

How to create trust in electronic voting over an untrusted platform

Gerhard Skagestein¹, Are Vegard Haug², Einar Nødtvedt³, Judith Rossebø⁴

¹University of Oslo, Dept. of Informatics
Box 1080 Blindern, N-0316 Oslo, Norway
gerhard@ifi.uio.no

²University of Oslo, Dept. of Political Science
Box 1097 Blindern, N-0317 Oslo, Norway
a.v.haug@stv.uio.no

³Senit rådgivning AS
Skogtunet 12, N-1369 Stabekk, Norway
einar@senit.no

⁴Norwegian University of Science and Technology, Dept. of Telematics,
and Telenor R&D
Snarøyveien 30, N-1331 Fornebu, Norway
Judith.Rossebo@telenor.com

Abstract: Casting electronic votes via an inherently unreliable channel like the Internet in an uncontrolled environment is controversial for two main reasons: The first one is of democratic nature and the second of technical nature. The democratic concerns are about the possible dangers of buying and selling votes and so called "family voting". The technical concerns are how to convince everybody involved that the votes will be anonymously and accurately recorded and counted, and that no votes will get changed or lost, and that no "fake votes" will be introduced, with the knowledge that any computerized system may contain bugs or may be hacked by evildoers.

In this paper, we will show how the principle of repeated vote casting may be used to alleviate both the democratic and the technical concerns above, and how hybrid cryptography makes it possible for the voter to inspect his votes as stored within the voting system.

1 Introduction

In 2004, the Ministry of Local Government and Regional Development in Norway mandated a working group to work out a recommendation concerning the future of electronic elections in the country. The result of this work is documented in the report [KRD2006]. The basic conclusions are that there is no need to rush into electronic voting and that electronic solutions should be introduced with great care, due to the current deficiencies in the technical platform. Yet the working group recommended the setup of a project group and a step-by-step introduction of e-voting for certain types of elections. However, we do not know when the solutions proposed in this paper will be turned into reality, or whether they will be realised at all.

The working group rather quickly arrived at the conclusion that it had little value to put electronic solutions into the polling places – the ultimate goal had to be to give the voter the possibility to vote in uncontrolled environments from his home or at work. The particular solutions described in this paper is recommended as the basis for a possible system for Internet-voting in uncontrolled environments, as an alternative to solutions built on trustworthy platforms which may show up in the future.

The working group has been very well aware of the Recommendation No. R (2004) 11 of the Committee of Ministers to member states on E-voting [Rec2004] (later in this paper referred to as “the Recommendation”), and maintain the point of view that the proposed solutions are compatible with its intentions, although perhaps not always with its lettering.

Readers familiar with the Estonian electronic voting system [Maaten2004] [NEC2004] will find a lot of similarities. However, it may be of interest to know that the working group did arrive at similar principles before obtaining detailed knowledge about the Estonian system.

2 Two important principles

The solution proposed in this paper relies heavily on two fundamental principles: The principle of two-phase voting and the principle of repeated vote-casting.

2.1 The principle of two-phase voting

Elections in Norway have for a long time been carried out in two phases: One advanced voting phase followed by the Election Day itself. Between the two phases there is a one or two day break. During this period, the voters who have voted during the first phase will be marked in the Voter register, so that this information is available for the election officials on Election Day.

We propose to continue with this two-phase setup. Electronic voting should be used only in the first phase, voting on the Election Day should be done in the traditional way by means of paper ballots. This gives the voters complete freedom in how to vote, electronically or by paper ballots. At some time in the future, electronic voting may become so popular that the efforts for setting up traditional elections will be reduced to almost nothing. This is feasible; however, it will be driven by the preferences of the voters, the politicians and the society in general, not by the technology.

