

Evaluating Security of Voting Schemes in the Universal Composability Framework

Jens Groth^{1,2}

¹ BRICS^{***}, University of Aarhus, IT-parken Åbogade 34, 8000 Århus C, Denmark

² Cryptomathic A/S[†], Jægergårdsgade 118, 8000 Århus C, Denmark
jg@brics.dk

Abstract. In the literature, voting protocols are considered secure if they satisfy requirements such as privacy, accuracy, robustness, etc. It can be time consuming to evaluate a voting protocol with respect to all these requirements and it is not clear that the list of known requirements is complete. Perhaps because of this many papers on electronic voting do not offer any security proof at all.

As a solution to this, we suggest evaluating voting schemes in the universal composability framework. We investigate the popular class of voting schemes based on homomorphic threshold encryption. It turns out that schemes in this class realize an ideal voting functionality that takes the votes as input and outputs the result. This ideal functionality corresponds closely to the well-known ballot box model used today in manual voting. Security properties such as privacy, accuracy and robustness now follow as easy corollaries. We note that some security requirements, for instance incoercibility, are not addressed by our solution.

Security holds in the random oracle model against a non-adaptive adversary. We show with a concrete example that the schemes are not secure against adaptive adversaries. We proceed to sketch how to make them secure against adaptive adversaries in the erasure model with virtually no loss of efficiency. We also briefly hint at how to achieve security against adaptive adversaries in the erasure-free model.

Keywords: Voting, homomorphic threshold encryption, universal composability.

1 Introduction

We consider the security of voting protocols. As time has progressed, more and more security requirements have been published in the literature. Examples of such requirements are privacy, accuracy, fairness, robustness, universal verifiability, incoercibility and receipt-freeness [1,2]. With this growing list of requirements, designers of voting protocols face two problems: if they do not know the

^{***} Basic Research in Computer Science (www.brics.dk),
funded by the Danish National Research Foundation.

[†] www.cryptomathic.com

literature well they may miss a security requirement, and even if they do cover all known requirements this does not guarantee that new yet to be discovered requirements are satisfied by their voting scheme.

To partially solve these problems we suggest evaluating voting schemes in the universal composability (UC) framework of Canetti [3]. In the UC framework, an execution of a multi-party computation protocol is compared to an execution where a trusted ideal functionality handles the data and produces the output. A protocol is said to be secure if an adversary operating in a real-life model can be simulated in the ideal process model with the ideal functionality. In the case of voting, the ideal functionality takes as input the votes and outputs the result of the election. This ideal functionality corresponds to the old method of voters marking their choice on paper and putting the ballot in a box, which is opened once the election is over.

Let us see how this solution addresses some of the properties that we mentioned. Privacy and accuracy are automatically satisfied since it is a part of the model that input to the ideal functionality is not revealed in any way to the adversary and the ideal functionality does compute the result correctly. Robustness follows too; in the UC framework, we can corrupt parties and still have a good simulation in the ideal process. Fairness follows from the fact that the ideal functionality does not reveal any partial tallies during the process.

Our approach has the advantage that it covers many security requirements in a single security model. This simplifies security proofs since we only need to prove universal composability to prove all these specific security requirements. Our approach is also pro-active in the sense that using a general security model may mean that security requirements yet to be discovered are covered.

We do not claim to solve all security issues with this approach. In particular, universal composability of a voting scheme does not guarantee universal verifiability, incoercibility, receipt-freeness or protection against hackers. However, considering that many security issues are dealt with, and considering that the properties dealt with are often defined vaguely in papers dealing with voting schemes, we do find that this application of the UC framework is worthwhile to investigate.

The UC framework allows for modular composition. In short, this means that if we take a hybrid protocol, where part of the protocol is specified by an ideal functionality, then we can freely plug in any protocol that securely realizes this ideal functionality. Most voting schemes presented in the literature make shortcuts. They assume we have a broadcast channel with memory or an anonymous broadcast channel. Often they also assume some public keys are set up and assume that voters are registered without specifying how this is done. We take this approach too and assume these things are provided through an ideal functionality. The modular composition theorem of the UC framework tells us that this is a sound approach and that we may later insert any protocol that realizes this functionality to get a full-blown election protocol.

The specific class of voting protocols we look at in this paper is based on homomorphic threshold encryption. Many such schemes have been proposed in

the literature [4,5,6,7], only the first one of these offers a security proof. We prove that indeed these schemes realize an ideal voting functionality when the adversary is non-adaptive. The schemes are not secure against adaptive adversaries, however, we propose a simple modification to make them secure against adaptive adversaries in the erasure model. Furthermore, in the full paper [8] we suggest another modification based on Paillier encryption that gives security against adaptive adversaries in the erasure-free model.

2 Preliminaries

In this section, we present the various tools used in the class of voting schemes we intend to investigate. Before doing so, we offer a brief introduction to the idea behind this class of voting protocols.

The idea behind voting based on homomorphic encryption. We assume that the parties have access to a message board where everybody may post messages, everybody can read the messages posted on the message board, messages cannot be deleted, and all messages are authenticated, for instance with digital signatures. All communication will take place through this message board. Public data pertaining to the election is also posted on the message board. In particular, a public key pk for a cryptosystem is posted.

In this example, we assume for simplicity that the voters only have two choices. We encode “yes” as 1, while “no” is encoded as 0. A voter casts his vote by encrypting the vote and posting it on the message board, i.e., posting $E_{pk}(0)$ or $E_{pk}(1)$. Since the messages are authenticated, everybody can check whether an eligible voter cast the vote.

