Web-Voting, Security and Cryptography

Recently E-voting, or voting using digital technologies, has often been in the news. It seems that many politicians are in favor of adopting E-voting (see eg. [1] for a recent example) but that most IT security experts and cryptographers are against it. By the name E-voting it is usually intended a traditional voting procedure aided by digital voting machines (also called Direct Recording Electronic, DRE). In this article we will discuss another form of digital voting, web-voting, that is voting through internet using a web-browser to cast a vote on a web-server.

We discuss web-voting starting from the characteristics of a voting process and discussing what we need to achieve to implement a web-voting system.

The traditional voting process

We start from a traditional voting process that can be divided in 4 steps (see fig. 1)

a) *Authentication*: Alice walks in a voting precinct and authenticates herself by showing some voting credentials; this step is public and verified by the officials present in the room; at the end of the authentication Alice is given a paper ballot to write her vote on.

b) *The Vote*: in a protected booth where she cannot be seen by anybody, Alice casts her vote by writing it with a pencil on the paper ballot; she then folds the paper ballot and puts it in the ballot box where all the ballots are mixed; since nobody can see what Alice writes and there are no marks on the paper ballots, Alice's vote is anonymous.

c) *Counting the Votes*: at the end of the voting time, the officials open the box containing the paper ballots and publicly count the votes; the results are then announced.

d) *Verification*: various types of verification are present or possible, most procedures are indeed public and overseen by representatives of competing parties: the opposite interests of the parties warrant for the first level of protection against fraud; recount is also possible if there is a presumption of fraud or error.

A Theoretical Voting Protocol

The requirements for a perfect voting protocol have been studied in details (see e.g. [2] and [3] for a review). Here we mention only the ones which are more evident and important for our discussion:

1. Unreusability (prevent double voting): no voter can vote more than once
2. Privacy: nobody can get any information about the voter's vote
3. Completeness: all valid votes should be counted correctly
4. Soundness: any invalid vote should not be counted
5. Eligibility: no one who is not allowed to vote, can vote
6. Fairness: nothing can affect the voting
7. Incoercibility and no-vote-selling: a voter cannot be coerced into casting a particular vote by a coercer nor can prove to a third person how she has voted.

We discuss now if and how a web-voting system can guarantee these requirements.

Web-voting, authentication and vote-selling

By web-voting we consider the very simple process of voting with a web-browser (either a common one like Internet Explorer, Firefox etc. or one specially developed for web-voting) and casting the vote on a web-server. For the moment we assume that the process is perfect and that there are no security problems in the software, network protocols etc.

Even in this theoretical situation, already at the first step (step a, the authentication) there is a problem. Indeed, since the voter is in a remote location, we cannot be sure that the voter is who she claims to be, unless we use some biometric authentication protocol. Without biometrics, Alice can sell or can be forced to sell her voting credentials to Eve without anybody realizing it.

Even assuming that biometrics is used to authenticate Alice, Alice and Eve can sit in front of the same PC, Alice can do the authentication and Eve can vote or check that Alice votes as she wants. If Alice wants to sell her vote, she can take a picture or make a movie of her voting and give it to Eve.

In any case, the remoteness of the voter makes requirement 7. impossible to fulfill for web-voting

(see e.g. [3] page 31). Notice that this is a difficult requirement also for traditional voting procedures and indeed there have been cases in political elections of pictures taken of paper ballots in the voting booth with mobile phones' cameras.

We are then forced to abandon requirement 7. and substitute requirement 5. with the weaker

5b. Weak-Eligibility: only eligible voters can get voting credentials from trusted authorities.

In practice this means that web-voting cannot be used in elections or polls where fraud by vote-selling or coercion is a concern, like political elections, whereas it can be appropriate for elections of associations' boards, users' polls etc.

Privacy and cryptography

After being authenticated, Alice casts her single vote in such a way as to maintain her privacy, that is the protocol must guarantee requirements 1. and 2. listed above. At first sight this seems to lead to a contradiction: we need to prevent double voting but at the same time guarantee that all voters are anonymous at the vote web-server. Here is where we first meet cryptography.

A simple, but not the only, solution to this problem is the following. At the end of the authentication step a), Alice is given one (and only one) digitally signed document (a *Digital Authorization*) by the authentication authority or server. The Digital Authorization is the equivalent of the blank official paper-ballot in traditional elections: it allows Alice to vote anonymously.

