How Could Snowden Attack an Election?

Douglas Wikström[1(✉)], Jordi Barrat[2], Sven Heiberg[3], Robert Krimmer[4], and Carsten Schürmann[5]

[1] KTH Royal Institute of Technology, Stockholm, Sweden
dog@kth.se
[2] University of Catalonia, Barcelona, Spain
jordi.barrat@urv.cat
[3] Smartmatic-Cybernetica Centre of Excellence for Internet Voting, London, UK
sven@ivotingcentre.ee
[4] Tallinn University of Technology, Tallinn, Estonia
robert.krimmer@ttu.ee
[5] IT University of Copenhagen, Copenhagen, Denmark
carsten@demtech.dk

Abstract. We discuss a new type of attack on voting systems that in contrast to attacks described in the literature does not disrupt the expected behavior of the voting system itself. Instead the attack abuses the normal functionality to link the tallying of the election to disclosing sensitive information assumed to be held by the adversary. Thus the attack forces election officials to choose between two undesirable options: Not to publish the election result or to play into the adversary's hand and to publicize sensitive information. We stress that the attack is different from extortion and not restricted to electronic voting systems.

1 Introduction

Existing paper-based voting systems are often considered to be the gold standard against which any other voting system is measured, despite that the classic systems have security weaknesses.

For example, a certain degree of errors when voters fill in ballots is sometimes accepted as long as the voters' intent can still be determined. Similarly, blank ballots may be allowed to give the possibility to vote for unlisted candidates, or as a last resort to counter attacks where ballot papers are stolen from a polling station.

Another example is how the results are reported, e.g., in Norway the results for voting districts that are deemed too small are reported at an aggregate level

We use "Snowden" as a placeholder for somebody in possession of sensitive information and do not in any way suggest that he has any intention to attack any elections. The recent presidential election in USA 2016 show that there may be other parties in possession of similar information with the intent to disrupt elections.

R. Krimmer—Work supported in part by the Estonian Research Council project PUT1361.

R. Krimmer et al. (Eds.): E-Vote-ID 2017, LNCS 10615, pp. 280–291, 2017.
DOI: 10.1007/978-3-319-68687-5_17

to preserve the privacy of voters. Even when a result for a given voting district is not needed to compute the distribution of seats, the result is often considered an important channel of information of the broader democratic system.

However, such weaknesses are typically well known and due to a careful tradeoff between several conflicting goals such as security, availability, cultural values and traditions, and economy.

In this paper we introduce a previously unknown type of attack that should be added to the list of threats to be considered in such tradeoffs. How serious the attack is depends strongly on the strategic value of the election and how well the election management body is prepared to handle it. Important factors include legal, procedural, and the strategic value of causing confusion or a delay in the tabulation of an election. The vulnerability to the attack of an election depends both on how ballot papers are designed and marked, and how the election is tallied.

2 Contribution

We first present a novel attack that can be executed on numerous existing voting systems with potentially far-reaching and serious implications. Then we identify the most important parameters of the attack and discuss how and to what extent the attack can be mitigated.

We hope that this paper will raise the awareness among researchers, governments, and other stakeholders. Short term, election owners must prepare plans and procedures to handle an attack. Modest improvements may also be applied to existing voting systems within current laws. Long term, each voting system should be studied carefully to see if it is possible to mitigate the attack in a way that is acceptable from a democratic point of view and election laws should be changed if needed. Due to the diversity of the details of voting systems, election schemes, legal frameworks, and democratic cultures, this is out of scope of this paper.

In this paper we focus on the mechanics of the attack at a high level. We do not consider the details of specific elections and voting systems to determine how vulnerable they are to the attack, assess the strategic value of carrying out the attack, and the threat model. Legal and political aspects are also out of scope of this paper, but we hope to inspire such research.

In an appendix we consider to what extent the attack can be applied to the particular voting systems of a handful of countries and informally propose a number of modest changes that could be deployed quickly to raise the cost to execute the attack and improve the chance to identify the perpetrator.