The cryptosystem should have a homomorphic property:

$$E_{pk}(m_1; r_1) \cdot E_{pk}(m_2; r_2) = E_{pk}(m_1 + m_2; r_1 + r_2).$$

When everybody has cast his vote we may therefore compute the product of all the ciphertexts and get an encryption of the number of “yes” votes.

Now the authorities must decrypt this ciphertext containing the result of the election. For this purpose, we assume that the cryptosystem has threshold decryption. The authorities each hold a secret share of the private key and if sufficiently many of them cooperate, they may decrypt the ciphertext. However, no coalition below the threshold value is able to decrypt any of the encrypted votes; this preserves privacy.

To prevent cheating we require that voters attach a non-interactive zero-knowledge proof that their ciphertext contains either 0 or 1. Otherwise, it would for instance be easy to cast 100 “yes”-votes by posting $E_{pk}(100)$. Standard non-interactive zero-knowledge proofs are too cumbersome to be used in practice, therefore this is typically done through a 3-move honest verifier zero-knowledge proof of correctness of a vote made non-interactive through the Fiat-Shamir heuristic.

In this section, we define Σ -protocols [9], the type of 3-move honest verifier zero-knowledge proofs that we use. We then note that these proofs in the random oracle model [10] can be transformed into non-interactive zero-knowledge proofs. We prove that in the random oracle model, we are dealing with a proof of knowledge, and for any prover there exists an emulator that also produces corresponding witnesses. This can be seen as a random oracle parallel of witness extended emulation as defined by Lindell [11]. Finally, we define the kind of homomorphic threshold encryption that we need.

Σ -protocols. A Σ -protocol is a special type of 3-move proof system. Say we have an element x and a language L . The prover P knows a witness w for $x \in L$ and wants to convince the verifier V that $x \in L$. We assume that both parties have access to a common reference string σ chosen with a suitable distribution. Some Σ -protocols do not require this, and in that case, we can of course just let σ be the empty string. The protocol goes like this: The prover sends an initial message a , receives a random challenge e and produces an answer z . V can now evaluate (σ, x, a, e, z) and decide whether to accept or reject the proof.

A Σ -protocol satisfies the following properties.

Completeness: Given (x, w) where w is a witness for $x \in L$ the prover will with overwhelming probability convince the verifier, if they both follow the protocol.

Special Soundness: There exists an efficient extractor that for any x given two acceptable proofs (a, e, z) and (a, e', z') with the same initial message but different challenges can compute a witness w for $x \in L$.

Special Honest Verifier Zero-Knowledge: There exists an efficient simulator that given x, e can create a “proof” (a, e, z) for $x \in L$, which is indistinguishable from a real proof with challenge e .

Non-interactive zero-knowledge proofs. Given access to a random oracle \mathcal{O} we can transform a Σ -protocol into a non-interactive proof system. To get the challenge e we form the initial message a , query \mathcal{O} with (x, a, aux) to get the challenge e and then compute the answer z .¹ The proof is then (a, z, aux) . To verify such a proof query \mathcal{O} with (x, a, aux) to get e and then run the verifier from the Σ -protocol.

Using standard techniques, we can prove that we get a non-interactive proof system with the following properties:

Completeness: Given (x, w) where w is a witness for $x \in L$ the verifier will accept if both the prover and the verifier follow the protocol.

Soundness: A dishonest prover cannot convince the verifier if $x \notin L$.

Zero-Knowledge: There exist a simulator $S^{\mathcal{O}}$ that given $x \in L$ can create a convincing proof (a, z, aux) indistinguishable from a real proof provided it has the following ability to modify the oracle. It may give (x, a, aux, e) to \mathcal{O}

¹ Typically, aux will contain the identity of the prover in order to prevent somebody else to duplicate the proof and claim to have made it.

and provided (x, a, aux) has not been queried before \mathcal{O} assigns the value e to be the answer to query (x, a, aux) .

The random oracle model is an idealization of the Fiat-Shamir heuristic, see [10]. In the Fiat-Shamir heuristic the prover uses a cryptographic hash-function to produce the challenge as $e = \text{hash}(x, a, aux)$.

Witness extended emulation in the random oracle model. A Σ -protocol is a proof of knowledge in the random oracle model. We formulate this in the form of witness extended emulation in the following way. Given some adversary that produces a vector of elements $x \in L$ and valid proofs of memberships of L , there is an emulator E_A that produces identically distributed elements together with the corresponding witnesses for memberships of L .

Theorem 1. *For all adversaries A there exists an expected polynomial time emulator E_A such that for all distinguishers D (even unbounded ones) we have*

$$\begin{aligned} & P[(\mathbf{x}, \mathbf{p}, s) \leftarrow A^{\mathcal{O}}(z) : (\mathbf{x}, \mathbf{p}) \in V \wedge D^{\mathcal{O}}(\mathbf{x}, \mathbf{p}, s, z) = 1] \\ & \approx P[(\mathbf{x}, \mathbf{p}, \mathbf{w}, s) \leftarrow E_A^{\mathcal{O}}(z) : (\mathbf{x}, \mathbf{p}) \in V \wedge (\mathbf{x}, \mathbf{w}) \in W \wedge D^{\mathcal{O}}(\mathbf{x}, \mathbf{p}, s, z) = 1], \end{aligned}$$

where z is some advice with length bounded by a polynomial in k , \mathcal{O} is a random oracle, V is the set of vector pairs (\mathbf{x}, \mathbf{p}) such that \mathbf{p} contains valid proofs for the elements in \mathbf{x} belonging to L , and W is the set of pairs (\mathbf{x}, \mathbf{w}) where \mathbf{w} contains witnesses for the elements of \mathbf{x} belonging to L .²

The theorem follows from standard rewinding techniques. A proof can be found in the full paper [8].

