

On Some Incompatible Properties of Voting Schemes

Benoît Chevallier-Mames¹, Pierre-Alain Fouque², David Pointcheval²,
Julien Stern³, and Jacques Traoré⁴

¹ DCSSI – `Benoit.Chevallier-Mames@sgdn.pm.gouv.fr`

² ENS – CNRS – INRIA – `{Pierre-Alain.Fouque,David.Pointcheval}@ens.fr`

³ Cryptolog International – `Julien.Stern@cryptolog.com`

⁴ Orange Labs – France Telecom R&D – `Jacques.Traore@orange-ftgroup.com`

Abstract. In this paper, we study the problem of simultaneously achieving several security properties, for voting schemes, without non-standard assumptions. More specifically, we focus on the universal verifiability of the computation of the tally, on the unconditional privacy/anonymity of the votes, and on the receipt-freeness properties, for the most classical election processes. Under usual assumptions and efficiency requirements, we show that a voting system that wants to publish the final list of the voters who actually voted, and to compute the number of times each candidate has been chosen, we cannot achieve:

- universal verifiability of the tally (UV) and unconditional privacy of the votes (UP) simultaneously, unless *all* the registered voters actually vote;
- universal verifiability of the tally (UV) and receipt-freeness (RF), unless private channels are available between the voters and/or the voting authorities.

1 Introduction

1.1 Motivations

A huge number of properties for voting schemes have been proposed so far: and namely, *universal verifiability* (UV), the *unconditional privacy/anonymity* of the votes (UP), receipt-freeness (RF), and incoercibility.

Some properties seem quite important because usual systems and/or paper-based systems achieve them, and some other seem more theoretical because they are not (efficiently) satisfied in existing schemes: people expect much more from electronic voting schemes than from paper-based systems: the best example is the universal verifiability, which is definitely not satisfied with the paper-based voting systems, since one can supervise one place only. On the other hand, an attack on an internet-based vote could be at a very large scale and thus much more damaging.

Furthermore, some properties are easily satisfied by using physical assumptions such as voting booths, while they are difficult if one can vote from home:

this is the case of incoercibility. Since cryptography is usually very powerful and makes possible some paradoxical things, one is tempted to build a system that achieves as many properties as possible, with as few assumptions as possible. But what is actually achievable?

1.2 Contributions

In this paper, we address this question: can we build a voting system that simultaneously satisfies several properties, without non-standard assumptions (such as physical assumptions)? More precisely, we focus on the large class of election systems that simply consist in counting the number of times that each candidate has been chosen (whatever the constraints on the choices may be) and want to be able to compute the list of the voters who actually voted. Such election rules are used in many countries (such as in France). On the one hand we study the universal verifiability (UV) and the unconditional privacy of the votes (UP), which is sometimes replaced by the unconditional anonymity of the voters. On the other hand, we consider the universal verifiability (UV) and the receipt-freeness (RF). In both cases, we show that we cannot simultaneously achieve the two properties without strong extra assumptions, such as secure channel between the voters and/or the authorities, which is unrealistic for efficient and practical protocols.

The universal verifiability and the unconditional privacy can actually be simultaneously satisfied if *all* the registered voters *do* vote; similarly the universal verifiability and the receipt-freeness can be simultaneously achieved if the voting transcript of a voter *does not* depend on the voter's vote, his secret, some personal possible private/random value, and additional public data only. It is well-known that using multi-party computation techniques a strongly secure voting scheme can be built, that achieves all the above ideal properties, but using secure channels between the parties (the voters and/or the authorities): efficient voting schemes that guarantee receipt-freeness or incoercibility [2, 4, 13, 17, 18, 21] use such secure channels.

In the standard model we adopt below, we assume algorithmic assumptions only, but no secret channels nor physical assumptions such as tamper-resistant devices [18]. In addition, while studying the security properties of voting schemes, we try to explain why the traditional schemes, based on blind signatures, mixnets or homomorphic encryption, satisfy these properties or not.