3 The Attack

We first observe that most voting systems provide a channel to voters that not only allow them to express their voting intents, but also to send arbitrary information through the voting system. More precisely, given a piece of information, one or more voters can use their right to vote to encode the information into the output of the voting system. In this context, the *output consists not only of the tally of the election, but also of all auxiliary information that is published or otherwise available*, e.g., the number of invalid votes and in what way they are invalid. Depending on how restricted the access to different parts of the output is, the attack is more or less feasible. In Sect. 4 we discuss several examples of how the information can be encoded depending on the specifics of the voting system.

Then we assume that the adversary has access to sensitive information that must not be published by the election authority. Secret information is clearly sensitive such as information published by WikiLeaks, but other information which is not particularly secret may also be sensitive. In Sect. 3.1 we consider different types of sensitive information.

We also assume that the adversary is able to publish information on "the Internet" in the sense that the data is made broadly available and can not be deleted. Today this is a very mild assumption due to the plethora of forums and servers that store information for free without deep authentication of the users.

Throughout we write $\mathsf{Enc}(k, m)$ to denote the encryption of a message m with a secret key k, and we denote by m the sensitive data. An example of a suitable cryptosystem is AES. We denote by H a cryptographic hash function such as SHA-3 that compresses an arbitrarily long input to a short digest. The basic attack proceeds as follows:

1. The adversary forms a ciphertext $c = \mathsf{Enc}(k, m)$ using the sensitive information m and a randomly chosen secret key k, and publishes it on the Internet anonymously.
2. She uses corrupted voters to submit votes that encode the secret key k in such a way that it can be easily derived from the output of the election after tallying.
3. She anonymously informs the relevant authorities, and possibly media or other parties, that if the result is tallied, then the sensitive data m will be published by the election authority.
4. She makes sure that her claim is credible, e.g., by revealing parts of the sensitive information to the owner of the election and chosen government agencies and media.

Example 1. Suppose that the attacker has access to Snowden's complete information m and consider an election that allows write-in votes, and that the contents of all votes are reported in the final result. Here the adversary picks a random party name p, hashes it to form the key $k = \mathsf{H}(p)$, encrypts the sensitive data to form the ciphertext $c = \mathsf{Enc}(k, m)$ which is published on the Internet. Then it submits p using a write-in vote. Then it informs the election authority and chosen media. If the election is tallied, then p appears in the result, $k = \mathsf{H}(p)$ can be computed, and the sensitive information $m = \mathsf{Dec}(k, c)$ is disclosed.

We stress that the attack is easy to execute completely anonymously, since the ciphertext c can be published through any electronic channel and the voting system itself provides privacy to the attacker when k is encoded into the votes.

The attack puts the owner of the election, e.g., a national election authority, in a situation where they de-facto become fully responsible for publishing the sensitive information and this is known by a wide audience. This immediately spawns a number of questions that demands answers such as:

- How vulnerable is a given system to the attack?
- What can we do to counter the attack?
- Is it legal to tally, or conversely refuse to tally, and can tallying be delayed?
- Should the election authority tally unconditionally?
- Who is politically and legally responsible for publishing the information?
- Can individuals or organizations demand damages for disclosed information?

3.1 Types of Sensitive Information

Before we consider how the attack differs from extortion we give a number of examples of sensitive information and discuss how the type of information influences the characteristics of the attack.

State Secrets. Imagine that a disgruntled officer in the military, or arms industry, decides to execute the attack. The obvious real world example is somebody like Edward Snowden, but with a more sinister agenda. The sensitive data may be worth billions and threaten the lives of many people if it is leaked.

The motivation may be political to punish the establishment, or at a national level to punish a foreign state. In the latter case, it may be clear that the attacker has no intention to leak the information, i.e., the goal is specifically to stop or delay the election.

We can even imagine an attack that is intended to look like an attack by an insider, but which in reality is an attack by a corrupt state. It is not far-fetched that elements of a country like USA or Russia sacrifices a measured amount of sensitive information and manufactures an insider attack on their own, or a foreign country, for political purposes. Consider the political pressure these countries can exert on small states to delay an election if needed.

The motive could also be economical. We would expect that the stock market reacts quickly if the tallying of the election is delayed. Trading on movements on the stock market in a covert way is not difficult and could result in huge revenues. A single individual with access to sensitive information and plausible deniability could today with little risk of detection execute this attack in several countries.