2.2 The principle of repeated vote-casting

The Recommendation [Rec2004], paragraphs 5 to 8, states the obvious democratic rule that a voter should give only one valid vote in each election event (one person – one vote). An electronic system may enforce this rule in two ways, either by invalidating the voter's credentials for further voting in the same election event, or by letting the vote-receiving server in some way keep track of the identity of the voter and reject multiple ballots from the same voter. The first solution is susceptible to conscious or unconscious errors and mistakes on the client side. Hence, it is better to let the server side handle the duplicate ballots from the same voter. We propose to let the vote-storage server store all the received ballots, rather than rejecting the second and the following ballots. At the end of the voting period, the election system will run through the ballots and only the last ballot received from each voter will be transferred to the electronic ballot box. Thus, the voter may effectively regret and cancel his vote just by casting another one at a later point in time.

As a final possibility for repeated vote-casting, the voter may show up on the Election Day asking to vote by means of a traditional paper ballot. In that case, the election officials will register with the vote-receiving server an instruction to throw away all the electronic ballots cast by the voter during phase one.

The principle of repeated vote-casting reduces significantly the well know democratic concerns connected with voting in uncontrolled environments [Maaten 2004]. There will be no market for buying and selling votes, since the buyer can never know whether the voter will cast another vote, maybe even on the Election Day. And the voter who feels subjected to coercion (e.g. "family voting") may cast another vote as soon as the coercer has disappeared. As we shall see, the principle also makes it possible to allow the voter to check the content of his electronic ballot as it is stored on the vote-storage server, since an observer can never know whether this ballot will be the final one.

3 Raising trust by securing the electronic voting system

Whenever communicating over an insecure channel, the demands for security must be built into the applications *using* the insecure channel. Basically, the sender may send a certain amount of redundant data with the message so that the receiver can check the consistency of the data and ask for retransmission if something seems to be wrong, or the receiver may reflect back to the sender its understanding of the message so that the sender can check that the receiver understood the message correctly.

The weakest point in an electronic election system based on voting in uncontrolled environments is probably the client machine, which may be infected by viruses and other malicious programs. The most difficult part to control is the very first part of the journey of a message from the keyboard to the program handling the input from the keyboard. We can not rule out the possibility that some illegal program is sitting between the keyboard and the rest of the system, faking correct looking screen images but sending completely incorrect data to the vote-receiving server. The only (almost) secure way to compensate for this threat is to have the user enter some redundant data via another completely separated and independent channel, for example via SMS on a mobile phone. The user friendliness in such a setup, however, is questionable.

It is more appealing to let the system reflect back to the voter so much data that the voter is convinced that the vote has been correctly registered. In this way, we utilise two different channels between the mind of the voter and the system: The typing on the keyboard and the visual observation of the reflected data on the screen.

It is, however, important that the reflection of the data is not done by an untrusted client machine, but by a trusted, well controlled server. In order to rule out the possibility that the client may intercept the reflected data and make it look right even if it isn't, the data may be returned to the voter via a completely different technical channel, for example SMS on a mobile phone.

The voting client in the system described in this paper is assumed to be a client machine. However, with the emergence of smartphones, GSM telephones equipped with WLAN access, 3G networks and more and more sophisticated mobile terminals, it is feasible that the voting client is a mobile handset. The advantage is that each of these is equipped with a GSM SIM card or a USIM card upon which the user's ID and PKI functions and key pair can be generated and safely contained. Note that in this case, access to the (U)SIM is secured by PIN and PUK, and the users private key never leaves the (U)SIM, see [THJ2004] for details regarding PKI on the (U)SIM.

3.1 The double envelope principle

In order to be able to allow recasting of votes, some kind of voter identity has to follow the ballot until the last, counting ballot is eventually dropped into the electronic ballot box. At the same time, in order to keep the vote secret, the identity of the voter and the content of the ballot must not be made available to anybody at any time. To ensure this, we propose to use an electronic double envelope setup similar to the one used in the Estonian election system [Maaten2004].

We employ two sets of key pairs for asymmetric cryptography. The first set consists of the public and private key of the voter¹. The second set consists of the public and private key for the election event. In addition, we will also employ a session key in a symmetric cryptographic process.

