A practical and secure coercion-resistant scheme for remote elections
(Extended Abstract)

Roberto Araújo[1], Sébastien Foulle[2], and Jacques Traoré[2]

[1] TU-Darmstadt, Hochschulstrasse 10, 64289 Darmstadt, Germany
rsa@cdc.informatik.tu-darmstadt.de
[2] Orange Labs, 42 rue des Coutures, BP 6243, 14066 Caen Cedex, France
s.foulle.ext@rd.francetelecom.com, jacques.traore@orange-ftgroup.com

Abstract. Juels, Catalano, and Jakobsson (JCJ) proposed at WPES 2005 the first scheme that considers real-world threats and that is more realistic for remote elections. Their scheme, though, has quadratic work factor and thereby is not efficient for large scale elections. Based on the work of JCJ, Smith proposed an efficient scheme that has linear work factor. In this paper we first show that the Smith's scheme is insecure. Then we present a new coercion-resistant election scheme with linear work factor that overcomes this and other flaws of the Smith's proposal. Our solution is based on the group signature scheme of Camenisch and Lysyanskaya (Crypto 2004).

1 Introduction

Remote electronic elections may provide many benefits to democratic societies. It may increase elections turnouts, afford convenience to the voters, and reduce costs, for instance. The risks inherent in such elections, though, can discourage its use in real world. The main threat is that coercion and vote-selling can be easily explored by adversaries. Remote elections, thus, must have ways to mitigate such problems.

Most existing proposals for remote elections rely on the receipt-freeness requirement and on assumptions to deal with coercion and vote-selling. Preventing receipts to be made, though, is not enough to counter these problems as the voter can be observed while voting, for example. The assumptions claimed (e.g. the voter cannot give away her private key material), in addition, are not realistic for remote scenarios.

Recently, Juels, Catalano, and Jakobsson (JCJ) [6] introduced a more complete requirement for remote elections called coercion-resistance. This concept considers not only the receipt-freeness requirement, but also real world attacks related to coercion (and vote-selling). Coercion-resistance takes into account that an adversary can force the voter to abstain from voting, can obtain private information from the voter and vote on his behalf, or can force the voter to send a randomly formed ballot as her vote.

Dagstuhl Seminar Proceedings 07311
Frontiers of Electronic Voting
http://drops.dagstuhl.de/opus/volltexte/2008/1295

In order to fulfill this new requirement, JCJ introduced the first coercion-resistant election scheme. The scheme basically mitigates coercive attacks by allowing the voter to deceive the adversary about her true vote intention and to vote again afterwards. It, though, cannot be used in large elections as it requires quadratic work factor (in number of votes) to compute the voting results. Particularly, the scheme relies on an inefficient comparison mechanism to determine the results.

Based on JCJ solution, Smith [7] proposed an improved coercion-resistant scheme. The scheme replaces the inefficient comparison mechanism of JCJ by a new one that computes the voting results in linear time. Weber et al. [8], however, pointed out weaknesses on Smith's proposal and fixed them through of a new scheme.

In this work we first present a weakness of Smith's mechanism of comparison that makes it insecure. The problem is also relevant to the scheme of Weber et al. as it employs the ideas of Smith. We then introduce a new coercion-resistant scheme. Our solution is based on JCJ idea, but it has linear work factor and does not rely on comparisons to compute the voting results.

The paper is organized as follows: in Section 2 we give an overview of JCJ scheme and recall the comparison mechanism of Smith; also, we present the weakness of Smith's solution. In Section 3, we describe our proposal of coercion-resistant scheme. Finally, we conclude our work in Section 4.

2 The coercion-resistant scheme of JCJ and the comparison mechanism of Smith

2.1 The scheme of JCJ

The proposal of Juels, Catalano, and Jakobsson [6] relies essentially on a mechanism of authentication through anonymous credentials to overcome coercive attacks. Specially, the voter receives a valid credential (e.g. an alphanumeric string) in a secure way and uses it when she want to cast her valid vote. A voter over coercion, though, is able to make fake credentials and use them to cast fake votes. An adversary has no way to distinguish between a valid and a fake credential in the tallying phase.

The scheme considers a registration phase free of adversaries and a bulletin board communication model. Also, it employs non-interactive zero-knowledge proofs, a probabilistic threshold public-key cryptosystem, and universal verifiable mix nets as cryptographic tools. In particular, the solution uses the plaintext equivalence test [5] as comparison mechanism. This mechanism takes two ciphertexts as input and returns a bit indicating if the corresponding plaintexts are equal or not. The scheme is briefly described as follows:

In the registration phase, a trustworthy authority issues a unique valid credential for the voter and publishes a probabilistic encryption of this credential on the bulletin board. The encryptions published on the bulletin board forms the list $L1$. During the voting phase, the voter sends to the bulletin board (through

an anonymous channel) a triple containing her encrypted vote, her encrypted credential, and proofs that the vote is for a valid candidate and that the voter knows the vote and the credential encrypted.