Homomorphic threshold encryption. A (t, n) -threshold cryptosystem is a public key cryptosystem where the secret key is shared between n authorities A_1, \dots, A_n . If t of them cooperate they may decrypt ciphertexts, but any group of less than t authorities cannot learn anything about the contents of a ciphertext.

We use a key generation algorithm K to generate the keys. In general, all elements of the cryptosystem, messages, randomness and ciphertexts belong to suitable groups. We write the ciphertext space with multiplicative notation and the other groups with additive notation. The key generation algorithm produces

² It is instructive to consider this theorem in connection with the cryptosystem TDH0 in [12]. TDH0 is a cryptosystem where a Σ -protocol made non-interactive with a random oracle is used to prove knowledge of the plaintext. Intuitively one might argue CCA2 security by saying that the adversary already knows the answer when submitting decryption requests. However, Gennaro and Shoup show that this argument fails since rewinding is used to get the plaintexts, and since decryption requests may depend on oracle queries made before several other oracle queries we risk an exponential blow-up when tracking back through the decryption requests. Our theorem does not solve this problem. What our theorem can be used to prove, however, is that TDH0 is non-malleable.

a public key pk which is used for encryption, secret keys sk_1, \dots, sk_n used for decryption, and verification keys vk_1, \dots, vk_n that are public and used for verifying that the authorities act according to the protocol.

Encryption works as usual. To decrypt a ciphertext the authorities use their secret keys to produce decryption shares. Given t decryption shares anybody can combine them to get the plaintext. The verification keys are used by the authorities to make a zero-knowledge proof that they have provided the correct decryption shares.

We require that the cryptosystem have the following properties.

Semantic security: The cryptosystem must be semantically secure.

Errorless decryption: With overwhelming probability, the key generation algorithm selects keys such that there is probability 1 for the decryption to yield the message encrypted.³

Homomorphicity: For all messages m_1, m_2 and randomizers r_1, r_2 we have $E_{pk}(m_1 + m_2; r_1 + r_2) = E_{pk}(m_1; r_1) \cdot E_{pk}(m_2; r_2)$.

Simulatability of decryption: There is an algorithm S that takes as input a ciphertext c , a message m and the secret shares of any group of $t - 1$ authorities and produces simulated decryption shares for all the remaining authorities that c decrypts to m . S must be such that even with knowledge of the corresponding $t - 1$ keys the simulated decryption shares are indistinguishable from real decryption shares.

3 Universal Composability

The universal composability framework is described in details in [3]. The main gist is to compare a real-life execution of a protocol with an ideal process. We say a real-life protocol π realizes an ideal functionality \mathcal{F} if an adversary \mathcal{A} in the real-life model cannot gain more than an adversary \mathcal{S} in the ideal process does. More precisely, we have an environment \mathcal{Z} that gives inputs to parties, sees outputs from parties and learns which parties are corrupted, and we say π securely realizes \mathcal{F} if \mathcal{Z} cannot distinguish the real-life protocol with adversary \mathcal{A} from the ideal process with simulator \mathcal{S} .

In the ideal process, the ideal functionality handles everything taking place in the protocol. The parties in the protocol hand their inputs from \mathcal{Z} directly and securely to \mathcal{F} . \mathcal{F} computes the parties outputs and sends it to them. When a party receives a message from \mathcal{F} , it outputs this message. \mathcal{S} is restricted to corrupting some of the parties and blocking messages from \mathcal{F} to the honest parties. On the other hand, in the real-life execution the parties carry out the protocol π to produce their outputs.

One main feature in this framework is security under modular composition. Let us say we have a protocol ρ that realizes the ideal functionality \mathcal{F} . Say that

³ Most known cryptosystems have this property. However, in the notion of deniable encryption [13] the goal is to make it possible to deny that a particular thing was encrypted by producing honest looking randomness for an entirely different plaintext.

ρ is used as a sub protocol in π and write this as π^ρ . We may then form the hybrid $\pi^{\mathcal{F}}$ where calls to ρ are replaced with calls to \mathcal{F} . It is a theorem that π^ρ securely realizes $\pi^{\mathcal{F}}$.

Key generation and message board hybrid model. We will take advantage of the modular composition theorem and work in a hybrid model where we assume we have protocols that realize the key generation and message board functionality described in Figure 1. For distributed key generation protocols refer to [14,15,16,17]. This enables us to concentrate on the voting protocol itself.

