SOME CONDITIONS FOR ELECTRONIC VOTING

McCarthy Smári, International Modern Media Institute, Iceland , smari@smarimccarthy.is

Abstract: The critique of electronic voting tends to stem from two primary factors: technical concerns about the security of electronic voting, and the level of technical understanding voters must be presumed to have in order to trust electronic voting methods. In this paper I will show that the issue of security regards guaranteeing two features: unlinkability and verifiability. I will address some of the conditions required to guarantee unlinkability and verifiability, show why they have never coexisted in any voting system, electronic or otherwise; and propose conditions which, if met, should guarantee simultaneous unlinkability and verifiability. In order to do that, I shall attempt to demonstrate that unlinkability and verifiability are necessary and sufficient conditions, given a practical definition of security.
Finally, I will try to address how a system implemented in such a way need not be technically incomprehensible to a reasonable voter.

Keywords: e-voting, security, unlinkability, verifiability, stability, calculability, nonfalsifiability

УСЛОВИЯ ЗА ЕЛЕКТРОННО ГЛАСУВАНЕ

Смари МакКарти, Международен институт за медерни медии, Исландия,
smari@smarimccarthy.is

Абстракт: Критиките на електронното гласуване имат две основни предпоставки: първата е от технически характер и е свързана със сигурността при електронното гласуване; втората е по-скоро психологическа, свързана с нивото на технически умения, които са необходими на един гласоподавател, за да се довери на методологията за електронно гласуване.
В тази статия ще покажа, че гарантирането на сигурността има две страни: (1) невъзможност за връзка (между гласоподавател и глас – б. пр.) и (2) възможност за проверка. Ще разгледам някои условия, които гарантират (1) и (2); ще покажа защо те не могат да бъдат едновременно изпълнени в никоя избирателна система (електронна или друга) и ще предложа условия, които, ако са изпълнени, ще гарантират наличието на (1) и (2) едновременно. За целта ще се опитам да покажа, че (1) и (2) са необходими и достатъчни условия в случай, че е дадена практическа дефиниция за сигурност.
Накрая ще се спра върху въпроса как система, в която са налице (1) и (2) не е непременно технически неразбираема за разумния гласоподавател.

Ключови думи: е-гласуване, сигурност, несвързаност, верифицируемост, стабилност, изчислимост, нефалшифицируемост

Статията е предоставена от Проекта "Насърчаване на е-демокрацията в България с най-добър опит от Исландия", изпълняван от Българско сдружение за насърчаване на гражданската инициатива, Бургаски свободен университет и Международен институт за модерни медии - Рейкявик с финансовата подкрепа на Програмата за НПО на ФМ на Европейското икономическо пространство.

Introduction

Voting is one of the central activities of democracy. It is the way by which democratic societies arrive at collective decisions. But voting is not a constant, well defined concept: it has many variations. Some of these variations are good, others bad. It ultimately comes down to the properties of the system at hand, and various properties have been shown to be more effective in building what Bergson [1] referred to as a social welfare function, such as Arrow's conditions [2], Gibbard's conditions, and others. While these are all important for the purposes of voting systems in general, they are insufficient to address many of the concerns regarding electronic voting.

Our aim here is to try to nail down further conditions for voting systems. In order to do so, we should define a voting system as a tuple S consisting of a ballot method B, a social welfare function f (SWF), and a set of alternatives C that an ordering or selection needs to be made from. For these, we shall use the notation

$$S = (B, f, C),$$

A ballot method is a protocol for acquiring a social choice from a person. For formal purposes, we can consider a balloting method B to be a function such that for a population A:

$$B(A) = V,$$

a set of votes. Depending on the context, one include various other aspects as parameters, such as the nature of the social welfare function, the mood of the population, or generally any external factors which might alter the population's social choices. For purposes of simplicity we shall simply reduce all of these factors, with the exception of the social welfare function f and the choice set C, to a single variable representing all external factors, e.

Then balloting method used to generate votes is:

$$B_f(A, C, e) = V.$$

A social welfare function (SWF) is the part of a voting system which establishes a social choice (i.e., the election result) given a set of social choices (expressed as ballots). It is simply a function f such that the social welfare function used to calculate results

$$f(V) = R,$$

where R is a possible choice.

