A Discussion of Cryptographic Protocols for Electronic Voting

Lelia Barlow
December 2003

Introduction

In this paper, we will consider several proposed ideas for improving electronic voting systems. Some, but not all, of these ideas utilize cryptographic protocols.

We will compare and contrast these ideas, based primarily on the following factors: accuracy, privacy, and usability. These factors are discussed below.

Accuracy

After some preliminary research into the subject of electronic voting [BARL2003], it became apparent that there may be reason to mistrust current electronic voting systems. Further, as I studied a variety of voting systems, I realized I had little confidence that any of them (even paper ballots) would be guaranteed to produce a perfectly accurate tally.

In my previous paper [BARL2003], I suggested the following definition: "An accurate voting system counts all valid votes with minimal processing error such that the intent of eligible voters is reflected in the final tally."

I believe it is important for a voting system to (a) minimize error, (b) count all valid votes, and (c) preserve the intent of eligible voters in the final tally. At least in theory, I believe that today's electronic voting systems are capable of counting votes with minimal processing error. Assuming that all parties involved in the process of creating an electronic voting machine perform ethically, and that enough time, money, and effort is expended to do the job correctly, it should be possible to create a system that processes each entry that was received by the system. However, even making those assumptions, it is not quite that simple. There may be many reasons that a valid entry is not received by the system. For example, voters may be confused by the systems and may not enter their selections into the machine they way they intended. The system may "crash" before reporting its results. An eligible voter may be denied access to voting machines during the election interval.

As far as I have been able to determine, the process for deciding the validity of a vote is subject to interpretation by human beings. Election officials are responsible for checking the voter's credentials and allowing them access to the voting machines if their credentials appear to be valid. An effort is made to protect the voting machines using physical security techniques such as locking mechanisms and guards. I will assume for the moment that the physical security cannot be breached, that election officials perform ethically, and that it is possible for election officials to determine quickly and easily whether an individual is eligible to vote based on credentials. Thus, my concern becomes how to demonstrate that a vote is valid once it is cast and the voter walks away from the polling location. In other words, the vote-counter needs a way to determine whether a cast vote is valid.

Let's say that a vote is valid if (a) the ballot was cast by an eligible voter, (b) the eligible voter did not vote more than once in the current election, and (c) the ballot was not altered in transit, between the realm of the voter and the realm of the vote-counter. Thus, the authenticity and integrity of the ballot are important in determining whether a vote is valid.

All this considered, I believe it is most important that the intent of eligible voters is reflected in the final tally. Indeed, if there is a way for each eligible voter to verify that his or her vote is counted correctly, then there is a way to check whether the electronic voting system counted the valid votes without error.

The potential exists for electronic voting systems to provide a method for eligible voters to check whether their votes were counted correctly in the final tally. If votes are stored electronically, it is much easier to publish these votes. (Compare this to publishing votes counted from paper ballots.)

If votes are published, then it becomes possible for everyone to see the results of an election. As a naive example, consider the following list of results from an election:

Jane	Yes
Joe	Yes
Jenny	No
John	Yes

Jane can check the list for her name and see that her vote was counted as a "Yes." Jane can also verify that a majority of voters voted "Yes," which agrees with the announced election results.

Of course, this obviously brings up the issue of voter privacy. How can Jane verify her vote anonymously?

Privacy

Each paper, article, and discussion I have read describes privacy as the most important aspect of a voting system. I completely agree. It is essential that votes be cast in secret, without a provable link between the voter and the cast ballot.

Without anonymity in the voting process, voters may be unwilling to reveal their real intentions. They may fear retribution if they do not vote with the majority, or even if their intentions are unpopular with a particular group. There may be financial incentive for voters to vote a certain way. The concept of "vote-buying" has existed for a long time.

I would urge readers to consider how their privacy is protected when they use current voting systems. For example, if a voter walks into a polling location to cast her vote, it is possible for technology and human beings to record information about this visit. Video cameras could be used to gather information about who voted at this location and when they were there. When the voter shows proof of eligibility to vote in order to gain access to voting machinery, the election official could remember this information. Poll workers or other voters could also gather information about voters who visit the polling location. This information could then be correlated with the results entered into voting machinery, and could give significant clues about the choices made by a voter.