As soon as the voter has finalized his electronic ballot and is ready to send it to the vote-receiving server, the client will generate a random session key and perform a symmetric encryption of the ballot. Then the session key is encrypted with the public key of the election event. The message consisting of the encrypted vote together with the encrypted session key corresponds to a paper ballot in a sealed inner envelope.

Normally, this two-step encryption process, called hybrid crypto, is used just for efficiency reasons, since symmetric crypto-algorithms are much quicker than the asymmetric ones. However, in our scheme, the hybrid crypto is also used for another purpose, as we will see.

Next, the client will digitally sign the message with the private key of the voter. This signed message corresponds to an outer envelope containing the already mentioned inner envelope. To the message, we attach some data which in some way gives the identity of the voter.

The whole package is then sent to the vote-receiving server, which will relay it to a firewall-protected vote-storage server where it will be written to a write-once-medium. Further ballots in double envelopes from the same voter will be written to the same medium, and not overwrite previous ballots. The same will happen with a message from an election official saying that all ballots from the voter should be cancelled. At the end of the election period, the election system will pick the last received ballot (if no cancelling message exists), remove the outer envelope by using the public key of the voter to check the signature on the data (the ballot) and, if verified, drop the inner envelope with the ballot in the electronic ballot box. From this point, there is no connection between the identity of the voters and the content of the ballots. The anonymous enveloped ballots will then be unsealed by decrypting the session key with the private key of the election event (which until then is kept secret inside a security module) and then decrypting the message with the session key.

¹ It is preferred that this key pair is used for much more than just electronic voting – the best solution is that the key pair is a part of an officially recognised PKI-system. This will reduce the possibilities for that the voter is selling the key pair.

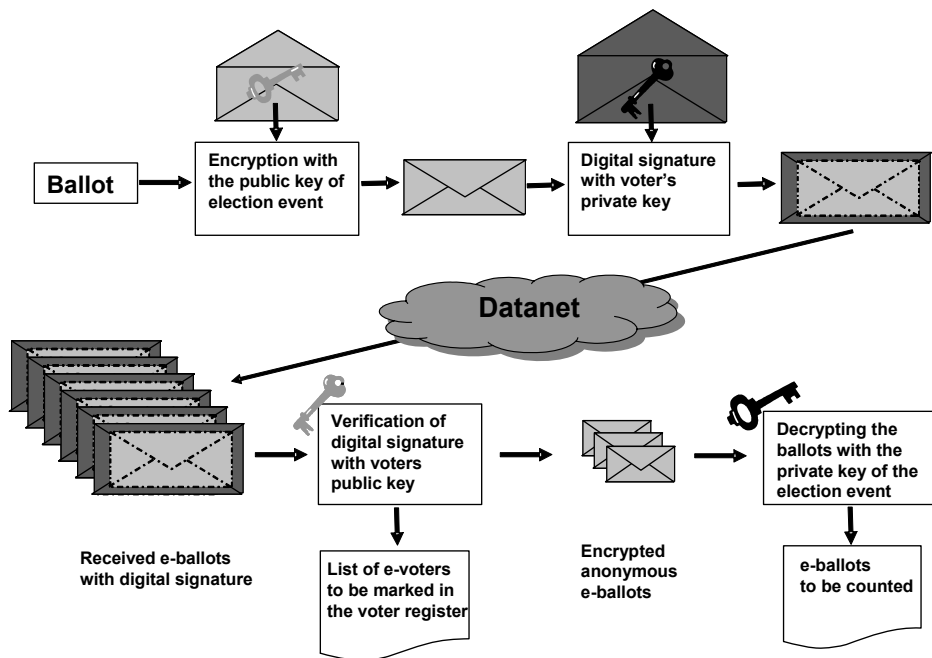


Figure 1: The double envelope principle

3.2 Letting the voter check the ballot

Returning to the question of how to convince the voter that his vote has been properly received, we propose that the doubly encrypted vote is returned to the voter client from the vote-storing server. In this case, the vote-storing server should sign it so that the user is convinced that the vote-storing server is actually storing the ballot. We also propose that the voter at any time during phase one of the election event may request the vote-storing server to send the doubly encrypted ballot to his client.