At the end of the election day, the talliers verify all proofs posted and exclude triples with invalid proofs. They then apply the plaintext equivalence test to the remaining encrypted credentials to remove duplicated votes. After removing the duplicates keeping the last posted votes, the remaining pairs of encryptions (vote and credential) forms the list $L2$ and this list is sent to a mix net. The mix net returns $L2'$. Then, the list $L1$ is sent to a different mix net that returns $L1'$. Now, the plaintext equivalence test is used a second time to compare the list $L1'$ with the list $L2'$. A vote is removed if its encrypted credential in $L2'$ does not match with an element of $L1'$. Finally, the votes with valid credentials are decrypted by the talliers.

Although the JCJ scheme fulfills the coercion-resistance requirement, the comparisons involving the plaintext equivalence tests makes it inefficient for large scale elections. Let N be the number of voters and V be the number of posted votes, one has $V \geq N$ and the overhead to perform the tests is quadratic in V.

2.2 The Smith's comparison mechanism

Based on the JCJ proposal, Smith [7] introduced a more efficient coercion-resistant scheme. The solution substitutes the previous comparison mechanism of JCJ for a new one that computes the voting results in linear time. In addition, it includes an additional mix step in the tallying phase and uses timestamps. The later improvements, though, are not relevant as pointed out by Weber et al. [8].

The mechanism of Smith performs a global blind comparison of ciphertexts instead of employing the costly plaintext equivalence test. In order to do this, the method makes deterministic fingerprints from probabilistic encryptions. This way, the fingerprints can be compared through hash tables efficiently. The method is shortly described as follows:

Let k be an ElGamal private key shared among the talliers and corresponding to a public key published on a bulletin board, $E[\sigma]$ an ElGamal ciphertext of a credential σ that is made with talliers' public key, and j a secret key shared by the talliers that is used jointly with the product kj (also shared) to make the fingerprints. The talliers cooperatively computes $E[\sigma^j]$ from $E[\sigma]$ by using the ElGamal malleability and obtains σ^j after the decryption process. For the comparison, the talliers uses half of the bits of σ^j.

The Smith's comparison method is efficient. However, we noted [1] that it is not secure: an adversary can use the ElGamal malleability to determine whether a coerced voter gave him a valid or a fake credential. In order to show this, we consider the following scenario:

Suppose an adversary forces the voter to reveal her credential σ and publishes two votes $(E[C], E[\sigma], P)$ and $(E[C], E[\sigma^2], P')$ on the bulletin board, where

[1] This problem was also observed by Michael Clarkson et al. [3].

$E[C]$ are encryptions of the same candidate C and P, P' are the respective proofs. After publishing the mix net output and applying the Smith's method to obtain σ^j and σ^{2j} (this is similar to the last steps of the JCJ tallying phase), the adversary is able to recover the two votes by testing if one fingerprint is square of another. Thus, if the two votes were removed by the talliers, the coercer learns that σ is an invalid credential.

3 Our coercion-resistant scheme

As we presented in the last Section, the scheme of JCJ is inefficient for large scale elections. Also, we showed that the comparison mechanism of Smith is insecure. We now introduce a new coercion-resistant voting scheme that employs some of the JCJ ideas and that computes election results in linear time.

Our solution does not rely on comparisons to identify valid credentials. Instead, we employ a particular mathematical structure to make the credentials and use a function to identify them apart. The structure makes hard for a coercer or a dishonest voter to forge new valid credentials, even after having seen several valid ones.

The new scheme has the following advantages: its security can be proved, it is a practical linear scheme (in the number of votes posted by the voters), one cannot link the votes of a given voter in different elections, and the construction of the credentials and the verification of their validity can be shared among several authorities. Thus, a single corrupted authority cannot give valid credentials to an attacker or tell to a coercer whether a credential is valid or not.

The credentials we employ are based on the group signature scheme of Camenish et al. [2] (but without bilinear maps). A credential is composed of two parts: a short one r which must be kept secret, and a long one (a, b, c). The first part of the credential (i.e. r) has around twenty ascii characters (this corresponds to 160 bits, the actual secure size for the order of generic groups), so a small piece of paper and a pen are sufficient to write r down. The other part can be stored in a device or be even sent by email to the voter without compromising the credential security.

The cryptographic tools required to realize the new scheme are: non-interactive zero-knowledge proofs, a probabilistic threshold public-key cryptosystem, and a universal verifiable mix net. In addition, the scheme assumes a bulletin board communication model. We describe our proposal as follows:

In a setup phase, the keys of the authorities as well as the elections parameters are generated. Specially, let G be a cyclic group with prime order p and E be an ElGamal encryption algorithm with a public key (g, h) and a secret key (s) shared among the authorities. The authorities publish their public key and a randomly chosen value $m \neq 1$ on the bulletin board. We assume that the Decision Diffie-Hellman problem is hard in G (excluding therefore groups equipped with an efficient bilinear map).