Private Information About Voters in the Election. Suppose that the election, or part of it, is performed using an electronic voting system. Due to lack of analysis and poor understanding of cryptography and computer security, several such systems have been broken [3,4,6].

Consider a political activist that has repeatedly pointed out vulnerabilities and tried to convince the authorities to not use the electronic voting system, and that she in despair decides to grab the secret key, or the votes of many voters, and use it as the sensitive information in the attack. Note how this differs from simply proving knowledge of the secret key where the government could dismiss the complaint with various explanations. Here the election cannot be tallied (as planned) and still preserve the privacy of voters.

We stress that the attacker has no intention to leak the information and has no incitement to claim otherwise. The goal is in fact to protect the privacy of voters.

This would of course be illegal, but it is also a form of whistleblowing on an election authority that ignores valid criticism through legitimate channels. Thus, we expect that many citizens would side with the attacker.

Information about how voters cast their votes could also be collected using something as simple as covertly filming voters in the polling station. The attacker would then cast her vote among the last in the election.

Illegal Information. Recall that the key feature of the attack is not that the sensitive information is secret, but that the election authority becomes responsible for publishing it. There are several examples of information that is sometimes publicly available, but not in aggregated form that allows, e.g., searching, and such information can be very difficult to collect without detection.

One example is an attacker that holds a large catalogue of child pornography. Publishing this information would not only be illegal, it could also seriously harm many children and people emotionally and constitute defamation leading to lawsuits.

Another example is sensitive user data from, e.g., forums, social media, infidelity websites, and perhaps more seriously, medical journals. Disclosing medical journals is not only problematic because it violates the privacy of people, it can cause people to lose their jobs and insurance policies. In the case of medical journals the goal could be to force the government to take action to improve the privacy properties of systems to protect the citizens.

In both latter examples, it could be clear that the attacker has no malicious intent and no intention of publishing the data on her own.

3.2 Is This Simply a Form of Extortion?

One may object that the attack is simply a form of extortion aiming to disrupt an election, i.e., the attacker could just as well simply explain that if the election is tallied, then it will publish the sensitive information. However, there are prominent features of the attack that distinguishes it from extortion.

An extortionist must convince the victim that the threat is credible, i.e., that she is willing to publish the data unless the victim stops the election. This is not the case in our attack. As illustrated in the examples above, it can be clear that the attacker has no intention to publish the data.

An extortionist can also change its mind. Thus, it is meaningful to negotiate with her and if she is captured in time, then the attack can be stopped. In our attack on the other hand, not even the attacker can stop the attack after it has been set in motion.

We believe that the distinction is of fundamental importance and changes the way governments can, and should, respond.

4 Encoding Data into the Output of the Election

A closer look at a typical voting system reveals that the bandwidth from the attacker to the output of the election is large. Below we give a non-exhaustive list of ways to encode information, but note that these may be combined if available. An additional factor is who is given access to the information and this is discussed in the next section.

4.1 Write-In Votes

There are two types of votes that are sometimes called write-in, but are quite different in our setting. Both assume that the voter can use a blank ballot and simply write on it the name of their favorite candidate.

Type I assumes that the candidate has been registered in advance, so in the election result such a write-in vote would be indistinguishable from votes cast using pre-printed ballots. A narrow channel of information is given by such ballots if available to the observers, since the candidate name may, e.g., be positioned differently on the ballot paper to encode information, but the ballot is difficult to spot even given access to the tallying.

Type II allows the voter to write anything on a blank ballot, and as long as it can be interpreted as something meaningful when it appears in the election result. This can be used directly to execute the attack if the vote is available to the observers, since the voter can simply write the secret key k used to encrypt the sensitive information as the candidate name. To make sure that the key seems meaningful the attacker can first come up with a randomly chosen name p and hash it to derive the actual secret key k as explained in Example 1.

4.2 Invalid Votes

In our setting invalid votes can be viewed as a form of write-in votes, but with limited information capacity. There are numerous ways to make a vote invalid and how they are processed depends on the type of election, so we can only give some examples to illustrate the problem. In all variations the observers must of course be able to record information about invalid votes.

In countries where detailed statistics about different types of invalid votes are disclosed they are truly a form of write-in votes of Type II in the eyes of the attacker.