Definition of security

Security is a contextual and frequently overloaded term. The factors determining the security of a system are dependent on the threats assumed to exist towards the system and the vectors available to an attacker to make good of those threats. When dealing with voting systems, we generally want to guarantee several independent aspects with regard to security. In the case of voting systems, we generally wish to guarantee that:

a) tampering with ballots is not possible (nonfalsifiability),

b) calculation of results can be conducted correctly (calculability),

c) result calculations are stable (stability),

d) no voter can be coerced into making a particular choice (noncoercion).

For the purposes of this paper, a voting system which is nonfalsifiable, calculable, stable and noncoercive is secure, but let’s define each in more detail.

Stability

A voting system is stable if the votes alone are sufficient to calculate the result, and that external factors at the time of counting do not affect the outcome of the count. If two people were to start calculating the outcome from the same set of ballots, using the same balloting method, at different times, and arrive at a different result, then the system is considered unstable.

From this, we arrive at

Definition 1: A voting system $S = (B, f, C)$ is stable if, given a set of votes V, $f(V) = R$, independent of any other factors.

Many voting systems, such as some implementations of single transferable vote, employ an aspect of randomness in order to resolve conditions of equality. While this kind of random resolution can be said not to reflect social choice, it can be argued that this merely implies that the social welfare function being employed to determine the results is not adequately capturing the social choice.

Regardless of whether this is true or not, the result is the same: two independent actors, given the ballots and the voting system, may arrive at different results in the case of a voting method which employs randomness. This alone makes verifiability substantially harder to achieve. If I wished to independently verify the official outcome, given the balloting method and the votes, I would have to first calculate all possible outcomes of the vote by varying the random factors to depletion. Only then could I even build a probability estimation that the random factor used was chosen fairly but I could still not verify that it had in fact been chosen fairly.

This variation can, however, be eliminated if all parties conducting counts agree upon a pseudorandom number generator function and an initial seed, preferably selected prior to conducting the election. For the purposes of this paper, we assume that stability of social welfare functions which employ a random factor is not eliminated by the existence of the random factor, by assuming that a pseudorandom number generator function and initial seed are specified a priori.

This does not mean that stability cannot be eliminated by other factors. In this case, the pseudorandom number generator and its initial seed are provided as external factors in the balloting mechanism.

Calculability

Calculability is in some ways the simplest of the conditions. A voting system is calculable if, given a set of social choices and a social welfare function, a result can be calculated. Under the conditions set forward by Arrow [2], calculability is guaranteed if the social welfare

function has an unrestricted domain. It is possible, however, that a social welfare function which does not have unrestricted domain may be used in a voting system which is nevertheless calculable, if restrictions are placed on the ballot method which are at least as limiting as the restrictions on the domain of the social welfare function. This form of restriction is generally called a limitation on the admissible set [2].

Definition 2: A voting system $S = (B, f, C)$ is calculable if, for any set of votes V that can be generated under B, $f(V)$ exists.

Nonfalsifiability

Nonfalsifiability is the property that, after each social choice is made by a person, that choice cannot be altered by a third party. This condition is met if it can be guaranteed that every social choice is an input to the social welfare function. While mathematically this does not seem like a difficult condition, in practice it is the hardest to guarantee by far.

In traditional pencilandpaper voting, the pathway a ballot takes from the voter in the voting booth to the point at which it is counted, and the result subsequently from there to the point where it is made public knowledge, is riddled with vulnerable points. Amongst some of these vulnerabilities are the exclusion of a legitimate voter from creating a ballot, a third party illegitimately creating and entering a ballot on the behalf of another voter, a third party adding illegitimate votes to the social choice set, the refusal to allow a legitimate voter to add a ballot to the social choice set, the destruction or replacement of a ballot box or particular ballots from it, a falsification in the tallying of the ballots, the elimination of social choices from the social choice set, the replacement of the legitimate social welfare function with an alternative function (including altering constants and other subsidiary inputs, such as a random seed), and the manipulation of the final result. More or less the same set of vulnerabilities apply to electronic methods, however, they are often less detectable in electronic voting systems. This is not as such due to an inherent insecurity of electronic systems, but rather the degree to which byzantine fault tolerance [3] has been introduced into pencilandpaper balloting systems. As every step of paper ballot systems with manual counting can be watched over in real time and inspected in detail by any number of people, a certain amount of nonfalsifiability is introduced into the system. Because of this, I refer to such systems as having byzantine verifiability if all other security conditions are met.

