

Handout 2 (E-Voting)

In security engineering, there are many counter-intuitive phenomena: for example I am happy (more or less) to use online banking every day, where if something goes wrong, I can potentially lose a lot of money, but I am staunchly against using electronic voting (let's call it e-voting for short). E-voting is an idea that is nowadays often promoted in order to counter low turnouts in elections¹ and generally sounds like a good idea. Right? Voting from the comfort of your own home, or on your mobile on the go, what could possibly go wrong? Even the UK's head of the Electoral Commission, Jenny Watson, argued in 2014 in a Guardian article that the UK should have e-voting. Her plausible argument is that 76% of pensioners in the UK vote (in a general election?), but only 44% of the under-25s. For which constituency politicians might therefore make more favourable (short-term) decisions is clear. So being not yet pensioner, I should be in favour of e-voting, no?

Well, it turns out there are many things that can go wrong with e-voting, as I like to argue in this handout. E-voting in a "secure way" seems to be one of the things in computer science that are still very much unsolved. It is not on the scale of Turing's halting problem, which is proved that it can never be solved in general, but more in the category of being unsolvable with current technology. This is not just my opinion, but also shared by many security researchers amongst them Alex Halderman, who is the world-expert on this subject and from whose Coursera course on Securing Digital Democracy I have most of my information and inspiration. It is also a controversial topic in many countries:

- The Netherlands between 1997–2006 had electronic voting machines, but "hacktivists" had found they can be hacked to change votes and also emitted radio signals revealing how you voted.
- Germany conducted pilot studies with e-voting, but in 2007 a law suit has reached the highest court and it rejected e-voting on the grounds of not being understandable by the general public.
- UK used optical scan voting systems in a few trail polls, but to my knowledge does not use any e-voting in elections.
- The US used mechanical machines since the 1930s, later punch cards, now DREs and optical scan voting machines. But there is a lot of evidence that DREs and optical scan voting machines are not as secure as they should be.
- Estonia used since 2007 the Internet for national elections. There were earlier pilot studies for voting via Internet in other countries.
- India uses e-voting devices since at least 2003. They use "keep-it-simple" machines produced by a government owned company.

¹In my last local election where I was eligible to vote only 48% of the population have cast their ballot. I was, I shamefully admit, one of the non-voters.

- South Africa used software for its tallying in the 1993 elections (when Nelson Mandela was elected) and found that the tallying software was rigged, but they were able to tally manually.

The reason that e-voting is such a hard problem is that we have requirements about the voting process that conflict with each other. The five main requirements for voting in general are:

- **Integrity**

- By this we mean that the outcome of the vote matches with the voters' intent. Note that it does not say that every vote should be counted as cast. This might be surprising, but even counting paper ballots will always have an error rate: people after several hours looking at ballots will inevitably miscount votes. But what should be ensured is that the error rate does not change the outcome of the election. Of course if elections continue to be on knives edges we need to strive for rather small error rates.
- There might be gigantic sums at stake and need to be defended against. The problem with this is that if the incentives are great and enough resources are available, then maybe it is feasible to mount a DoS attack against the voting server and by bringing the system to its knees, change the outcome of an election. Not to mention to hack the complete system with malware and change votes undetectably.

- **Ballot Secrecy**

- Nobody can find out how you voted. This is to avoid that voters can be coerced to vote in a certain way (for example by relatives, employers etc).
- (Stronger) Even if you try, you cannot prove how you voted. The reason for this is that you want to avoid vote coercion, but also vote selling. That this can be a problem is proved by the fact that some jokers in the recent Scottish referendum tried to make money out of their vote.

- **Voter Authentication**

- Only authorised voters can vote up to the permitted number of votes (in order to avoid the "vote early, vote often").

- **Enfranchisement**

- Authorised voters should have the opportunity to vote. This can, for example, be a problem if you make the authorisation dependent on an ID card, say a driving license. Then everybody who does not have a license cannot vote. While this sounds an innocent requirement, in fact some parts of the population for one reason or another just

do not have driving licenses. They are now excluded. Also if you insist on paper ballots you have to have special provisions for blind people. Otherwise they too cannot vote.

- **Availability**

- The voting system should accept all authorised votes and produce results in a timely manner. If you move an election online, you have to guard against DoS attacks for example.

While these requirements seem natural, the problem is that they often clash with each other. For example

integrity vs. ballot secrecy
authentication vs. enfranchisement

If we had ballots with complete voter identification, then we can improve integrity because we can trace back the votes to the voters. This would be good when verifying the results or when recounting. But such an identification would violate ballot secrecy (you can prove to somebody else how you voted). In contrast, if we remove all identification for ensuring ballot secrecy, then we have to ensure that no “vote-stuffing” occurs. Similarly, if we improve authentication by requiring to be present at the polling station with an ID card, then we exclude absentee voting.