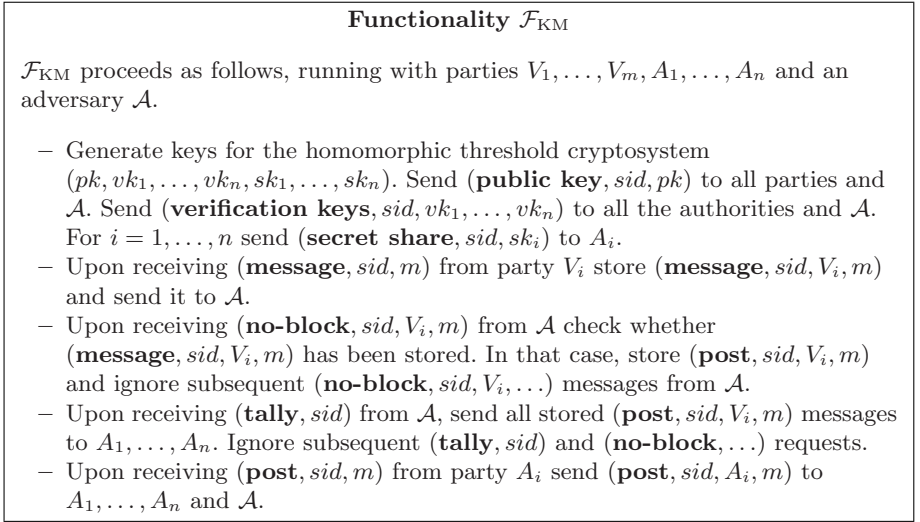


Fig. 1. The key generation and message board functionality, \mathcal{F}_{KM} .

We note that in \mathcal{F}_{KM} we allow \mathcal{A} to block voters' messages. This is to cover all the benign and malicious failures that may occur when voters try to cast their vote; everything from the Internet connection being unstable to an adversary deliberately cutting the cables to groups of voters with a particular political opinion. A typical requirement of a voting system is that it should be available, i.e., voters wanting to vote should have access to vote. This covers protecting against denial of service attacks, etc., but is not part of what the cryptographic protocol can accomplish. Therefore, we specifically allow the adversary to block votes. We quantify over all adversaries in the security proof, so in particular the security proof also holds for non-blocking adversaries that do not block messages, i.e., it holds for voting systems with the availability property. In contrast, for simplicity we do not allow the adversary to block inputs from the authorities. This choice is reasonable since any voting system must have appropriate back-up procedures to ensure that all authorities can communicate as needed.

Another remark pertains to resubmission of votes. Depending on the requirements, sometimes dictated by law, it may or may not be allowed for voters to change their votes. For simplicity, we treat the case where voters cannot change their mind, and therefore we only allow a single message not to be blocked. Security can be proved quite similarly in the case where we allow voters to change their mind.

Functionality $\mathcal{F}_{\text{VOTING}}$

$\mathcal{F}_{\text{VOTING}}$ proceeds as follows, running with parties $V_1, \dots, V_m, A_1, \dots, A_n$ and an adversary \mathcal{S} .

- Upon receiving **(vote, sid, V_i, v)** from V_i store it and send **(vote, sid, V_i)** to \mathcal{S} . Ignore future **(vote, sid, \dots)** messages from V_i .
- Upon receiving **(no-block, sid, V_i)** from \mathcal{S} check whether some **(vote, sid, V_i, v)** has been stored. In that case, add v to the result and ignore subsequent **(no-block, sid, V_i)** messages from \mathcal{S} .
- Upon receiving **(result, sid)** from \mathcal{S} compute the result and send **(result, sid, result)** to \mathcal{S} and A_1, \dots, A_n and halt.

Fig. 2. The voting functionality, $\mathcal{F}_{\text{VOTING}}$.

Voting protocol. Before describing the protocol that we use to realize the ideal voting functionality in Figure 2, we need to discuss how to encode the voters' choice as a plaintext to be encrypted. In [5,6,7] this is done by assigning each candidate a number $j \in \{0, \dots, L-1\}$ and encoding the candidate as M^j , where M is a strict upper bound on the number of votes any candidate can receive. Adding many such votes gives a result on the form $\sum_{j=0}^{L-1} v_j M^j$ where v_j is the number of votes on candidate number j . Votes and result can be embedded in a message space on the form \mathbb{Z}_N provided $N \geq M^L$. More generally we require that there is an encoding such that:

- Each valid vote v can be encoded as $\text{Encode}(v)$.
- The sum of the encodings yields an encoding of the result, $\text{Encode}(\text{result})$.
- It is possible to efficiently extract the result from an encoding.
- The encodings can be embedded in the message space of the cryptosystem.

We describe the voting protocol based on homomorphic threshold encryption in Figure 3. Examples of such voting protocols can be found in [4,5,6,7].

Ideal process adversary. To prove security of the voting protocol we need to provide an ideal process adversary \mathcal{S} that fares as well as \mathcal{A} does in the \mathcal{F}_{KM} -hybrid model. \mathcal{S} is described in Figure 4.

Theorem 2. *The voting protocol hybrid $\pi_{\text{VOTING}}^{\mathcal{F}_{\text{KM}}}$ securely realizes $\mathcal{F}_{\text{VOTING}}$ for the class of non-adaptive adversaries that corrupt less than t authorities.*

Voting Protocol $\pi_{\text{VOTING}}^{\mathcal{F}_{\text{KM}}}$

The voting protocol for voters V_1, \dots, V_m and authorities A_1, \dots, A_n with access to ideal functionality \mathcal{F}_{KM} and random oracle \mathcal{O} is as follows.