Registration phase

In this phase, the voter obtains a unique and valid credential from the registration authorities. The authorities share the secret keys (x, y) and issue cooperatively to the voter her secret credential $\sigma = (r, a, b, c)$, where a is a random number in G (with $a \neq 1$), r is a random number in Z_p, $b = a^y$ and $c = a^{x+rxy}$. In addition, the authorities furnish the voter with a proof of well-formedness for σ. Note that if (r, a, b, c) is valid, then for all r the credential (r, a^l, b^l, c^l) is a valid one too. This property is used by the voter to change the values a, b, c after receiving them.

The LRSW assumption (see [2] for details) insures that even if an attacker has many genuine credentials (r_i, a_i, b_i, c_i), it is hard for him to forge a new and valid credential (r, a, b, c), with $r \neq r_i$ for all i. The security of our scheme relies heavily on this assumption, which is known to hold for generic groups. In addition, the voter cannot prove to anyone else whether (r, a, b, c) is a valid credential or not, under the DDH assumption. This way, a voter over coercion can make a fake r (and also make fakes a, b, c) to deceive an adversary who will not be able to distinguish between a fake and a valid r.

Voting phase

The voter casts her vote by sending $(E[C], a, E[a^r], E[a^{x+rxy}], m^r, P)$ through an anonymous channel to the bulletin board, where C is the chosen candidate, (r, a, a^r, a^{x+rxy}) is her credential, and P is a list of non-interactive zero-knowledge proofs which insure that this vote is well-formed, in particular the r in $E[a^r]$ and the r in m^r are the same. Recall from the previous paragraph that the values a and $c = a^{x+rxy}$ have been changed by the voter and are therefore different from the ones she received from the authorities. In addition, note that the voter does not need the part $b = a^y$ of her credential to make her vote.

The value m^r is used to detect duplicates and guarantees that only one vote will be counted per voter. Otherwise, a dishonest voter could vote several times without being detected.

Tallying phase

In order to tally the votes, the talliers first verify the values m^r to remove duplicates and check the proofs P. Votes with invalid proofs are excluded and only the last posted (duplicated) votes are considered based on the order of posting on the bulletin board. The votes that passed the verification have their values m^r and P deleted, and their second component (i.e. a) replaced by the ElGamal ciphertext $E[a] = (1, a)$. This way, only the values $E[C], E[a], E[a^r], E[a^{x+rxy}]$ are processed in the next step.

The talliers then send $E[C], E[a], E[a^r], E[a^{x+rxy}]$ to the mix net and publish the output on the bulletin board. Let the quadruple $(t, u, v, w) = (E[C]', E[a]', E[a^r]', E[a^{x+rxy}]')$ be the mix net output. For each quadruple the talliers choose a random string $\alpha \neq 0 \mod p$ and compute $(wu^{-x}v^{-xy})^\alpha$. The result is equal

to 1 if and only if the credential is a valid one. Note that if the credential is invalid, just computing $wu^{-x}v^{-xy}$ may give some information to an attacker, so the random exponent α is necessary.

Finally, the talliers cooperatively decrypt the first component of each quadruple (i.e. $E[C]$) with valid credential and count the votes.

4 Conclusion

The scheme introduced by Juels, Catalano, and Jakobsson (JCJ) considers realistic threats for remote elections and is more appropriate for such scenarios. Unfortunately their scheme is inefficient for large scale elections and the improvement proposed by Smith is insecure. Aiming at overcoming the drawbacks of these previous solutions, we have presented a new coercion-resistant scheme.

Our solution inherits some ideas from the JCJ scheme as the use of anonymous credentials. However, it does not rely on comparisons to identify valid credentials. Instead, we employ a special kind of credentials and use a secure function to identify them apart. Our credentials are based on the work of Camenish et al. [2]. We have another proposal based on Boneh et al.[1] protocol, but we leave it to a forthcoming paper.

The scheme that we have showed is efficient and secure. Although we have not proved these properties, they will be presented in a more complete paper.

References

1. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Franklin [4], pages 41–55.
2. Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Franklin [4], pages 56–72.
3. Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: A secure remote voting system. Technical Report TR2007-2081, Cornell University, May 2007.
4. Matthew K. Franklin, editor. *Advances in Cryptology - CRYPTO 2004, 24th Annual International CryptologyConference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*. Springer, 2004.
5. Markus Jakobsson and Ari Juels. Mix and match: Secure function evaluation via ciphertexts. In Tatsuaki Okamoto, editor, *ASIACRYPT*, volume 1976 of *Lecture Notes in Computer Science*, pages 162–177. Springer, 2000.
6. Ari Juels, Dario Catalano, and Markus Jakobsson. Coercion-resistant electronic elections (extended abstract). In *ACM Workshop On Privacy In The Electronic Society 2005 (WPES'05)*, pages 61–70, November 2005.
7. Warren D. Smith. New cryptographic voting scheme with best-known theoretical properties. In *Workshop on Frontiers in Electronic Elections (FEE 2005)*, Milan, Italy, September 2005.
8. Stefan G. Weber, Roberto Araújo, and Johannes Buchmann. On coercion-resistant electronic elections with linear work. In *2nd Workshop on Dependability and Security in e-Government (DeSeGov 2007) at 2nd Int. Conference on Availability, Reliability and Security (ARES'07)*, pages 908–916. IEEE Computer Society, 2007.