Such a system should not be assumed to be nonfalsifiable: if there is no explicit way to prove a posteriori that tampering did not occur, it must be assumed that tampering may have occurred.

A nonfalsifiable method is therefore one which makes it demonstrably impossible for any party to alter any aspect of the voting system, the social choice set, or the resulting calculations in a way which is undetectable.

Definition 3: A voting system $S = (B, f, C)$ is nonfalsifiably $f(V) = R$ if:

1) $B_f(A, C, e) = V$;
2) B , f , A , C , e are known before V is calculated;
3) V is publicly known after it is calculated;
4) for the voting population A, $\forall\, a \in A : B_f(\{a\}, C, e) \in V$;

5) anybody having B , f , A , C , e can independently arrive at R.

This is a relatively complicated formulation, and should perhaps be formulated in terms of zero knowledge proofs, and using temporal logic. At any rate, this suffices for our purposes, as long as we acknowledge there to be in any election four active moments:

1) initialization, where the balloting method B, the social welfare function f, and the voters A are clearly defined;

2) preballoting, where the choice set C and the external factors e are defined;

3) balloting, where V is calculated;

4) postballoting, where R is calculated.

None shall be declared out of turn, and at every point before these moments, before, between and after them, it should be possible for anybody to independently verify the state of the overall system.

Noncoercion

The notion of noncoercion focuses not on the details of the voting system and the social choice set, but rather on the question of whether a social choice entered by a person represents their real choice. Here, we try to distinguish between tactical voting on the one hand, where a person decides to misrepresent her social choice in order to attempt to increase presumed or perceived benefit from the social welfare function, and on the other hand situations where a person is in some way made to misrepresent her social choice in a way which negatively impacts the person's presumed benefit, or the total benefit derived from the social welfare function. This would typically occur for the purpose of maximizing the perceived benefit of another person or group of people.

It is not selfevident that a social choice set which contains a ballot created through a process of coercion is necessarily Paretounoptimal.

We regard coercion as a negative impact on the voting system as it eliminates a legitimate voter's social choice from the social choice set and replaces it with a fraudulent social choice. In this sense it is equivalent to falsification, except that it happens prior to the balloting action by a given person, and is therefore undetectable as falsification under the conditions of nonfalsifiability.

In reality, a common scheme for achieving this is where a person intending to rig the elections acquires through forgery, theft or extraction from the balloting site a single ballot. This ballot is prefilled with the selection said person wishes to make. When another person comes to the precinct to vote, they are stopped by the rigger, who takes a copy of their ID and hands them the ballot. The voter, under duress, enters the precinct, takes an empty ballot, enters the balloting box, then returns and casts the prefilled ballot, exiting the precinct with an empty, unfilled ballot as proof of having obeyed. Disobedience often leads to violence, under this scheme.

This particular scheme could be defeated by verifying that voters do not have ballot sheets with them when entering the precinct, but that is invasive and difficult. Even if it were

practical to do so, it is relatively easy to construct a host of similar schemes for both pencilandpaper voting and electronic voting, which all suffer the same failing point: once a person has been forced in some way to cast a vote under duress, the person has no way to undo the damage.

A solution to this would be to allow voters to cast multiple votes, with each subsequent vote annulling the previous vote. This is difficult to accomplish in practice. It requires that each vote is linked to a voter in a unique way which allows for the vote's correct identification and removal in the case where it has been eliminated. This quickly becomes messy at scale: when processing millions of ballots, having to check each ballot for elimination, while simultaneously not accidentally falsifying the election by either eliminating an incorrect vote or not eliminating an invalid vote, can be uneconomical. Electronic voting makes this easier.

Yet remains the possibility that a person is coerced to vote, and is then eliminated before the person can recast the vote without duress. Such elimination is only useful to a potential voterigger in the case where a person's vote continues to be valid even after they are deceased. It is a grim notion to have to consider, but it may work to the benefit of a voter to require that all votes during counting belong to living voters. This too may be difficult to enforce in practice, and raises questions about temporary or permanent voter debilitation, which is outside the scope of this paper. Let us include a simple version of this for now.