Having a clear view of which sets of properties are achievable has a practical significance: one can easily conceive that the properties required for a national election or for an internal company board vote are different. For instance, the unconditional privacy (UP) of the vote will be important (if not required) for national elections, while the receipt-freeness (RF) will not be as critical as it may be difficult to buy votes on a very large scale without detection. For a board vote, a few number of voters typically have a very large number of shares, while the rest have a small number of shares. The major voters choices are often not private (let alone unconditionally private) because they can be inferred from the result of the vote. However, it may be tempting for a dishonest important

voter, which could already have 40% of the shares, to buy the missing 10% to safeguard a majority. The receipt-freeness property is therefore more critical in that case.

1.3 Organization

The paper is organized as follows: first, in section 2, we give formal definitions to the above UV, UP and RF security notions. Then, we show the incompatibility results in section 3.

1.4 Notation

We use the following notation in the rest of the paper:

- L represents the list of the registered voters,
- V_i is a voter, who casts his ballot,
- \mathbf{V} is the list of the voters, who cast their ballots,
- v_i is the vote of voter V_i ,
- \mathbf{v} is the set of votes,
- r_i is the random coins of voter V_i ,
- \mathbf{r} is the set of the random coins,
- B_i is the transcript of V_i (that is the interactions between voter V_i and the voting authority, assumed to be public),
- \mathbf{B} is the set of transcripts, also known as the bulletin-board,
- T is the tally of the vote (the vector of the number of times that each candidate has been chosen),
- w, w' will denote the witnesses in some \mathcal{NP} - relations R and R' ,
- f, f', f'', g and h will be some functions.

Since we won't assume any private channel, any interaction can be assumed public, and also through the authority, and then included in the public transcript available on the bulletin-board. Furthermore, for practical reasons, the vote-and-go approach is often preferable, which excludes any complex interaction, but with the authorities only.

2 Security Notions

In this section, we formally define the most usual security notions: universal verifiability, unconditional privacy, and receipt-freeness.

2.1 Universal Verifiability of the Tally

This security notion tries to prevent dishonest voting authorities from cheating during the computation of the tally.

For example, voting schemes using blind-signature [8, 16, 20] cannot achieve this property since the authority can add some ballots and bias the tally. On

the other hand, schemes using mix-nets [1, 9–12, 14, 19, 22] and/or homomorphic encryption [3, 6, 7] may provide it.

First, in order to universally check the validity and the correctness of a vote, one has to guarantee that a voter has not voted twice. Consequently, one needs to authenticate the transaction in some way. To this end, one needs to be able to verify both the link between the list of the registered voters L and the list of the transcripts \mathbf{B} (or the bulletin-board) in order to validate the vote, and the link between the bulletin-board and the computation of the tally T .

Definition 1 (Voting Scheme). *For a voting scheme to be practical and sound, it must hold the following properties.*

- Detection of individual fraud. *From a partial list of transcripts \mathbf{B} produced by $V_1, \dots, V_n \in L$, the voting authority should be able to determine whether a new transcript B produced by V_{n+1} is valid (well-formed and does not correspond to a double vote). More formally, there exists a boolean function f that can determine this fact,*

$$\begin{aligned} & \forall n, \forall V_1, \dots, V_n, V_{n+1} \in L, \\ & \forall \mathbf{B} \leftarrow V_1, \dots, V_n, B \leftarrow V_{n+1}, \\ & f(\mathbf{B}, B) = \left\{ \begin{array}{l} 0, \text{ if } V_{n+1} \in \{V_1, \dots, V_n\} \\ 1, \text{ if } V_{n+1} \notin \{V_1, \dots, V_n\} \end{array} \right\} \wedge B \text{ valid.} \end{aligned}$$

We thus denote by \mathcal{L} the language of the bulletin-boards \mathbf{B} which are iteratively valid.

- Computation of the tally. *From the transcripts, the voting authority should be able to compute the tally, that is a vector of the number of selections for each candidates: there exists an efficient function f' that, from the bulletin-board \mathbf{B} , outputs the tally T ,*

$$\forall \mathbf{B} \in \mathcal{L}, f'(\mathbf{B}) = \sum_i v_i = T.$$

- Computation of the list of the voters. *From the transcripts, the voting authority should be able to determine the list of the voters who actually casted their ballots: there exists an efficient function f'' that, from the bulletin-board \mathbf{B} , extracts the sub-list \mathbf{V} of the voters,*

$$\forall \mathbf{B} \in \mathcal{L}, f''(\mathbf{B}) = \mathbf{V}.$$

When one wants the universal verifiability, everybody should be able to check the correctness/validity of the votes and of the computation of the tally and the voters: the bulletin-board \mathbf{B} , the tally T and the list of the voters \mathbf{V} should rely in an \mathcal{NP} language \mathcal{L}' , defined by the relation R : there exists a witness w which allows an efficient verification. Furthermore, for any \mathbf{B} , the valid T and \mathbf{V} should be unique:

Definition 2 (Universal Verifiability (UV)). *Let R be the \mathcal{NP} -relation for the language \mathcal{L}' of the valid ballots and valid computation of the tally. A voting*

scheme achieves the universal verification property if only one value for the tally and the list of the voters can be accepted by the relation R , and the witness w can be easily computed from the bulletin-board \mathbf{B} using a function g :

$$\begin{aligned} \forall \mathbf{B} \in \mathcal{L}, \exists! (T, \mathbf{V}) \text{ s.t. } \exists w \text{ s.t. } R(\mathbf{B}, T, \mathbf{V}, w) = 1 \\ \forall \mathbf{B} \notin \mathcal{L}, \forall (T, \mathbf{V}, w) \ R(\mathbf{B}, T, \mathbf{V}, w) = 0 \\ \forall \mathbf{B} \in \mathcal{L} \ R(\mathbf{B}, f'(\mathbf{B}), f''(\mathbf{B}), g(\mathbf{B})) = 1. \end{aligned}$$

Note that g is a function private to the authorities, to compute a short string (the witness) that allows everybody to check the overall validity, granted the public relation R .

The functions f , f' , f'' and g may be keyed according to the system parameters: g is clearly private to the voting authority, while f and f'' may be public (which is the case in schemes based on homomorphic encryption). The function f' is likely to be private.

2.2 Unconditional Privacy

First, one should note that this notion can not be achieved in a very strong sense: if all voters vote identically, the tally reveals the vote of each voter. Consequently, privacy means that nobody should learn more information than what is leaked by the tally. By unconditional privacy, we thus mean that nobody should be able to learn any additional information even several centuries after the voting process.

In voting schemes based on homomorphic encryption [3, 6, 7] privacy relies on computational assumptions, and is thus not unconditional. When mix-nets are used, this is the same, since the latter applies on asymmetric encryptions of the votes. On the other hand, voting schemes based on blind signatures can achieve this strong security notion, but under the assumption of anonymous channels, which are usually obtained with asymmetric encryption: unconditional privacy vanishes!

Definition 3 (Unconditional Privacy (UP)). *A voting scheme achieves the unconditional privacy if*

$$\mathcal{D}(\mathbf{v} \mid T, \mathbf{B}) \stackrel{p,s}{\equiv} \mathcal{D}(\mathbf{v} \mid T).$$

This equation means that the distribution of the votes, given the bulletin-board and the tally T is the same as without any additional information to the tally. The distance between these two distributions can be perfect or statistical, hence the s and p . But we of course exclude any computational distance.

2.3 Receipt-Freeness

The receipt-freeness property means that a voter cannot produce a proof of his vote to a third party. In such a security notion, interactions with the third party

are allowed before and after the vote. Furthermore, if the vote is performed outside a booth, we can also assume that the third party has access to the channel between the voter and the voting authority: he has knowledge of the transcript, but also of all the information known to the voter, as well as the public information.

A receipt would thus be a proof of the vote v_i , by the voter V_i to a third party: a proof (a witness w') that shows that the bulletin-board contains the vote v_i for voter V_i . The proof must be sound, which means that several proofs are possible, but all for the same statement v_i for a given voter V_i :

Definition 4 (Receipt-Freeness). *A receipt is a witness w' which allows a third party to verify, in an unambiguous way, the vote of a voter $V_i \in \mathbf{V}$:*

$$\exists! v_i, \text{ s.t. } \exists w' \text{ s.t. } R'(\mathbf{B}, V_i, v_i, w') = 1.$$

A voting scheme achieves the receipt-freeness property if there is no such a relation R' , or the witness w' is hard to compute.