To be able to decrypt his ballot, the voter must have stored the session key used during voting somewhere. He may have written it down (not very likely), or stored it on a removable storage unit like a memory stick. To store it on the hard disk of the voting client is not to be recommended, for obvious security reasons. With the session key, it is possible for the client machine to open the two envelopes and show the voter the content of his ballot. The outer envelope is opened by decrypting with the public key of the voter, the inner envelope is opened by decrypting with the session key (we are of course not interested in the encrypted session key). The sceptical voter may do this on a client machine different from the one he used for voting – the likelihood that some evildoer may have managed to infect both machines with malicious software that even must show a consistent behaviour, is very small. In the future, this decrypting process may even be done by a mobile phone, so that the voter can use different technical channels for voting and for checking the ballot.

It may well be argued that this functionality is in conflict with paragraph 51 of the Recommendation [Rec2004], stating that "A remote e-voting system shall not enable the voter to be in possession of a proof of the content of the vote cast." If the voter, for any reason, wants to do it, he may show the content of his ballot to anybody, print it or e-mail it, just as he wish. The answer to this objection is of course that the voter may choose to cast another (and maybe completely different) vote at a later stage. Hence, seeing a copy of a ballot stored on the vote-storing server says next to nothing about how the voter finally is going to vote.

The second part of building the voters trust, namely that the final ballot will be dropped in the electronic ballot box, kept anonymous and properly counted, must be solved in a completely different way. The solution here is to use carefully designed and programmed software; verified and certified by an accredited certification institution. Additionally, if deemed necessary, the whole process may be run in parallel on different machines with different software developed by different developers with different methods, and compare the results – so called N-version systems [Liburd2004]. This is possible because this part of the election process can be run on a very limited number of machines in a heavily controlled and secured environment.

3.3 Keeping the votes anonymous

The anonymity of the votes (the impossibility of connecting the content of the ballot to the identity of the voter) rests on the principle that the double enveloped ballots and the private key of the election event should never be available to any person at the same time. Since it is difficult to keep the distribution of the double envelopes stored on the vote-receiving server under complete control (they may be logged for security reasons, or perhaps even copied by a hacker misusing the available functionality for checking the ballot), the solution is to handle the private key of the election event very carefully. It should be stored in a security module (separate hardware container) until it is time to open the inner envelopes, and it should be disposed of as soon as this task is done. In this case, a pin code may also be required in order to enable use of the key.

The degree of anonymity possible with a traditional paper ballot system cannot be guaranteed by an electronic voting system, however, these and other technical means can be employed to guarantee anonymity as far as possible. A security audit is essential to be able to track whether or not the election event key is being misused at any time.

If this solution does not look trustworthy, additional security may be achieved by using voter pseudo-identities. This, however, complicates the task of getting hold of the public key of the voter when opening the outer envelope and the latter solution has therefore not been recommended by the working group.

4 An overall picture of the architecture

Figure 2 depicts the overall architecture of the voting system. In the complete report written by the working group [KRD2006], the functionality of each module is described by means of UML Use cases.