To add one last point of complexity: the ability to replace votes implies the ability to identify votes to voters, which may also be the basis for schemes leading to violence. Therefore we must maintain that paradoxically a vote should not be linkable to the voter who cast it. We shall deal with this paradox shortly, but for now our complete definition is:

Definition 4: A voting system $S = (B, f, C)$ is noncoercive if

a) A voter $a \in A$ can cast votes $v_1, v_2, v_3, \dots, v_n$, such that $v_n \in V$ but $v_i \notin V$ for $i < n$;

b) Deceased voters votes are removed from V prior to the calculation of R;

c) V is unlinkable to A.

Unlinkability

Traditional voting schemes assume that noncoercion can be achieved by enforcing that a voter is alone in a regulated voting booth, and that from the point in time when the voter enters his ballot into the ballot box, the ballot is anonymous. This does not necessarily hold true. Depending on the design of the ballot, the design of the ballot box, and the conditions of the polling station, it is possible to violate this assumption. The ubiquity of mobile phones with cameras and the fact that although cameras are typically forbidden in polling stations, the ban is not strictly enforced, implies that a person can be coerced to privately take a picture of his ballot before submitting it; this can be used to demonstrate compliance. Often ballots are marked in some way to guarantee authenticity and uniqueness, or to indicate origin. This can be used to violate voter anonymity, or at least substantially reduce the set of possible voters who could have cast a particular ballot.

These assumptions, apart from not being sufficient to guarantee anonymity, are in a sense attempting to solve the wrong problem. The issue is not specifically anonymity , i.e. whether the identity of the voter can be obscured, but rather the special case of whether a ballot or a social choice can be linked back to a particular voter. This is a special case of anonymity which I call unlinkability .

Formally, given a voter a in a set of voters A, and a set of social choices V, if there exists a function L such that $L(V_a) = a$ but $L(V_m) \neq a$ for $\forall\, m \in A \setminus \{a\}$, then the voting system is linkable. If no such function exists, the system is unlinkable. Put simply: if the exclusion of a vote from the set of votes is sufficient for the identification of the voter, then the voter can be identified.

Unlinkability and verifiability

The security properties defined above are in fact conditions of verifiability : without them, it is impossible to independently verify the outcome of an election. It can be done to a substantial degree while violating some of the conditions, as is done in the common practice of byzantine verifiability, which in practice is simply the act of adding more people with supposedly different interests and allegiances into the process of conducting a vote until everybody is satisfied that the chance of any abuse has been made statistically insignificant.

However, as verifiability requires noncoercion, and noncoercion requires simultaneously that votes can be recast in a way that invalidates previous votes, and that votes cannot be linked back to the voter who cast them, we are faced with the interesting scenario where verifiability and unlinkability have never coexisted in any voting system.

Let’s take a moment to fully grasp the importance of that. In traditional pencilpaper voting systems, the vote being cast is (generally speaking and outside of exceptional circumstances) not linkable to the person who cast the vote from the moment the vote was put into the ballot box onwards. Assuming all has gone well up until this point in time, the voter now believes that his vote will be included in the determination of R. Most of the time, this will hold true. But as long as at least one vote with an identical value to that chosen by a given voter is included, the voter cannot independently verify that her ballot is included after it has been cast. It may be possible to reason that some votes may not have been included on inspection, but this does not help to identify where the problem occurred.

Adding more information to the ballot could help. If b bits of information are provided on each ballot, then there is a high probability that 2 b/2 individuals could identify their ballots based solely on their values. This is however perhaps impractical for large populations, as the number of aggregate options on each ballot would have to equal the size of the population. This is untenable in practice, but perhaps it would be possible, in an electronic setting, to add more information to the point of unpredictability, in a way that does not influence the ballot sheet itself.

Conditions for coexistence

There has been much recent excitement about the blockchain mechanism, developed for Bitcoin [4]. It is in effect nothing more than a cryptographically verifiable, distributed, appendonly log: anybody can add an entry, nobody has full control over the log, and anybody can verify that the log has not been tampered with. This gives us the basis for an interesting feature.