1. Invoke \mathcal{F}_{KM} to give each voter V_1, \dots, V_m the public key and give each authority A_1, \dots, A_n all the verification keys and its own secret key.
2. Each voter V_i with a public key pk on the incoming message tape and a valid vote v_i on the input tape computes $c_i \leftarrow E_{pk}(\text{Encode}(v_i))$. He creates a proof p_i for the vote being correct using a Σ -protocol with \mathcal{O} 's answer on (c_i, a_i, pk, sid, V_i) as the challenge e_i .
He sends (**message**, sid, c_i, p_i) to \mathcal{F}_{KM} .
3. Authority A_j with the public key and the verification keys on its tape and a secret share of the private key on its tape does the following. When receiving a bunch of broadcast votes it computes C as the product of all the votes with valid proofs. Then it computes the decryption share ds_j . It also forms a proof p_j for the decryption share being correct using the verification key vk_j . The challenge in this proof is computed with \mathcal{O} .
It sends (**post**, sid, ds_j, p_j) to \mathcal{F}_{KM} .
4. Each authority picks the first t decryption shares with valid proofs that it receives and computes the plaintext of C .
It interprets the plaintext as $\text{Encode}(\text{result})$ and outputs (**result**, sid, result).

Fig. 3. The voting protocol $\pi_{\text{VOTING}}^{\mathcal{F}_{\text{KM}}}$.

Proof. We will take a walk one step at a time from the \mathcal{F}_{KM} -hybrid model to the ideal process. In doing so we will use expected polynomial time algorithms and rewind the environment. This is all right as long as we do not do this in the \mathcal{F}_{KM} -hybrid model or the ideal process itself.

Exp₁. Define Exp_1 to be the following modification of the \mathcal{F}_{KM} -hybrid model. After \mathcal{A} has submitted the command (**tally**, sid) to \mathcal{F}_{KM} we use the honest authorities' secret shares to decrypt the encrypted votes with valid proofs sent by \mathcal{A} on behalf of corrupt voters. We look at the tapes of the honest voters and if they are not blocked by \mathcal{A} , we add their votes to the corrupt voters' votes. This gives us the result of the election.

By the simulation property of the threshold cryptosystem, we may now simulate the honest authorities' decryption shares such that they fit with the result. To do this simulation we do not need knowledge of the honest authorities' secret shares. Using our ability to control the random oracle, we may also simulate proofs of these decryption shares being correct.

$\text{HYB}_{\pi, \mathcal{Z}, \mathcal{A}}^{\mathcal{F}_{\text{KM}}} \approx P_1$. We define P_1 to be the probability of \mathcal{Z} outputting 1 in Exp_1 . It is not possible for \mathcal{Z} to distinguish whether it is running in the \mathcal{F}_{KM} -hybrid model or experiment Exp_1 . The result is the same in both cases and indistinguishability follows from the zero-knowledge property of the proofs and the simulation property of the threshold cryptosystem.

Ideal process adversary \mathcal{S}

\mathcal{S} operates in the ideal process with dummy voters $\tilde{V}_1, \dots, \tilde{V}_m$ and dummy authorities $\tilde{A}_1, \dots, \tilde{A}_n$. It has input z . It controls the random oracle \mathcal{O} in the sense that it may assign a response e to a query q . This means that it can simulate proofs.

\mathcal{S} runs a simulated \mathcal{F}_{KM} -hybrid execution with simulated adversary \mathcal{A} . We write V_1, \dots, V_m and A_1, \dots, A_n to denote simulated parties.

- \mathcal{S} forwards all messages between \mathcal{A} and \mathcal{Z} .
- \mathcal{S} simulates the invocation of \mathcal{F}_{KM} . Having done this it knows the secret shares of the private key of all the authorities, in other words \mathcal{S} may decrypt messages encrypted under the public key.
- Suppose \mathcal{A} on behalf of a corrupt V_i sends (**message**, sid, c_i, p_i) and sends (**no-block**, sid, V_i, c_i, p_i) to \mathcal{F}_{KM} . \mathcal{S} checks whether the proof is valid and in that case it decrypts c_i to get a vote v_i . It submits (**vote**, sid, V_i, v_i) to $\mathcal{F}_{\text{VOTING}}$ on behalf of \tilde{V}_i and sends (**no-block**, sid, V_i) to $\mathcal{F}_{\text{VOTING}}$.
- Upon receiving (**vote**, sid, V_i) from $\mathcal{F}_{\text{VOTING}}$ it knows that \tilde{V}_i got (**vote**, sid, V_i, v_i) as input from \mathcal{Z} . It does not know the actual vote v_i . As long as V_i has not received the public key for the election \mathcal{S} ignores the problem, but if V_i has or gets the public key for the election \mathcal{S} must simulate V_i trying to cast a vote. It forms $c_i = E_{pk}(0)$ and simulates a proof p_i for c_i containing a valid vote. It simulates V_i sending (**message**, sid, V_i, c_i, p_i) to \mathcal{F}_{KM} and sends the resulting (**message**, sid, V_i, c_i, p_i) to the copy of \mathcal{A} . If it later receives (**no-block**, sid, V_i, m) from \mathcal{A} , \mathcal{S} simulates \mathcal{F}_{KM} receiving this message, and it sends (**no-block**, sid, V_i) to $\mathcal{F}_{\text{VOTING}}$.
- Upon \mathcal{A} sending (**tally**, sid) to \mathcal{F}_{KM} , \mathcal{S} lets the simulated \mathcal{F}_{KM} send the list of stored messages (**post**, sid, V_i, c_i, p_i) to A_1, \dots, A_n . It sends (**tally**, sid) to $\mathcal{F}_{\text{VOTING}}$ and learns the result. Let C be the product of all the c_i 's. \mathcal{S} uses the simulation property of the threshold cryptosystem to simulate shares ds_j for the honest A_j 's such that C decrypts to the actual result. Furthermore, it also simulates proofs p_j of the shares being correct.
- After \mathcal{A} has delivered both the keys and the messages to honest A_j , \mathcal{S} simulates that A_j sends the decryption share (**post**, sid, A_j, ds_j, p_j) to \mathcal{F}_{KM} .
- When A_j has received both the public keys and t decryption shares, then \mathcal{S} delivers the (**result**, sid, result) message from $\mathcal{F}_{\text{VOTING}}$ to \tilde{A}_j .