To tackle the problem of e-voting, we should first have a look into the history of voting and how paper-based ballots evolved. Because also good-old-fashioned paper ballot voting is not entirely trivial and immune from being hacked. We know for sure that elections were held in Athens as early as 600 BC, but might even date to the time of Mesopotamia and also in India some kind of republics might have existed before the Alexander the Great invaded it. Have a look at Wikipedia about the history of democracy for more information. These elections were mainly based on voting by show of hands. While this method of voting satisfies many of the requirements stipulated above, the main problem with hand voting is that it does not guaranty ballot secrecy. As far as I know the old Greeks and Romans did not perceive this as a problem, but the result was that their elections favoured rich, famous people who had enough resources to swing votes. Even using small coloured stones, which were also used at that time, did not really mitigate the problem with ballot secrecy. The problem of authorisation was solved by friends or neighbours vouching for you to prove you are eligible to vote (there were no ID cards in ancient Greece and Rome).

Starting with the French Revolution and the US constitution, people began to value a more egalitarian approach to voting and electing officials. This was also the time where paper ballots started to become the prevailing form of casting votes. While more resistant against voter intimidation, paper ballots need a number of security mechanisms to avoid fraud. For example you need voting booths for being able to fill out the ballot in secret. Also transparent ballot boxes

are often used in order to easily detect and prevent vote stuffing (prefilling the ballot box with false votes).



Another security mechanism is to guard the ballot box against any tampering during the election until counting. The counting needs to be done by a team potentially involving also independent observers.

One interesting attack against completely anonymous paper ballots is called *chain vote attack*. It works if the paper ballots are given out to each voter at the polling station. Then an attacker can give the prefilled ballot to a voter. The voter uses this prefilled ballot to cast the vote, and then returns the empty ballot paper back to the attacker who now compensates the voter. The blank ballot can be reused for the next voter.

To sum up, the point is that paper ballots have evolved over some time and no single best method has emerged for preventing fraud. But the involved technology is well understood in order to provide good enough security with paper ballots.

E-Voting

If one is to replace paper ballots by some electronic mechanism, one should always start from simple premise taken from an Australian white paper about e-voting:

“Any electronic voting system should provide at least the same security, privacy and transparency as the system it replaces.”

Whenever people argue in favour of e-voting they seem to be ignoring this basic premise.

After the debacle of the Florida presidential election in 2000, many voting precincts in the US used Direct-Recording Electronic voting machines (DREs) or optical scan machines. One popular model of DREs was sold by a company called Diebold. In hindsight they were a complete disaster: the products were inadequate and the company incompetent. Direct recording meant that there was no paper trail, the votes were directly recorded on memory cards. Thus the voters had no visible assurance whether the votes were correctly cast. Even if there is



Figure 1: Direct-Recording Electronic voting machines above; an optical scan machine below.

a printout provided; it does not give any guaranty about what is recorded on the memory card.

The machines behind these DREs were “normal” windows computers, which could be used for anything, for example for changing votes. Why did nobody at Diebold think of that? I have no idea. But that this was eventually done undetectably is the result of the determination of ethical hackers like Alex Halderman. His group thoroughly hacked Diebold’s DREs showing that election fraud with them is easily possible. They even managed to write a virus that infected the whole system by having only access to a single machine.

What made matters worse was that Diebold tried to hide their incompetency and the inferiority of their products, by requiring that election counties must not give the machines up for independent review. They also kept their source secret. This meant Halderman and his group had to obtain a machine not through the official channels. They then had to reverse engineer the source code in order to design their attack. What this all showed is that a shady security design is no match to a determined hacker.

Apart from the obvious failings (for example no papertrail), this story also told another side. While a paper ballot box need to be kept secure from the beginning of the election (when it needs to be ensured it is empty) until the end of the day, electronic voting machines need to be kept secure the whole year. The reason is of course that one cannot see whether somebody has tampered with the program a computer is running. Such a 24/7 security is costly and often even impossible, because voting machines need to be distributed usually the day before the election to the polling stations. These are often schools where the voting machines are kept unsecured overnight. The obvious solution of

putting seals on computers did not work: in the process of getting these DREs discredited (involving court cases) it was shown that seals can easily be circumvented. The moral of this story is that election officials were incentivised with money by the central government to obtain new voting equipment and in the process fell prey to pariahs which sold them a substandard product. Diebold was not the only pariah in this area, but one of the more notorious ones.

Optical scan machines are slightly better from a security point of view but by no means good enough. Their main idea is that the voter fills out a paper ballot, which is then scanned by a machine. At the very least the paper ballot can serve as a paper trail in cases an election result needs to be recounted. But if one takes the paper ballots as the version that counts in the end, thereby using the optical scan machine only as a device to obtain quickly preliminary results, then why not sticking with paper ballots in the first place?

An interesting solution for e-voting was designed in India. Essentially they designed a bespoke voting device, which could not be used for anything else. Having a bespoke device is a good security engineering decision because it makes the attack surface much smaller. If you have a full-fledged computer behind your system, then you can do everything a computer can do...and that is a lot, including a lot of abuse. What was bad about the devices in India was that these machines did not have the important paper trail: that means if an election was tampered with, nobody would find out. Even if they had by their bespoke design a very small attack surface, ethical hackers were still able to tamper with them. The moral with Indian's voting machines is that even if very good security design decisions are taken, e-voting is very hard to get right.