In countries where envelopes are used and the observers may witness the counting, the attacker can simply put, or not put, post-it notes of different colors to encode a sequence of bits. Post-it notes stand out in the counting and are easy to spot.

4.3 Bundled Races

Countries that artificially bundle together multiple races create ballots that can be exploited by encoding the key as a list of components of a few bits, where each such component represents a choice in a race. For example, a bundled ballot with three races containing two candidates each can encode three bits. To be of use a larger number of races and/or multiple candidates is needed, but it is not merely the number of possible votes that is important. It is the size of the space of possibilities expected to remain unused by legitimate voters that determines the feasibility of the attack.

4.4 Ranked Elections

In ranked elections a single ballot is used with a large number of different possible votes corresponding to the possible permutations of the available candidates. Variable-basis representations of integers are easily converted to and from more natural representations of permutations and a key may be viewed as an integer, so an arbitrary key can be cast as a vote. These ballots cannot in general be tallied except by revealing a large part of each vote.

4.5 Supporting Evidence

Most voting systems have embedded features for auditing. The auxiliary information provided for auditing can provide a channel for the attacker even if the rest of the election output does not. Thus, the election output must be understood as consisting of all information and all physical artifacts resulting from the tallying of an election. A concrete example could be images of the ballots scanned by ballot scanning machines that could embed information using tiny markings, placement of text, or steganography.

4.6 Elections with a Fixed Set of Candidates

Even in single-seat elections where the election output consists only of the reported election results, the attack may be feasible, but at a higher cost to the adversary in terms of the needed number of corrupted voters. This is best explained by an example.

Example 2. Consider an election with three fixed candidates where the election result is reported per voting district among a large number of voting districts. Assume that the first two candidates get almost all of the votes so that the third candidates get zero votes in most voting districts.

Here an adversary that controls n voters throughout the voting districts that typically receives zero votes for the third candidates can encode bits by simply casting, or abstaining to cast, votes for the third candidate in those districts. A somewhat more expensive encoding with more cast votes can add error correction. A randomized encoding where zero and one are instead encoded as, say more or less than two votes, respectively, gives plausible denial for every individual vote. This may protect the attacker against sanitation of the result under governing laws, since votes of legitimate votes cannot be eliminated.

Note that the example does not require the attacker to register new candidates, but the attack is of course facilitated if this is possible, since it almost guarantees the existence of a candidate that can be expected to get very few votes. In some countries this is unlikely to be the case due to requirements for registering new parties or candidates.

The critical weakness of the election is how the result is reported. If there are only a few large voting districts, then the attack is infeasible.

4.7 Multiple Elections

It is important to understand that the above encodings can not only be combined with each other, but also for multiple elections. If the adversary is unable to encode the needed number of bits into one election, then she may still be able to encode a fraction of the bits in each election. The semantics are changed slightly with this approach since when a key is partially disclosed outside parties may be able to recover the remainder of the key using algorithmic methods.

4.8 Preventing Sanitation

An attacker may worry that authorities sanitize the output of the election in a controlled environment to mitigate the attack. This may be possible apriori depending on who has immediate access and governing laws. To circumvent any such procedures the attacker can use a proof of work to make sure that nobody can recover the key except after a certain suitable amount of time. A trivial way to accomplish this is to only encode part of the secret key k used for encryption. This means that a brute force search for the missing bits of the key is needed to decrypt. This variation shows that there is little value in attempting to sanitize the output of the election by trying to identify the encoding of k, since this can only be done long after the result must be published.

4.9 Access to the Output of the Election

A necessary condition for the attack to succeed is that the sensitive information is revealed to parties that must not have access to it. However, this is a not a black or white property. For example, national security and military secrets *should not* be disclosed to anybody, but it *must not* be disclosed to unfriendly foreign states. Similarly, child pornography can safely, and should be, disclosed to the Police, but must not fall into the hands of the general public.

Thus, to properly analyze the value of the attack and capabilities of the adversary in a given election, we need a comprehensive and detailed understanding of the voting system. This is important, since it is likely to be infeasible to unconditionally mitigate the attack for many election schemes.