Fig. 4. The ideal process adversary \mathcal{S} .

Exp₂. Define Exp_2 as the following modification of Exp_1 . We look at the execution in the interval between key generation having been done and \mathcal{A} not yet having submitted (**tally**, sid) to \mathcal{F}_{KM} . After the key generation, we may for each honest voter and each possible vote it can get as input pre-generate the (**message**, sid, c_i, p_i) message.

Let A be an algorithm that takes as input the tapes of \mathcal{A} , \mathcal{Z} and the pre-generated encrypted votes. It runs the entire execution in this interval, and in

the end, it outputs the views of \mathcal{A} and \mathcal{Z} . From the views, we may read off the states of \mathcal{A} and \mathcal{Z} , restart them, and continue the experiment.

According to Theorem 1 we may replace A with an expected polynomial time algorithm E_A that indistinguishably outputs the same as A , but in addition provides the witnesses for the proofs made by corrupt voters. These witnesses include the votes of these corrupt parties and therefore we do not need to decrypt anything with the honest authorities' secret shares of the private key.

$P_1 \approx P_2$. We define P_2 as the probability that \mathcal{Z} outputs 1 at the end of experiment Exp_2 . It follows from Theorem 1 that $P_1 \approx P_2$.

Exp₃. Define Exp_3 the following way. Instead of letting the honest voters encrypt their votes and proving in zero-knowledge that the ciphertexts contain correct votes, we let them encrypt 0 and simulate the proofs of correctness. For each possible vote that \mathcal{Z} could give to an honest voter V_i , we construct such a 0-vote and feed A with these ciphertexts and simulated proofs.

$P_2 \approx P_3$. Let P_3 be the probability that \mathcal{Z} outputs 1 after experiment Exp_3 . In Exp_3 , we still use the real votes to fit the result in the end, and we do not at any point use the honest authorities' shares of the private key. Therefore, by the semantic security of the cryptosystem, the result is the same and \mathcal{Z} cannot distinguish the two experiments. Neither does it allow us to distinguish the views of \mathcal{A} and \mathcal{Z} that A produces, so these transcripts must still look like correct views of \mathcal{A} and \mathcal{Z} acting according to their programs.

Exp₄. We define Exp_4 as a modification of Exp_3 where we go back to using decryption to get \mathcal{A} 's votes. Instead of using the votes supplied by E_A , we decrypt the corrupt voters' ciphertexts with valid proofs and use these votes. We may now replace E_A with A since we do not need the votes directly. By definition, A produces valid transcripts of how \mathcal{A} and \mathcal{Z} behave with these inputs and we may therefore replace A with the execution of \mathcal{A} and \mathcal{Z} .

$P_3 \approx P_4$. By Theorem 1 we may shift back from E_A to A without being able to tell the difference. Since A produces two good transcripts for how \mathcal{A} and \mathcal{Z} work we may now go back to using \mathcal{A} and \mathcal{Z} also in the interval between key generation and \mathcal{A} submitting (**tally**, *sid*) to \mathcal{F}_{KM} .

$P_4 \approx \text{IDEAL}_{\mathcal{F}_{\text{VOTING}}, \mathcal{Z}, \mathcal{S}}$. The ideal process and Exp_4 are actually the same experiment. In Exp_4 we submit 0-votes on behalf of honest parties and simulate the proofs, just as \mathcal{S} does. When \mathcal{A} submits (**vote**, *sid*, V_i , c_i , p_i) on behalf of an honest voter we check the proof and decrypt just as \mathcal{S} does. To create something that looks as decryption shares that produce the result we simulate this just as \mathcal{S} does. \square

Recycling keys. One could ask whether the keys can be reused for several elections. The security proof fails in this case for the same reasons as described in [12] and Footnote 2. The problem is that we can prove non-malleability of the cryptosystem used to encrypt votes but not prove security with respect to general adaptive chosen ciphertext attacks. If we use the same keys in several elections, we give the adversary access to a decryption of the ciphertexts containing the results and therefore an adaptive chosen ciphertext attack. While we see no way to use this attack in practice, we cannot guarantee security.

If we really want to use the keys for several elections that is possible though. We can simply demand that the voter makes a proof of knowledge where votes can be straight-line extracted. For instance, the voter can encrypt votes under a second public key and prove that this has been done correctly. Then votes may be extracted directly from this ciphertext and no rewinding is needed. The authorities tally the votes by stripping away the extra proof and ciphertext and carrying out the usual tallying procedure with the remaining ciphertext.

4 Adaptive Adversaries

An adaptive adversary is one that decides during the execution of the protocol which parties to corrupt. After corruption of a party, the adversary may learn some data from earlier computations. To guard against such problems we may specifically specify in protocols that parties should erase certain data. We call this the erasure model. Sometimes the more strict erasure-free security model is preferred. In this model, the party's entire computational history is revealed upon corruption.