This brings us to the case of Estonia, which held in 2007 the worlds first general election that used Internet. Again their solution made some good choices: for example voter authentication is done via the Estonian ID card, which contains a chip like on credit cards. They also made most of their source code public for independent scrutiny. Of course this openness means that people (hackers) will look at your fingers and find code such as this snippet.

```
#!/usr/bin/python2.7
# -*- coding: UTF8 -*-

"""
Copyright: Eesti Vabariigi Valimiskomisjon
(Estonian National Electoral Committee), www.vvk.ee
Written in 2004-2013 by Cybernetica AS, www.cyber.ee

This work is licensed under the Creative Commons
Attribution-NonCommercial-NoDerivs 3.0 Unported License.
To view a copy of this license, visit
http://creativecommons.org/licenses/by-nc-nd/3.0/.
"""

def analyze(ik, vote, votebox):

    # TODO: implement security checks
```

```
# such as verifying the correct size
# of the encrypted vote

return []
```

If you want to have a look their code can be downloaded from their github repository.² Also their system is designed such that Internet voting is used before the election: votes can be changed an unlimited amount of times, always the last vote is tabulated, you can even change your vote on the polling day in person. This is an important security mechanism guarding against vote coercion, which of course is an important problem if you are allowed to vote via Internet.

However, the weak spots in any Internet voting system are the voters' computers and the central server. Unfortunately, their system is designed such that they needs to trust the integrity of voters' computers, central server components and also the election staff. In 2014, a group of independent observers around Alex Halderman were able to scrutinise the election process in Estonia. They found many weaknesses, for example careless handling of software updates on the servers. They also simulated an election with the available software and were able to covertly manipulate results by inserting malware on the voters' computers. Overall, their recommendation is to abandon Internet voting and to go back to an entirely paper-based voting process. In face of state-sponsored cyber-crime (for example NSA), Internet voting cannot be made secure with current technology. They have a small video clip with their findings at

<https://estoniaevoting.org>

This brings us to the question, what could be a viable electronic voting process in *theory* with current technology? In the literature one can find proposals such as this one:

1. Alice prepares and audits some ballots, then casts an encrypted ballot, which requires her to authenticate to a server.
2. A bulletin board posts Alice's name and encrypted ballot. Anyone, including Alice, can check the bulletin board and find her encrypted vote posted. This is to make sure the vote was received by the server.
3. When the election closes, all votes are shuffled and the system produces a non-interactive proof of a correct shuffling—correct in the sense that one cannot determine anymore who has voted for what. This will require a shuffling procedure based on zero-knowledge-proofs.
4. After a reasonable complaint period, let auditors check the shuffling, all shuffled ballots are decrypted, and the system provides a decryption proof for each decrypted ballot. Again this will need a zero-knowledge-proof-type of method.

²<https://github.com/vvk-ehk/evalimine/>

5. Perform a tally of the decrypted votes.
6. An auditor can download the entire (shuffled) election data and verify the shuffle, decryptions and tally.

As you can see the whole process is not trivial at all and leaves out a number of crucial details (such as how to best distribute public keys for encryption). It even depends on a highly sophisticated process called *zero-knowledge-proofs*. They essentially allow one to convince somebody else to know a secret without actually revealing what the secret is. This is a kind of cryptographic “magic”, like the Hellman-Diffie protocol which can be used to establish a secret even if you can only exchange postcards with your communication partner. We will look at zero-knowledge-proofs in a later lecture in more detail.

The point of these theoretical/hot-air musings is to show that such an e-voting procedure is far from convenient: it takes much more time to allow, for example, for scrutinising whether the votes were cast correctly. Very likely it will also not pass the benchmark of being understandable to Joe Average. This was a standard a court rules that needs to be passed in the German election process.

The overall conclusion is that an e-voting process involving the Internet cannot be made secure with current technology. Voting has just too high demands on integrity and ballot secrecy. This is different from online banking where the whole process is designed around authentication. If fraud occurs, you try to identify who did what (somebody’s account got zero; somewhere the money went). Even if there might be more gigantic sums at stake in online banking than with voting, it can be solved. That does not mean there are no problems with online banking. But with enough thought, they can usually be overcome with technology we have currently. This is different with e-voting: even the best have not come up with something workable yet.

This conclusion does not imply that in some special cases of Internet voting cannot be made to work securely. Just in a general election where stakes are very high, it does not work. For example a good-enough and workable in-lecture online voting system where students’ votes are anonymous and students cannot tamper with the outcome, I am sure, can be implemented.

If you want to know more about e-voting, I recommend the highly entertaining online course by Alex Halderman at Coursera.

<https://www.coursera.org/course/digitaldemocracy>

There is also an entertaining TEDtalk by Barbara Simons called “Why can I bank online but not vote online?”

<https://www.youtube.com/watch?v=Wv3VuGZzdK8>

At the beginning she describes the complete break-in by the group of Alex Halderman at the try-out voting at Washington D.C. Halderman’s amusing paper including pictures is at

<https://jhalderm.com/pub/papers/dcvoting-fc12.pdf>