The attack relies on the transfer of responsibility. Suppose election workers perform their duties in a closed room and the encoded key only appears in the room. Then if the key is disclosed we can argue that the election workers are culpable and not the election authority or government. This way of looking at the attack may be more or less realistic depending on the nature of the sensitive information.

5 Mitigating the Attack

The best we can hope to achieve may seem to be a voting system that outputs who won the election in a single seat race, or correspondly the distribution of seats in a multi-seat election, but a closer look at democratic systems shows that this is view is naive. The role of an election is not only to distribute seats, but also to communicate the voice of the voters in a broader sense such as fringe opinions and the geographic distribution of supporters of different candidates.

Thus, a more modest goal is that the voting system outputs the election result in the form it is currently output in most voting systems. This can clearly not be achieved if write-in votes are reported as part of the result without prior registration. The number of bundled races and cardinality of ranked elections combined with the number of candidates must also remain small. Furthermore, the result can only be reported for subsets of voters such that the number of votes for each candidate is large enough to hide encoded information in statistical noise provided by the votes of honest voters.

In addition to the above requirements, it must be ensured that no additional part of the output is leaked to the wrong parties. The specifics of this is inherently tied to particular elections, but we can make some general observations.

In elections with a voting envelope we can not allow the counting to be done in public. It is far too simple to insert arbitrary paper content into an envelope. However, it is probably fine to randomly select people from the general population to audit the counting and inform them to not leak any information except that they can dispute the counting.

Statistics about invalid votes should be kept to a minimum and reported in aggregate form and not per voting district or other small regions. The detailed statistics and information should be considered secret.

6 Variations

The attack could possibly be combined with a deliberate manipulation of the election result, and used to dissuade the authorities from publishing information that would indicate the manipulation. The key may be encoded not in the outcome, but in the evidence of the correctness of the outcome leading to a situation where the government is unable to allow a normal audit. Examples in an electronic voting system includes logging information such as timing of various events, flawed inputs, etc.

7 Future Work

There are two natural directions for future work. Firstly, understanding vulnerabilities and developing techniques and procedures that increase the cost of executing the attack is certainly possible, both for traditional and electronic voting systems.

There are also natural theoretical questions to be investigated. A function for which the adversary provides some of the inputs may be viewed as a channel in an information theoretical sense and we could demand that its capacity is low in the worst case, or average case, over the choice of the other inputs. Similarly to the discussion above, in a multiparty computation of the function, we must consider the output to be the complete view of the parties interested in the communicated information.

Acknowledgments. We thank a number of researchers that took part in our early discussions. We also thank civil servants and politicians in government organizations in several countries for their valuable feedback.

A Situation in Selected Countries

To make things more concrete we briefly discuss how serious the attack is in a handful of countries.

A.1 Australia

Many Australian elections allow each voter to rank many candidates, so each ballot may have about 100! different possibilities. Furthermore, tallying by Single Transferable Vote (STV) generally needs knowledge of most of each permutation—there is no easy way to split up the vote when tallying. Many Australian electoral authorities make complete voting data available on the web, for the very good reason that third parties may independently redo the count.

These sorts of voting systems are also vulnerable to a coercion attack sometimes called the “Italian attack”, in which voters are coerced into casting a particular voting pattern. The attack presented in this paper uses a similar feature, namely the large number of possible votes, but in a different way. Hence there is already some literature on how to compute a verifiable STV tally using cryptographic methods without revealing individual votes [2]. These mechanisms would also address the attack described in this paper, though they remain computationally intensive and not integrated into the Australian electoral process.

A.2 A.2 Estonia

A discussion related to the attack took place in Estonia in 2011 when an invalid i-vote was experienced for the first time in the history of Estonian i-voting system. The discussion is presented in [5] 3.1 Case: Invalid I-vote. Executive summary follows. One of the i-votes was registered invalid by the system during the tabulation phase of the Parliamentary Elections on March 6th, 2011.

The analysis of the system error logs showed that the invalid i-vote appeared to be correctly encrypted with the election public key. The reason behind the invalid i-vote could have been a bug in some of the components of the i-voting system, human mistake in the system setup or somebody could have intentionally cast an invalid i-vote (by implementing their own voting client or interfering with the existing one).