Then Alice presents the Digital Authorization to the vote web-server which checks the digital signature verifying that it has been made by the authentication authority/server, and also that this is the first time that it is presented to it for voting. If these two conditions are met, Alice is allowed to cast her vote on the vote web-server. To guarantee that Alice's vote is anonymous, the Digital Authorization does not contain any reference to her identity, indeed its content can be a random number.

To really guarantee Alice's privacy, it is necessary that nobody, included the authentication authority/server, can trace the Digital Authorization to the voter. Again cryptography can come to our rescue. For example there are various cryptographic protocols that allow to digitally sign a

document without knowing its contents, as the blind signature scheme [4]. Unfortunately the more secure these protocols are, the more cumbersome are to implement. Here we face for the first time a typical cryptographic usage dilemma: often the more theoretically secure a cryptographic protocol is, the more difficult is for the average user to follow it correctly without really understanding what she is doing. This is a real problem and in many cases it is preferable to adopt different, non cryptographic approaches to obtain the same goals, like strict separation of duties and other procedural steps.

From our analysis it follows that if the authentication is done on-line with a web-server, this must be a different server than the vote web-server. This is just because the authentication server knows who the voters are, whereas the vote server must not know who is voting. Another comment is that some researchers have taken a different approach to web-voting by keeping together identity of the voter and vote until the third step c), counting the votes, when identity of the voters and votes are decoupled, the votes are mixed and then counted. In the case of web-voting, we believe this procedure to be difficult to implement correctly, and to make more difficult to guarantee the privacy of voters. Indeed once you have some information it is not easy to delete it, whereas if you do not have it at all from the beginning the problem of information disclosure cannot arise.

Privacy and anonymous networks

But to guarantee the privacy of the voter when connected to the vote web-server we also need to consider the network connection to it. For example the IP address of Alice's PC must be concealed from the vote web-server. In theory it is possible to design perfect anonymous networks [5,6], in practice this is not so easy. Currently the Onion Router (TOR) [7] is a world network allowing web navigation in anonymous way at a cost of a slow down of the connection due to the extra path and encryption involved. Still TOR is not perfect, for example it is well known that the exit server from the TOR network knows the IP address of the voter and of the vote web-server. A rogue exit server from the TOR network can easily trace the connections instead of make them anonymous, but it cannot learn or modify the vote if the connection is encrypted with SSL/TLS.

But this is not all: standard browsers leak information about themselves and the voter's PC. Indeed

all standard browsers send basic information about themselves, the version of the operating system etc., to the web-server. Usually this leak of information is not crucial, but in some cases it could still give some hints on who the voter is. Today the only way to remove this information is to prepare a custom-made web-browser only for web-voting. On the other hand a voter must trust that a custom-made web-browser built for web-voting does not record, modify or divulge her vote.

Counting the votes

Once the vote web-server receives a vote, it has to store it securely until the time all votes are counted. Measures must be implemented to prevent manipulation of the votes while stored. The basic measure is to encrypt each vote with the public key of the electoral committee, in this way the votes can be decrypted only with the corresponding private key which must be kept safe until the moment of counting the votes. Again cryptographic procedures should be put in place to assure that only votes correctly received by the vote web-server are accepted and counted.

When data is stored in a computer, it is usually allocated sequentially. This means that votes are usually recorded in the order as they are casted. Instead the box in which the paper-ballots are put effectively mixes them up independently of their ingress order. There is indeed the risk of loss of privacy by correlating the order of authentication of the voters, which is a public information, with the order in which the votes are recorded. This risk is often reduced by the anonymous networks used to access the vote web-server which usually introduce delays and change slightly the order of connections. We need anyway to shuffle the encrypted votes to guarantee the privacy. There is an extensive cryptographic literature on algorithms to securely shuffle data which are based on the concept of *mixnet* first formulated in [6]. The basic idea is to make a random permutation of the data preserving its meaning but shuffling completely its origin.

Finally the voting authorities receive the encrypted anonymous votes, decrypt them and count the votes.

Vote Verification

In traditional voting Alice cannot directly verify that her vote has been counted correctly. Instead

she trusts the vote officials and the parties' representatives for the integrity and the correctness of the procedure. In case of complaints, assuming that a correct chain of custody of the paper ballots is in place, all or a part of the paper ballots can be recounted.

In digital voting most of the human control and verification is just not possible, all is done by machines that we should trust. But this means trusting that the people who designed and built the hardware, designed and wrote the software, installed and configured the machines, all have done everything correctly. Moreover we know very well that absolute security does not exist and that any complex digital system has bugs. So as voters we cannot trust a fully digital voting system to be correct.