The blockchain alone is insufficient though. Another useful mechanism is a type of mathematical proof known as a zero knowledge proof. It allows the verification of a statement without there being any information exchange. For Bitcoin, an extension has been suggested called Zerocoin [5], which strongly anonymizes cryptographic tokens. It is vaguely akin to money laundry: anonymous tokens are exchanged randomly in unpredictable ways, each time mediated through a zero knowledge proof to ensure that no “paper trail” is left behind.

Using these two mechanisms, we can create a voting system like so: Each eligible voter gets, by some mechanism, a unique cryptographic token – for instance a set of points on an elliptic curve generated by them and signed by an electoral authority. The electoral authority publishes the voter roster and their signature. Voters use a mechanism similar to Zerocoin to “launder” their voting tokens, rendering them unidentifiable, even to the electoral authority. Voters use their newly laundered anonymous tokens to sign their votes and append them to the block chain.

With this, it is possible to verify that every vote belongs to a legitimate person, but impossible to say who (unless the private part of that person's key is published). If a voter casts multiple votes, there will be multiple votes signed with the same key. In that case, the most recent vote can be considered valid, and older votes discarded.

This fulfills, in theory, all of the requirements for unlinkability and verifiability. In particular, anybody can verify that:

1. Votes are signed by a token that was legitimately issued to a legitimate voter
2. The system is calculable
3. Calculations are stable
4. No falsification can occur
5. Repudiation is impossible
6. Coercion is impossible

The process of establishing such a voting system requires that first, the system specifications, including the balloting method, the social welfare function, the set of ballot options, and the set of voters, are published at the beginning of the blockchain. Second, our e value, consisting of any external factors, is also published at the beginning of the blockchain. A rule can state that e and other parameters cannot be altered after the first vote has been appended, and shall be disregarded if such an event is to occur.

A word on delegable voting

With this system, it is relatively easily to demonstrate that any balloting method and social welfare function will work. The blockchain mechanism does not impose any conditions which would prevent any known balloting method or social welfare function from workign. However, in order to support liquid democracy style vote delegation [6], or proxying, a further complication can be added in the form of derived keys. Vote delegation is a process by which any third party can be nominated to participate in the election on one’s behalf, although at any time it is possible to revoke the nomination. These processes essentially transform the set of voters into a directed acyclic graph.

A voter can generate a new derived key which is linked to their private key in a verifiable way. They can share this key with the person they are delegating to, who can then cast votes on their behalf using it. If the voter wishes to revoke the proxying, they sign a revocation certificate for the derived key using the parent key. This however only works if we can construct a key derivation scheme with these properties.

Such a scheme can be easily established by making a distinction between delegation keys and voting keys , where the tokens which are swapped using the zero knowledge mechanism are delegation keys . These keys can be used to sign voting keys , the most recent of which for each delegation key is valid. The voting keys can then be established with a third party via a standard DiffieHellman key exchange [7].

Conclusion

This paper has not been particularly rigid in its approach, providing no proofs, as such. There is much work to be done in this direction. However, we have established a rough draft of conditions required, and suggested a mechanism which may work. Further work is needed to prove the veracity of these statements, in particular the unlinkability statement, which currently stands on relatively shaky ground.

References

[1] BERGSON A., A Reformulation of Certain Aspects of Welfare Economics. Quarterly Journal of Economics, 52(2), February 1938

[2] ARROW K. J., Social Choice and Individual Values. 1951

[3] LAMPORT L., SHOSTAK R.; PEASE M., The Byzantine Generals Problem. ACM Transactions on Programming Languages and Systems, 4 (3), pp.382–401, 1982.

[4] NAKAMOTO S., Bitcoin: A PeertoPeer Electronic Cash System., 2008

[5] MIERS I., GARMAN C., GREEN M., RUBIN A. D., Zerocoin: Anonymous Distributed Ecash from Bitcoin. IEEE Symposium on Security and Privacy, Computer Society Conference Publishing Services. pp. 397–411 , 2013

[6] McCARTHY S., The End of Artificial Scarcity. Free Beer, 2008

[7] DIFFIE W., Hellman M., New directions in cryptograph”. IEEE Transactions on Information Theory, 22 (6). pp. 644–654, 1976.