Only human mistake in the setup procedures could be excluded without decrypting the i-vote, so the National Electoral Committee (NEC) decided to decrypt the invalid i-vote and examine its contents in hopes to find out the root cause of the problem. The time window between the decision and the planned action gave an opportunity to consider invalid i-vote as a possible attack. If the attacker was aiming for publicity, then the simple scenario allowing manipulation would be used by the attacker himself to decoy the election officials to show whether the NEC – contrary to their claims – can find out who did cast the vote from the contents of the ballot.

If some more sophisticated technique to invalidate the ballot would have been applied, then the contents of the ballot could have been anything from the personal identification of the attacker or personal identification of someone not involved at all to a well formed ballot with an invalid candidate number.

After considering the matter of ballot secrecy and the possibility of an attack against i-voting as such, the NEC reached the conclusion that it would be better not to create a precedent of decrypting one i-vote separately from others. The decision from April 1st was reverted on April 8th.

A.3 A.3 Sweden

In Sweden the elections for parliament, county councils, and municipalities all take place at the same time, but using three distinct ballots and envelopes. Thus, it is not a bundled election. A voter picks a ballot paper with a pre-printed party name and a list of persons. He may make a single mark in front of one of the persons to increase her chances of getting a seat. This is called a "personröst" (person vote).

Votes are then counted and sieved for invalid votes at several levels and all counting is open for the public. The ballot papers are first taken out of their envelopes in the polling station by the election workers. Ballots that are deemed invalid are put back into their envelopes and put in a separate stack. There are exceptions, but broadly speaking a ballot is invalid if it is not formed as described above. The votes are then recounted by another authority before the final result is announced. During the first counting only party votes are counted and the person votes are ignored.

The voting system in Sweden has been reformed in several ways in preparation for the 2018 elections. Fortunately, a side effect of these changes is that the attack presented in this paper is harder to execute. Before the reform a voter could cast a write-in vote for a party or person. As of 2018 all parties and persons must be registered and acknowledge that they are willing to serve if they are elected.

We remark that parties such as “Kalleankapartiet” (Donald Duck party) would always receive a couple of votes and the results from the 2014 election are available at [1]. Although there are no longer any write-in votes (of Type II as defined in Sect. 4.1), an attacker can demand to see invalid votes and she could use post-it notes of multiple colors, corrupt a handful of voters and execute the attack in this way. There is also a fair number of fringe parties that only get a handful of votes and even more individuals listed for the parties that get even fewer votes. Thus, there is plenty of room to encode a key.

The system could be substantially hardened by replacing the public counting with counting in the presence of a randomly selected set of citizens and by not reporting results for parties that receive a small number of votes, or reporting them in aggregated form at a national level if the number of votes increases notably by doing this. Furthermore, a threshold could be introduced to register a party whereby it must be made plausible that it will receive, e.g., a few thousand votes. Such thresholds are already in place in several countries. A similar approach could be used for person votes.

References

1. Swedish Election Authority: Election results 2014. http://www.val.se
2. Benaloh, J., Moran, T., Naish, L., Ramchen, K., Teague, V.: Shuffle-sum: coercion-resistant verifiable tallying for STV voting. IEEE Trans. Inf. Forensics Secur. **4**(4), 685–698 (2009)
3. Halderman, J.A., Pereira, O. (eds.): 2012 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE 2012), Bellevue, WA, USA, 6–7 August 2012. USENIX Association (2012)
4. Halderman, J.A., Teague, V.: The New South Wales iVote system: security failures and verification flaws in a live online election. In: Haenni, R., Koenig, R.E., Wikström, D. (eds.) VOTELID 2015. LNCS, vol. 9269, pp. 35–53. Springer, Cham (2015). doi:10.1007/978-3-319-22270-7_3
5. Heiberg, S., Laud, P., Willemson, J.: The application of i-voting for Estonian parliamentary elections of 2011. In: Kiayias, A., Lipmaa, H. (eds.) Vote-ID 2011. LNCS, vol. 7187, pp. 208–223. Springer, Heidelberg (2012). doi:10.1007/978-3-642-32747-6_13
6. Khazaei, S., Terelius, B., Wikström, D.: Cryptanalysis of a universally verifiable efficient re-encryption mixnet. In Halderman and Pereira [3]