One-More-RSA-Inversion Problems and the Security of Chaum's Blind Signature Scheme. *Journal of Cryptology*, 16(3):185–215, 2003.

[3] David Chaum and Torben P. Pedersen. Wallet Databases with Observers. In *CRYPTO*, pages 89–105. Springer-Verlag, 1993.

[4] Ivan Damgård and Kasper Dupont. Efficient Threshold RSA Signatures with General Moduli and No Extra Assumptions. In *Public Key Cryptography*, pages 346–361. Springer-Verlag, 2005.

[5] Yvo Desmedt and Yair Frankel. Threshold Cryptosystems. In *CRYPTO*, pages 307–315. Springer-Verlag, 1989.

[6] Shimon Even, Oded Goldreich, and Silvio Micali. On-Line/Off-Line Digital Schemes. In *CRYPTO*, pages 263–275. Springer-Verlag, 1990.

[7] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Robust Threshold DSS Signatures. *Inf. Comput.*, 164(1):54–84, 2001.

[8] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure Distributed Key Generation for Discrete Logarithm Cryptosystems. 2003. Updated from *EUROCRYPT*, 1999.

[9] Hugo Krawczyk and Tal Rabin. Chameleon Signatures. In *Proceedings of the Network and Distributed System Security Symposium*, pages 143–154, 2000.

[10] John Kubiatowicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishna Gummadi, Sean Rhea, Hakim Weatherspoon, Westley Weimer, Chris Wells, and Ben Zhao. OceanStore: An Architecture for Global-Scale Persistent Storage. In *Proceedings of ACM Architectural Support for Programming Languages and Operating Systems*, Novemeber 2000.

[11] Ralph Merkle. Protocols for Public Key Cryptosystems. In *IEEE Symposium on Security and Privacy*, pages 122–134, April 1980.

[12] Sean Rhea, Patrick Eaton, Dennis Geels, Hakim Weatherspoon, Ben Zhao, and John Kubiatowicz. Pond: The OceanStore Prototype. In *Proceedings of the Conference on File and Storage Technologies*. USENIX, 2003.

[13] Sean Rhea and John Kubiatowicz. The OceanStore Write Path. http://roc.cs.berkeley.edu/retreats/summer_02/slides/srhea.pdf, June 2002.

[14] Mendel Rosenblum and John K. Ousterhout. The Design and Implementation of a Log-Structured File System. In *ACM Transactions on Computer Systems*, volume 10, pages 26–52, February 1992.

[15] Chris Ruemmler and John Wilkes. UNIX Disk Access Patterns. In *USENIX Winter 1993 Conference Proceedings*, January 1993.

[16] Adi Shamir and Yael Tauman. Improved Online/Offline Signature Schemes. In *CRYPTO*, pages 355–367. Springer-Verlag, 2001.

[17] Victor Shoup. Practical Threshold Signatures. In *EUROCRYPT*, pages 207–220, 2000.

[18] Zhiyong Xu, Yingwu Zhu, Rui Min, and Yiming Hu. Achieving Better Load Balance in Distributed Storage System. In *International Conference on Parallel and Distributed Processing Techniques and Applications*, June 2002.

[19] Lidong Zhou, Fred B. Schneider, and Robbert van Renesse. COCA: A Secure Distributed Online Certification Authority. *ACM Trans. Computer Systems*, 20(4):329–368, 2002.