But in digital voting there is a new possibility: to give to each voter a receipt that allows her to verify that the vote has been counted correctly (a final diagram of a possible web-voting protocol is in fig. 2). It is not easy to build a vote receipt which guarantees the voter that her vote has been counted correctly: we need to create a receipt which is unique for each vote, which cannot be repeated identical for different voters and which does not contain any reference to who the voter is. To build such receipts we need again cryptography typically by using one-way hash or zero-knowledge protocols.

Another desired property of the receipt is that it should be difficult to use as proof of vote to someone else. Even if we said that it is impossible for web-voting to prevent coercion and vote-selling due to the physical location of the voter, a receipt which cannot be used to prove how one has voted would make fraud more difficult. Receipts of this kind are called *Secret Receipts* and are difficult to create. Extensive work has been done on cryptographic protocols for secret receipts (see e.g. [3] for a review) but the main problem with all these protocols is that they require the active role of the voter: one thing is to check on a bulletin board if the vote appears with the correct receipt code next to it, another is to actively take part in a complex procedure without comprehending its meaning.

Real security of IT web-voting systems

Up to now we have assumed that IT systems are perfect without bugs or security issues. But in

reality things are quite different: web-servers and web-browsers are known to be software components susceptible to many security issues. If it is possible to obtain some assurance that servers are configured correctly and kept up-to-date, the current situation with the security of voters' PC is very bad. For example at a panel at the World Economic Forum in Davos in January 2007 [8], Vint Cerf, one of the fathers of internet, said that up to one quarter of the PCs connected to internet, which makes for 100 to 150 millions, are controlled by cyber criminals in the so-called botnets without the owner knowing it. Since these criminals remotely fully control the PCs, they can intercept and modify any activity done by the user of the PC, including a web-vote. In this case any kind of security we can devise for the overall web-voting system is at the end defeated at the voter's premises.

Thus it is important for web-voting that we adopt protocols which allow human independent verification of the results and are not purely based on digital procedures.

User interface, cryptography and concluding remarks

We have seen that web-voting is difficult, has some intrinsic limitations and real security issues. Research and development is still ongoing to find new protocols which guarantee reasonably secure, trusted and usable implementations. Cryptography must be seen as a tool, a building block necessary but not sufficient to build a web-voting system, also because cryptography is often difficult to use. We have seen that one of the most crucial points in a web-voting protocol is to give to the voter sufficient trust that the procedure is correct and her vote is counted. It cannot be expected that average users understand the subtleties of today's cryptography, nor they can be asked to follow complex procedures (see for example the recommendations in [9] in the case of DREs).

Since we cannot trust unconditionally digital systems to guarantee the correctness of a web-voting protocol, the problems of user interface and user trust are very important. Today it is somehow more important to renounce to some properties of a web-voting protocol, for example giving out a plain receipt which can be easily used for vote-selling (which in any case is possible), but make the voter understand and trust the process. This limits the applicability of web-voting protocols, but it makes it easier to implement and to manage the associated risks.

On the other side, web-voting can play a really important role in the democracy of our life. Very often it happens that we cannot participate in a decision because we are not physically present at the moment of the election or vote. Web-voting, associated with on-line webcasting of the meeting, can allow a much larger participation in many decision processes. For example a prototype for a web-voting system built according to what we have discussed here has been built [10] and has been tested in the election of the board of a professional association. The typical number of participants in person to these elections is between 10 and 20% of the members. In the election done using the web-voting system, 20% of the members participated in person and another 40% of the members voted from their offices or home for a total direct participation of 60% of the members of the association. We believe this to show the potential applications of web-voting.

Andrea Pasquinucci

PhD CISA CISSP

http://www.ucci.it/

References

[1] Announcement by the UK Department for Constitutional Affairs on January 29, 2007; see http://www.gnn.gov.uk/environment/fullDetail.asp?ReleaseID=260071&NewsAreaID=2

[2] E.Gerck, *Voting System Requirements*, The Bell, 2001

[3] B.Adida, *Advances in Cryptographic Voting Systems*, PhD thesis, 2006, MIT

[4] D.Chaum, *Security without identification; transaction systems to make big brother obsolete*, Comm. ACM 28(19) 1985, 1030

[5] M.K.Reiter, A.D.Rubin, *Anonymous web transactions with Crowds*, Comm. ACM 42(2) 1999 32

[6] D.Chaum, *Untraceable electronic mail, return address and digital pseudonyms*, Comm. ACM 24(2) 1981, 84