3 Incompatible Properties

In this section, we show that a voting scheme cannot provide

- the universal verifiability and the unconditional privacy of the votes, simultaneously, unless all the voters actually vote;
- the universal verifiability and the receipt-freeness, simultaneously, if the transcript of a voter depends on the voter, his vote, his own random, and public values only.

3.1 Universal Verifiability and Unconditional Privacy

Theorem 5. *In the standard model, it is impossible to build a voting scheme that simultaneously achieves the universal verifiability and the unconditional privacy unless all the voters actually vote.*

Proof. Assume we have a *universally verifiable* voting scheme. Then, we want to prove that the *unconditional privacy* cannot be achieved.

Because of the universal verifiability, there exists a public \mathcal{NP} -relation R such that $R(\mathbf{B}, T, \mathbf{V}, w) = 1$, where w is a witness, for a unique tally T and the unique list of voters. Because of the existence of f' , f'' and g , a powerful adversary can guess $\mathbf{V}' = f''(\mathbf{B}')$, $T = f'(\mathbf{B}')$ and $w = g(\mathbf{B}')$ for any $\mathbf{B}' \in \mathcal{L}$: excluding one transcript from \mathbf{B} to build \mathbf{B}' , this adversary can get the name of the excluded voter V' , and the new tally T' , which leaks the vote $v' = T - T'$ of the voter V' .

With an exhaustive search among all the sub-parts of \mathbf{B} , one can then get the vote of a specific voter. \square

This proof strongly relies on the latter sentence. And therefore, the contradiction comes from the above relation R that applies whatever the size of \mathbf{B} is, which allows us to exclude one transcript and use the universal-verifiability relation R .

If the transcripts of *all* the registered voters in L were required in R , the contradiction would not hold anymore, even if it is not clear whether a counter-example exists or not. Anyway, requiring all the registered voters to actually vote is not realistic. A denial of service would become very likely.

In [15], Kiayias and Yung propose a voting scheme in which the privacy is maintained in a distributed way among all the voters. There is no voting authority. They prove that the scheme provides the perfect ballot secrecy which does not correspond to our notion of unconditional privacy: it means that the security of a vote is guaranteed as long as the size of a coalition is not too large and of course according to the tally result and coalition votes. However, in their scheme, each ballot is encrypted using a public-key encryption scheme, that thus requires a computational assumption for the privacy.

In [5], Cramer *et al.* propose a voting scheme that guarantees the unconditional privacy, by using unconditionally secure homomorphic commitments, but only with respect to the voters, and not to the authorities, which would be able to open each individual vote if they all collude.

3.2 Universal Verifiability and Receipt-Freeness

Theorem 6. *Unless private channels are available, the universal verifiability and the receipt-freeness properties cannot be simultaneously achieved.*

Proof. Because of the universal verifiability, v_i is uniquely determined by B_i specific to the voter V_i . Since we exclude private channels, B_i can only be a function of V_i , his vote v_i , some input r_i private to V_i , and public data P_i : $B_i = h(V_i, v_i, r_i, P_i)$. Therefore, r_i is a good witness, and thus a receipt: the scheme is not receipt-free. \square

If the transcript is more intricate, and namely includes some private interactions between the voters and/or the authorities [13], then it may be possible to achieve the two properties simultaneously: B_i is no longer available to the third-party, and thus r_i is no longer a witness either. But such an assumption of private channels is not reasonable in practice.

Conclusion

As a conclusion, we have shown that voting systems with usual features cannot simultaneously achieve strong security notions: we cannot achieve simultaneously universal verifiability of the tally and unconditional privacy of the votes or receipt-freeness.