David Chaum, in [CHAUM] states,

> "Because of current surveillance technology, such as sensors like miniature cameras and emanations receivers, as well as memory and transmitters, the confidentiality of what transpires in voting booths cannot in practice be held to any absolute standard."

And,

> "Moreover, technical provision of privacy has its limits in voting, some examples of which are as follows: most US voter addresses and party affiliations are a matter of public record; the more help a device gives a voter the harder it is to keep it from learning who they vote for (though here the devices need not be able to retain data between votes); even the "gold standard" of voting systems, manual paper ballots, is subject to marking or ballot-number recording and automatically captures fingerprints; and theoretical limits are believed generally to force a choice in cryptographic systems between unconditional integrity and unconditional privacy."

Perhaps it is not necessary for a voter to visit a polling location in order to cast a vote. Couldn't we just allow everyone to vote on the Internet? It turns out that this is an even more difficult problem. Cryptographer Bruce Schneier claims that, "A secure Internet voting system is theoretically possible, but it would be the first secure networked application ever created in the history of computers." Computer scientist Peter Neumann warns, "the Internet is not safe for elections, due to its vast potential for disruption by viruses, denial-of-service flooding, spoofing, and other commonplace malicious interventions." [MERC2002]

Voters in the state of Oregon vote by mail. Paper ballots are sent through the postal service to eligible voters. The idea is that voters make their selections in private, then voters either mail in the ballots or drop the ballots in secure ballot boxes. Note that the voter, in the privacy of her home, can show her ballot to whomever she chooses. The voter could take her voted or un-voted ballot anywhere she chooses, and show it to whomever she chooses. In fact, in an extreme example, someone could be standing over the voter while she votes, pointing a gun at her head to ensure that she votes in a certain way. The voter could also copy the ballot – digital scanners are inexpensive and easy to acquire. A good copy might enable the voter to vote more than one ballot. A copy could also be used to prove that the voter made certain selections, so that the voter could sell her vote. A clever voter could make a copy of the ballot before filling it in, then fill in the ballot multiple ways and sell her vote multiple times.

I have yet to discover any voting method that completely and totally satisfies the privacy requirements. That said, I'm not sure that such a method would be at all practical. Somehow, the eligibility of the voter must be determined prior to every election, since voter eligibility may change over time. I believe we can all agree that a voting system falls apart if anyone and everyone can vote as often as they would like, with no oversight. However, when we allow eligibility to be verified it is likely that we will give up some degree of privacy.

Therefore, I believe the important question to ask is: how much privacy is "enough" privacy?

Usability

The literature describes what voters do when they use an electronic voting system. I am interested in going a step beyond "Joe pushes a button" to ask questions such as "How does Joe know which button to push?" and "What happens if Joe (accidentally or deliberately) fails to push the correct button?"

I believe that any time we attempt to create a system that involves human users, we must consider usability. Let's consider three "flavors" of usability: (1) the impact of human interaction on the security of the system, (2) the simplicity or complexity of a system as perceived by its human users, and (3) the time and effort a human user must expend in order to use the system.

The most technically sophisticated system can fail due to interaction between the technology and its human users. Frequently, human beings are the weak link in an otherwise strong system. Any time a human being is allowed to touch a system, the system may be exposed to an attack. Therefore, it is necessary to make it as difficult as possible for actions by human users (whether accidental or deliberate) to weaken the overall security of the system.

Furthermore, voting systems will be used by a variety of people. We cannot assume that these people think in the same ways as scientists and engineers. (In my previous work experience, it was surprising to me to see how often engineers made this assumption – without even realizing they had made an assumption!) Recall that in the 2000 election a number of American voters were confused by "butterfly ballots." Therefore, it is important to keep the human interactions with technology as simple and straightforward as possible.

We must also consider the time and effort voters are willing to expend in order to cast a vote. For example, if it requires two hours, nine transactions, specialized knowledge, buying new equipment or upgrading existing equipment, and repeated inquiries to technical support, it is unlikely that most voters would go to the trouble of voting. The barrier to the average voter can be made too high for the system to be practical.