The voting schemes are not adaptively secure. The schemes [4,5,6,7] are in fact not secure in the adaptive setting, even when we allow erasures. Let us sketch a counter-argument for the case of a yes/no election using the scheme in [4] with 2 voters, 3 authorities and a threshold $t = 2$. We refer the reader to [4] for a description of the scheme.

Consider an environment \mathcal{Z} and adversary \mathcal{A} , where \mathcal{A} forwards everything it sees to \mathcal{Z} and follows instructions from \mathcal{Z} on how to behave. \mathcal{Z} first asks \mathcal{A} to activate the key-generation step of \mathcal{F}_{KM} and to deliver all the keys to the relevant parties. Then \mathcal{Z} selects at random that all voters should vote yes or all voters should vote no. It lets the first voter post its vote and then it flips a coin to decide whether to block the second voter or not. If both voters were allowed to post their votes, \mathcal{Z} carries out the entire election according to the protocol. If only the first voter was allowed to post his vote, \mathcal{Z} lets \mathcal{A} activate A_1 to obtain its decryption share. Then it flips a coin and corrupts either A_2 or A_3 . From the secret share it obtains it may now compute the result of the election. If everything works out OK, \mathcal{Z} outputs 1. If we are operating in the real-life model everything will work out OK and \mathcal{Z} will output 1 with 100% probability.

To finish the argument we will show that any \mathcal{S} cannot make \mathcal{Z} accept with more than 50% probability. First, \mathcal{S} must provide public keys $g, h = g^s$ for an

ElGamal cryptosystem. Second it must provide verification keys $h_1 = g^{s_1}, h_2 = g^{s_2}, h_3 = g^{s_3}$ to the authorities. Here s, s_1, s_2, s_3 may or may not be known to \mathcal{S} and may or may not be chosen according to the protocol. Having given these keys to \mathcal{Z} \mathcal{S} must now produce the vote (x, y) for the first voter. At this point it cannot know the result since if it queries $\mathcal{F}_{\text{VOTING}}$ for the result, then \mathcal{Z} has 50% probability of letting the second voter vote, and then the result will be wrong and \mathcal{Z} will be able to distinguish. From now on, we look at the case where (x, y) has been produced without knowledge of the result, and where this is the only vote to be cast. \mathcal{S} must try to make it look like (x, y) decrypts to the result. First, it must produce a decryption share w_1 for the first authority. Then depending on \mathcal{Z} 's coin-flip, it must give either s_2 or s_3 to \mathcal{Z} according to which authority \mathcal{Z} decides to corrupt. To make \mathcal{Z} accept with more than 50% probability, \mathcal{S} must be able to make it look like (x, y) decrypts to the result in both cases. In other words, we have

$$G^{\text{result}} = y/w_1^{\lambda_{1,\{1,2\}}} x^{s_2 \lambda_{2,\{1,2\}}} = y/w_1^{\lambda_{1,\{1,3\}}} x^{s_3 \lambda_{3,\{1,3\}}},$$

where the Lagrange coefficients are $\lambda_{1,\{1,2\}} = 2, \lambda_{2,\{1,2\}} = -1, \lambda_{1,\{1,3\}} = 3/2, \lambda_{3,\{1,3\}} = -1/2$. This implies that we can compute $w_1 = x^{2s_2 - s_3}$ and $y = G^{\text{result}} x^{3s_2 - 2s_3}$. However, since (x, y) was chosen before the result was known to \mathcal{S} there is at least 50% probability that \mathcal{S} could not have done this. \mathcal{Z} only has 50% probability of outputting 1 in the ideal process and it can therefore distinguish.

Adaptive security in the erasure model. We can deal with an adversary that may adaptively corrupt voters quite easily. The voters simply erase the plaintext vote and the randomness after they have computed the encrypted vote. This way an adaptive adversary does not learn anything by corrupting a voter. We find the erasure model to be somewhat reasonable since a good voting system should specify that voters delete the randomness and the vote used in order to give some rudimentary receipt-freeness.

To guard against adversaries that adaptively corrupt authorities we can use techniques from [18,19,20]. Let us briefly sketch how to do this. All the homomorphic cryptosystems in [4,6,5,7] require that in the decryption process we raise the ciphertext C or part of the ciphertext to a secret value s . In the abovementioned schemes we share s using a polynomial f of degree $t - 1$, and give each authority a share $s_i = f(i)$. Lagrange interpolation can then be used to perform the decryption. As we saw before, this technique causes trouble in the adaptive setting. However, if we instead use a linear secret sharing of s , i.e., select s_1, \dots, s_{n-1} at random and $s_n = s - \sum_{i=1}^{n-1} s_i$, then we can cope with an adaptive adversary. To recover if an authority fails to provide its decryption share, we also use polynomials f_1, \dots, f_n of degree $t - 1$ to secret share s_1, \dots, s_n . I.e., $f_i(0) = s_i$ and $s_{i,j} = f_i(j)$. Authority j knows all the shares $\{s_{i,j}\}_{i=1,\dots,n}$. The verification keys now also include trapdoor commitments, for instance Pedersen commitments, to the $s_{i,j}$'s. In the simulation, we pick all the shares s_1, \dots, s_n at random. When the first honest authority is about to compute its share, it

computes the share such that it fits with the result and all the other authorities' shares, and it simulates a proof of correctness. The authorities have to go through a more complicated protocol to compute the result and anybody wishing to verify the result also has to do more work, but it is still well within what is practical. The voters do not pay any performance penalty when having to use this type of voting scheme instead of the original type of voting scheme, for them the protocol looks the same.