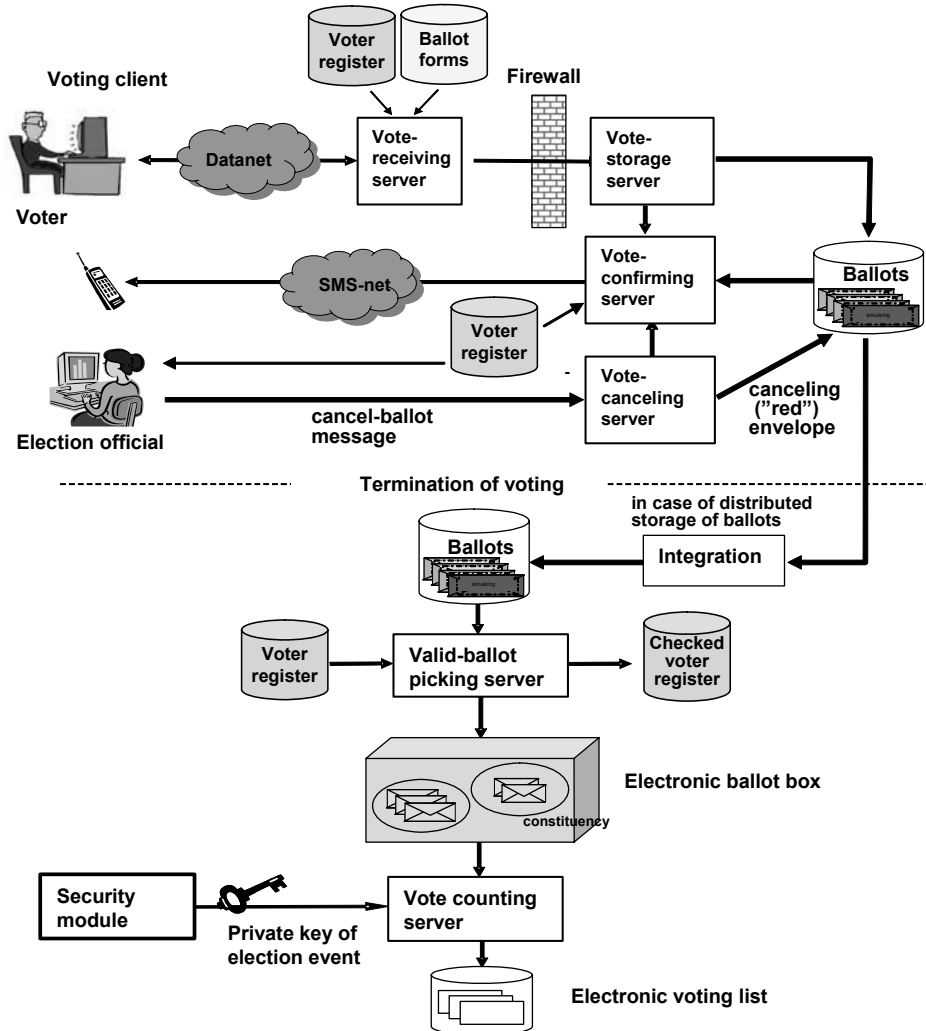


Figure 2: The architecture

5 Other security issues

Securing the availability of the vote-receiving server and other entities involved in the process is essential for the electronic voting to be conducted fairly. It is crucial to ensure that the vote-receiving server does not go down – in particular on the last day. Actually, last year the electronic tax return system in Norway failed to deliver on the last day [Ryv2005], and as a consequence, the deadline for submitting tax returns was extended by one day. For that service this was acceptable to the public, but would the citizens accept such a delay in e-voting? An issue related to this is protecting the vote-receiving server from so called denial of service (DoS) attacks aimed at making the vote-receiving server unavailable. Therefore we envision that a complete election system will encompass several vote-receiving servers, located in different geographical areas. On the other hand, we may have only one vote-storage server, but with well developed backup facilities.