References

1. M. Abe and M. Ohkubo, A Length-Invariant Hybrid Mix, *Proceedings of Asiacrypt'01*, volume 1976 of LNCS, pages 178-191, Springer-Verlag, 2001.
2. R. Aditya, B. Lee, C. Boyd and E. Dawson, An efficient mixnet-based voting scheme providing receipt-freeness, *Proceedings of TrustBus'04*, volume 3184 of LNCS, pages 152-161, Springer-Verlag, 2004.
3. O. Baudron, P.-A. Fouque, D. Pointcheval, G. Poupard and J. Stern, Practical multi-candidate election system, *Proceedings of the 20th ACM Symposium on Principles of Distributed Computing*, pages 274-283, ACM Press, 2001.
4. J. Benaloh and D. Tuinstra, Receipt-free secret ballot elections, *Proceedings of STOC'94*, volume 1976 of LNCS, pages 544-553, 1994.
5. R. Cramer and M. Franklin and B. Schoenmackers and Moti Yung. Multi-Authority Secret-Ballot Elections with Linear Work. Eurocrypt '96, LNCS 1070, pp. 72-83, 1996.
6. R. Cramer, R. Gennaro and B. Schoenmakers, A Secure and Optimally Efficient Multi-Authority Election Scheme, *Proceedings of Eurocrypt'97*, volume 1233 of LNCS, pages 113-118, Springer-Verlag, 1997.
7. I. Damgard and M. Jurik, A Generalisation, a Simplification and Some Applications of Paillier's Probabilistic Public-Key System, *Proceedings of PKC'01*, volume 1992 of LNCS, pages 119-136, Springer-Verlag, 2001.
8. A. Fujioka, K. Ohta and T. Okamoto, A practical Secret Voting Scheme for Large Scale Elections, *Proceedings of Auscrypt'92*, volume 718 of LNCS, pages 248-259. Springer-Verlag, 1992.
9. J. Furukawa, Efficient, verifiable shuffle decryption and its requirement of unlinkability. *Proceedings of PKC'04*, volume 2947 of LNCS, pages 319-332, Springer-Verlag, 2004.
10. J. Furukawa and K. Sako, An Efficient Scheme for Proving a Shuffle. *Proceedings of Crypto'01*, volume 2139 of LNCS, pages 368-387, Springer-Verlag, 2001.
11. P. Golle, S. Zhong, D. Boneh, M. Jakobsson and A. Juels, Optimistic Mixing for Exit-Polls, *Proceedings of Asiacrypt'02*, volume 2501 of LNCS, pages 451-465, Springer-Verlag, 2002.
12. J. Groth, A verifiable secret shuffle of homomorphic encryptions. *Proceedings of PKC'03*, volume 2567 of LNCS, pages 145-160, Springer-Verlag, 2003.
13. M. Hirt and K. Sako. Efficient Receipt-Free Voting Based on Homomorphic Encryption. Eurocrypt '00, LNCS 1807, pp. 539-556, 2000.
14. M. Jakobsson, A. Juels, and R. Rivest, Making Mix-Nets Robust for Electronic Voting by Randomized Partial Checking, *Proceedings of the 11th Usenix Security Symposium, USENIX '02*, pages 339-353, 2002.
15. A. Kiayias and M. Yung. Self-tallying Elections and Perfect Ballot Secrecy. PKC 2002, LNCS 2274, pp. 141-158, 2002.
16. K. Kim, J. Kim, B. Lee and G. Ahn, Experimental Design of Worldwide Internet Voting System using PKI, *SSGRR2001*, L'Aquila, Italy, Aug. 6-10, 2001.
17. B. Lee, C. Boyd, E. Dawson, K. Kim, J. Yang and S. Yoo, Providing receipt-freeness in mixnet based voting protocols, *Proceedings of ICICS'03*, volume 2971 of LNCS, pages 245-258, Springer-Verlag, 2003.
18. B. Lee and K. Kim, Receipt-free electronic voting scheme with a tamper-resistant randomizer, *Proceedings of ICICS'02*, volume 2587 of LNCS, pages 389-406, Springer-Verlag, 2002.

19. A. Neff, A verifiable secret shuffle and its application to e-voting, *ACM CCCS 2001*, pages 116-125, ACM Press, 2001.
20. M. Ohkubo, F. Miura, M. Abe, A. Fujioka and T. Okamoto, An Improvement on a Practical Secret Voting Scheme, *Information Security'99*, volume 1729 of LNCS, pages 225-234, Springer-Verlag, 1999.
21. T. Okamoto, Receipt-free electronic voting schemes for large scale elections, *Workshop on Security Protocols'97*, volume 1361 of LNCS, pages 25-35, Springer-Verlag, 1998.
22. K. Peng, C. Boyd and E. Dawson, Simple and efficient shuffling with provable correctness and ZK privacy, *Proceedings of CRYPTO'05*, volume 3621 of LNCS, pages 188-204, Springer-Verlag, 2005.