Adaptive security in the erasure-free model. To obtain a protocol security against adaptive adversaries in the erasure-free model we can use the UC threshold cryptosystem of Damgård and Nielsen [21]. One problem in this scheme is that they use the UC commitments of [22], which require that each voter receive an individual commitment key. [23] suggested to solve this problem using non-malleable commitments, and better efficiency can be obtained if we use simulation sound commitments [24]. We can combine the protocols with zero-knowledge proofs of the type in [7] to prove correctness of the votes. Making it all non-interactive using the Fiat-Shamir heuristic, we obtain a fairly efficient voting scheme, which is secure against adaptive adversaries in the erasure-free setting. More details on this scheme are offered in [8].

Acknowledgment. Thanks to Ivan Damgård for asking whether the schemes based on homomorphic threshold encryption are secure in the universal composability framework.

References

1. Burmester, M., Magkos, E.: Towards secure and practical e-elections in the new era. In Gritzalis, D., ed.: Secure Electronic Voting. Kluwer Academic Publishers (2003) 63–72
2. Lambrinouidakis, C., Gritzalis, D., Tsoumas, V., Karyda, M., Ikonomopoulos, S.: Secure electronic voting: The current landscape. In Gritzalis, D., ed.: Secure Electronic Voting. Kluwer Academic Publishers (2003) 101–122
3. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS 2001. (2001) 136–145 Full paper available at <http://eprint.iacr.org/2000/67>.
4. Cramer, R., Gennaro, R., Schoenmakers, B.: A secure and optimally efficient multi-authority election scheme. In: proceedings of EUROCRYPT '97, LNCS series, volume 1233. (1997) 103–118
5. Damgård, I., Jurik, M.J.: A generalisation, a simplification and some applications of paillier's probabilistic public-key system. In: 4th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2001, LNCS series, volume 1992. (2001)
6. Baudron, O., Fouque, P.A., Pointcheval, D., Poupard, G., Stern, J.: Practical multi-candidate election scheme. In: PODC '01. (2001) 274–283
7. Damgård, I., Groth, J., Salomonsen, G.: The theory and implementation of an electronic voting system. In Gritzalis, D., ed.: Secure Electronic Voting. Kluwer Academic Publishers (2003) 77–100

8. Groth, J.: Evaluating security of voting schemes in the universal composability framework. Cryptology ePrint Archive, Report 2002/002 (2002)
9. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: proceedings of CRYPTO '94, LNCS series, volume 893. (1994) 174–187
10. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security 1993. (1993) 62–73
11. Lindell, Y.: Parallel coin-tossing and constant round secure two-party computation. In: proceedings of CRYPTO '01, LNCS series, volume 2139. (2001) 408–432
12. Shoup, V., Gennaro, R.: Securing threshold cryptosystems against chosen ciphertext attack. *Journal of Cryptology* **15** (2002) 75–96
13. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable encryption. In: proceedings of CRYPTO '97, LNCS series, volume 1294. (1997) 90–104
14. Boneh, D., Franklin, M.K.: Efficient generation of shared rsa keys. In: proceedings of CRYPTO '97, LNCS series, volume 1294. (1997) 425–439
15. Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Secure distributed key generation for discrete-log based cryptosystems. In: proceedings of EUROCRYPT '99, LNCS series, volume 1592. (1999) 293–310
16. Fouque, P.A., Stern, J.: Fully distributed threshold rsa under standard assumptions. In: proceedings of ASIACRYPT '01, LNCS series, volume 2248. (2001) 310–330
17. Algesheimer, J., Camenisch, J., Shoup, V.: Efficient computation modulo a shared secret with application to the generation of shared safe-prime products. In: proceedings of CRYPTO '02, LNCS series, volume 2442. (2002) 417–432
18. Canetti, R., Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Adaptive security for threshold cryptosystems. In: proceedings of CRYPTO '99, LNCS series, volume 1666. (1999) 98–115
19. Jarecki, S., Lysyanskaya, A.: Adaptively secure threshold cryptography: Introducing concurrency, removing erasures. In: proceedings of EUROCRYPT '00, LNCS series, volume 1807. (2000) 221–242
20. Lysyanskaya, A., Peikert, C.: Adaptive security in the threshold setting: From cryptosystems to signature schemes. In: proceedings of ASIACRYPT '01, LNCS series, volume 2248. (2001) 331–350
21. Damgård, I., Nielsen, J.B.: Universally composable efficient multiparty computation from threshold homomorphic encryption. In: proceedings of CRYPTO '03, LNCS series, volume 2729. (2003) 247–264
22. Damgård, I., Nielsen, J.B.: Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In: proceedings of CRYPTO '02, LNCS series, volume 2442. (2002) 581–596
23. Damgård, I., Groth, J.: Non-interactive and reusable non-malleable commitment schemes. In: STOC '03. (2003) 426–437
24. Garay, J.A., MacKenzie, P.D., Yang, K.: Strengthening zero-knowledge protocols using signatures. In: proceedings of EUROCRYPT '03, LNCS series, volume 2656. (2003) 177–194 Full paper available at <http://eprint.iacr.org/2003/037>.