The availability of the underlying network must be also be secured – attacks on parts of the network to take down network segments in order to prevent voters from being able to cast their votes should be anticipated. In a close election, targeting specific neighbourhoods that are known to favour the opposing side by e.g. flooding the local network to prevent votes from being cast electronically is easily carried out. This risk can be mitigated by the solution presented in this paper under the assumption that it is difficult to ascertain when voters will cast their votes. However, traffic analysis over several elections might reveal information on how likely it is that voters cast their votes ahead of time instead of waiting until the last day – one can predict that most people will wait until the last day to cast the final vote. In this case it should be expected that sabotage by creating denial of service attacks targeting the voting traffic may be widespread. Preventing this type of sabotage will be challenging, as known attacks may be easy to prevent, but new and effective attacks may come as a surprise, making it difficult to even mitigate the attack. We only have to look at the example of sabotage of the “Get out the vote” operations regarding organized phone jamming during the 2002 Senate race in New Hampshire in the United States [Coh2006] to get a flavour of how easy it may be to attack the underlying network. A very big concern with this type of sabotage carried out in a broadband network (both fixed and mobile) is that it may be difficult to discover, and the extent of the sabotage may not be uncovered until long after the election results have been certified.

This is only one example of electronic vote sabotage. For the system described in this paper, sabotaging the electronic vote system by attacking the underlying untrusted network should be considered carefully. For example, in a broadband IP-based network it may be easy to prevent users from voting electronically or prevent the ballots from arriving. This type of attack is of course easily discovered by the user, but if the user does not anticipate that this may be a problem and waits to the last minute to cast his/her vote electronically, he may be forced to go the polling place the over next day.

The attack scenarios discussed here show that it is difficult to ensure that voters will have completely equal access to the electronic voting system.

6 Conclusions

We are of the opinion that an e-voting system based on the principles described in this paper has the potential of being universally trusted by the voters, the election administrators, the politicians and the society in general. The principle of repeated vote-casting alleviates the well known democratic concerns with electronic voting in uncontrolled environments. At the same time, it allows for the voter to inspect his ballot as it is stored on the vote-receiving server without threatening the secrecy of the final and counting vote. In order to make it possible for the voter to decrypt the doubly enveloped ballot, the session key used during vote casting must have been stored on some medium, for example a memory stick. In order to build trust to the part of the system which is picking the valid votes and counting them, this part of the system should be designed and programmed very carefully, and verified and certified by an accredited certification institution. If deemed necessary, the whole process may be run in parallel on different platforms and the results compared (N-version system). However, it will still be difficult to ensure that voters will have completely equal access to the e-voting system.

References

- [Coh2006] Cohen, A.: *A small time crime with hints of big time connections lights up the net*. http://www.nytimes.com/2006/04/17/opinion/17mon4.html?_r=3&oref=login&pagewanted=print&oref=slogin
- [KRD2006] Rapport: Elektronisk stemmegivning – utfordringer og muligheter. Kommunal og regionaldepartementet 2006. (In Norwegian – an English version will follow.) http://odin.dep.no/krd/norsk/dok/andre_dok/rapporter/016051-220022/dok-bn.html
- [Liburd2004] Liburd, Soyini: *An N-version Electronic Voting System*. Caltech/MIT Voting Technology Project Working paper # 17, July 2004 http://vote.caltech.edu/media/documents/wps/vtp_wp17.pdf
- [Maaten2004] Maaten, Epp: *Towards remote e-voting: Estonian case*. In Prosser & Krimmer (Eds.): *Electronic Voting in Europe – Technology, Law, Politics and Society*. Proceedings, Gesellschaft für Informatik 2004. <http://www.e-voting.cc/files/E-Voting-in-Europe-Proceedings/>
- [NEC2004] The National Election Committee: *E-Voting System – Overview* <http://www.vvk.ee/elektr/docs/Yldkirjeldus-eng.pdf>
- [Rec2004] Recommendation No. R (2004) 11 of the Committee of Ministers to member states on E-voting
- [Ryv2005] Ryvarden, E.: *Skatte-servere tålte ikke trykket*. Digi.no, April 30 (2005) (in Norwegian)
- [THJ2004] Johannessen, Tor Hjalmar: *On the mobile, its security issues and applicability potentials*. Teletronikk, 100 (1), ISSN 0085-7130, 2004.

Session 5: Redesigning Workflows for Electronic Voting