[7] R.Dingledine, N.Mathewson, P.Syverson, *Tor: the second-generation onion router*, http://tor.eff.org/

[8] http://news.bbc.co.uk/2/hi/business/6298641.stm

[9] R.G.Saltman, *Independent Verification: essential actions to assure integrity in the voting process*, preliminary review for NIST, 2006

[10] http://eballot.ucci.it/

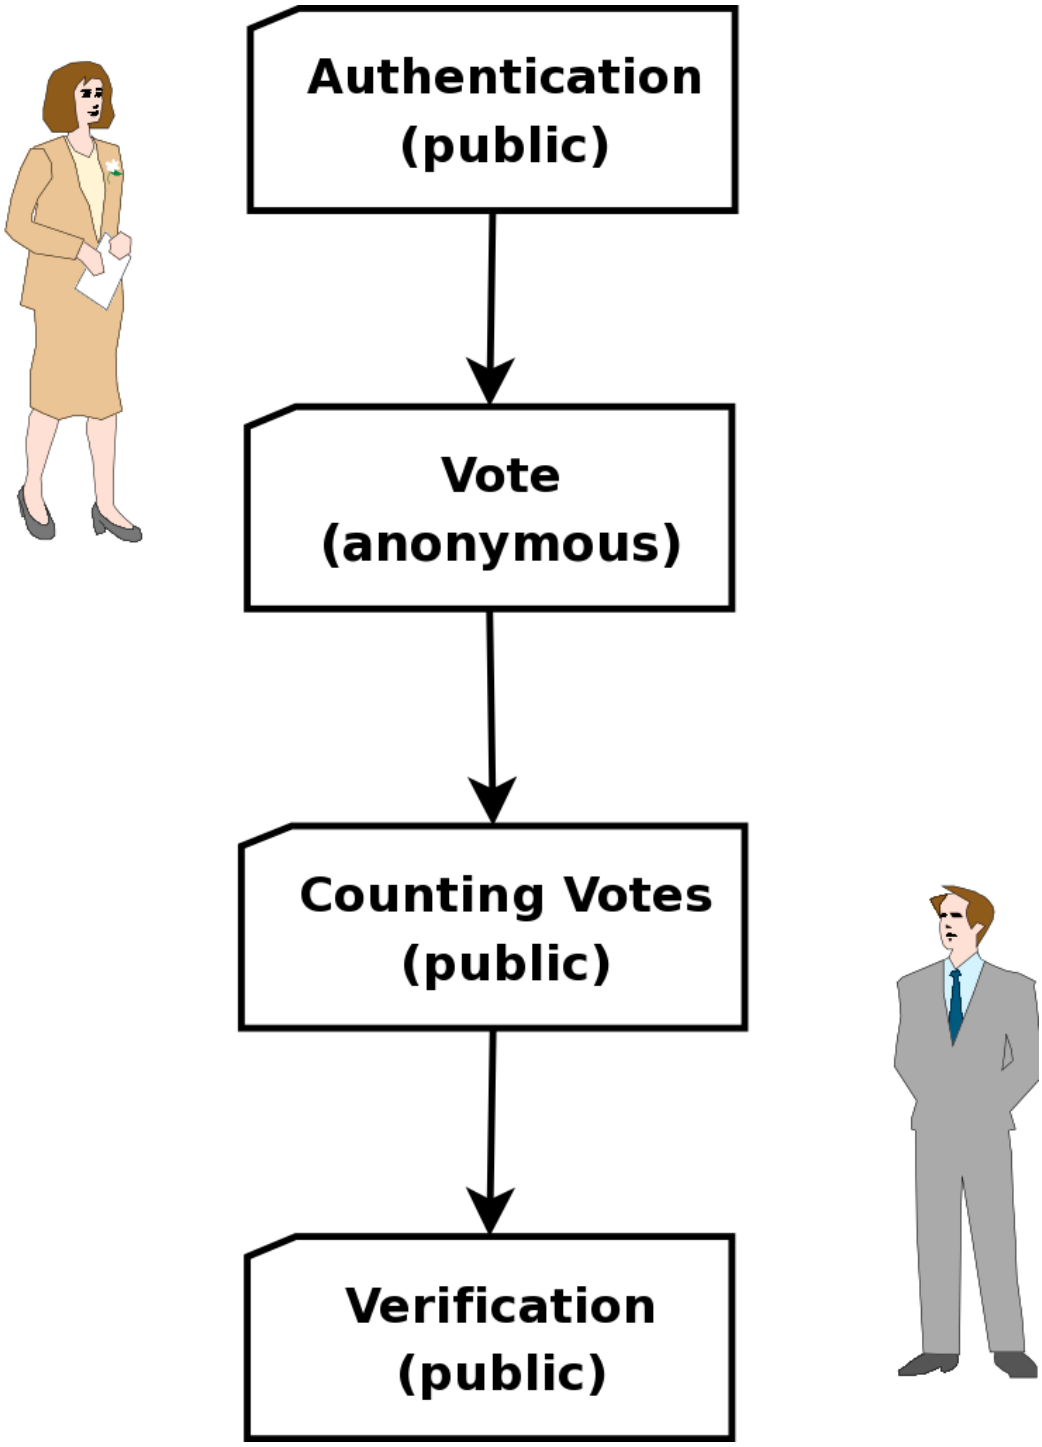

Fig.1 The 4 main steps of a theoretical voting protocol

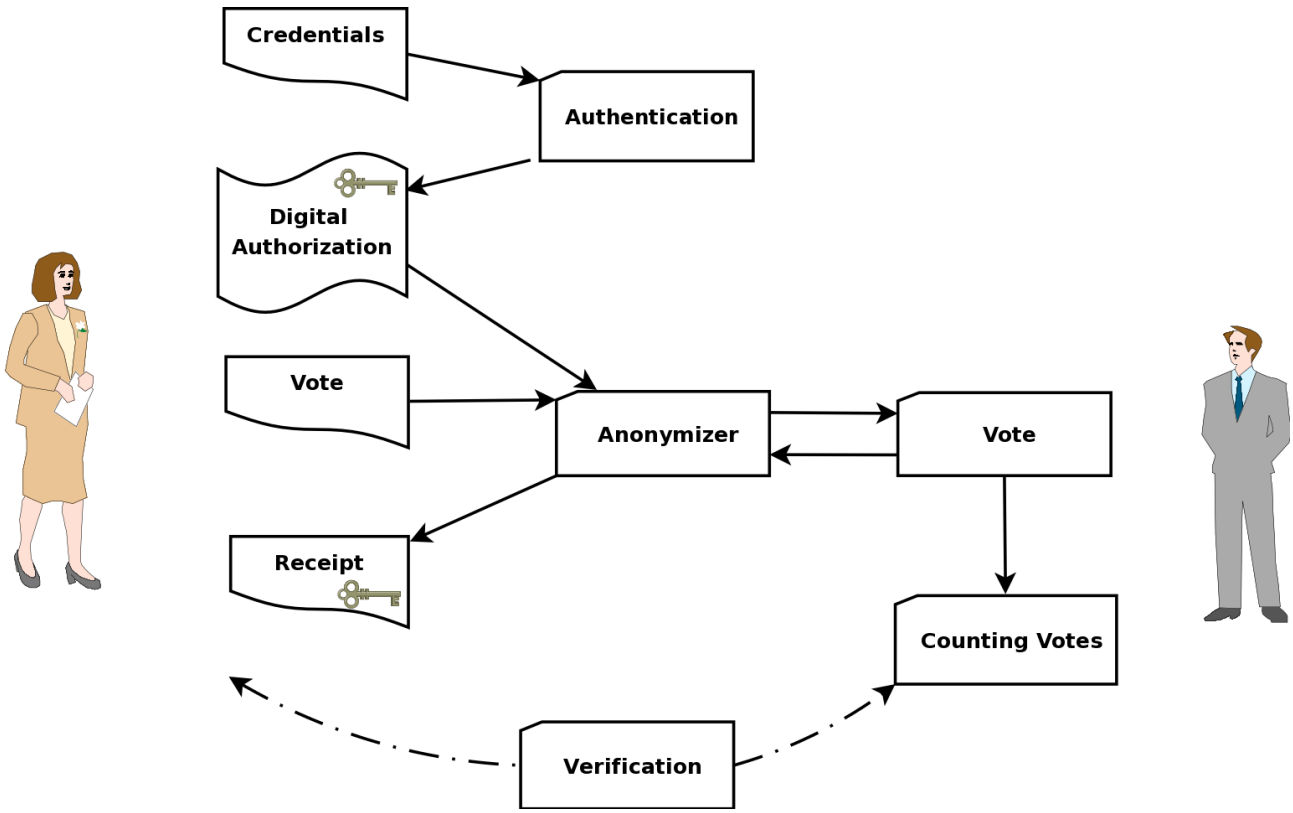

Fig.2 A theoretical web-voting protocol