Problems and Solutions

We now see that creating an idealized voting system is very difficult. We must have a system that produces an accurate result, while protecting the privacy of voters. All eligible voters should be able to use the system without excessive inconvenience, and without compromising the security of the system.

Fundamentally, I see this as an interesting engineering design problem. It may not be possible to achieve the theoretical maximum for all desired characteristics of the system simultaneously. Yet, I believe it is possible to improve on the current methods for voting.

Philosophically speaking, a methodology is an improvement only if it does not create any new problems while trying to solve existing problems. For example, if, through an application of an interesting new technology, we end up with an overall system that is less accurate or trustworthy or useable than, say, voting with paper ballots, then we can not consider this an improvement. Technology has great potential for improving many problems that human beings face; however, it also has the potential to generate new problems. Simply throwing technology at our problems is not necessarily a good solution.

David Dill, a professor of Computer Science at Stanford University and an active participant in electronic voting debates, said in a recent interview [PITT2003],

> "...being an engineer involves making choices about the appropriate use of technology. It is not using the highest tech solution to every problem, whether it's appropriate or not. It's focused on solving the problem by the best means that are available. The best engineers will use the best means that are available even if they don't involve any significant technology at all. I think it's the responsibility of everybody in technology to weigh in with their opinions about the appropriate use of technology and the inappropriate use of technology."

With this in mind, we begin a discussion of some proposed solutions.

Proposed Cryptographic Protocols for Electronic Voting

There are many and various approaches for improving electronic voting systems. I will choose only a few that I think are interesting. I would enjoy doing a more complete literature search as time allows, and perhaps in the course of this search I will uncover more ideas that are noteworthy.

The Mercuri Method

I would like to begin this discussion with a relatively low-tech idea. Rebecca Mercuri created a method to physically verify voters' ballots, which was first published in 1992 and later incorporated into her doctoral dissertation in 2000. A description also appeared in [MERC2002]. Dr. Mercuri has deliberately not patented this concept, so it could be freely incorporated into election systems. [MERC2003]

The Mercuri Method requires that the electronic voting system print a paper ballot containing the voter's selections. The voter examines this ballot for correctness through a glass or screen, and then the ballot is mechanically deposited into a sealed ballot box. This eliminates any chance that the ballot will be removed from the polling location. If the paper ballot does not match the voter's intended choices displayed on the electronic voting system screen, an election official can be summoned. The official can void the ballot, and the voter has another opportunity to vote. After the election, the tally produced by the electronic voting system provides a fast preliminary result. However, official results must be obtained by counting the paper ballots. These paper ballots could be optically scanned by marksense machines, or counted by hand. (See Figure 1, below.)

Dr. Mercuri argues that this type of system is cost effective because ballots are computer-generated. This means blank ballots no longer need to be prepared in advance.

I think that this method is a good first step toward improving current electronic voting systems. As far as accuracy, the method certainly does no worse than current methods (such as marksense scanning). Although no provision is made for voters to verify that their intended choices are included in the final tally, voters can at least verify that their intended choices are printed correctly on paper ballots. The paper ballots provide a physical record of the election. In case of a dispute or recount, the paper ballots can be inspected and processed again, perhaps by a neutral third party.

Except in the case where the electronic voting machine incorrectly produces a printed ballot and an election official is called to intervene, the privacy of the voter is maintained at currently acceptable levels.

The user needs to perform the extra step of verifying the printed ballot, but neither should this be very difficult to do, nor should it be inconvenient. The voter is not able to tamper with the printed ballot because it is inaccessible to the voter. Perhaps the voter could discover a way to get the machine to print an incorrect ballot, which could temporarily distract an election official from other responsibilities, but this seems unlikely to cause significant problems.

However, it should be noted that this method essentially downgrades the function of the electronic voting system from "official vote counter" to "preliminary result generator." The electronic voting system merely generates a ballot, guides a voter as he fills in the ballot (observing any mistakes such as over-voting or under-voting), prints out an official copy, and reports unofficial results.

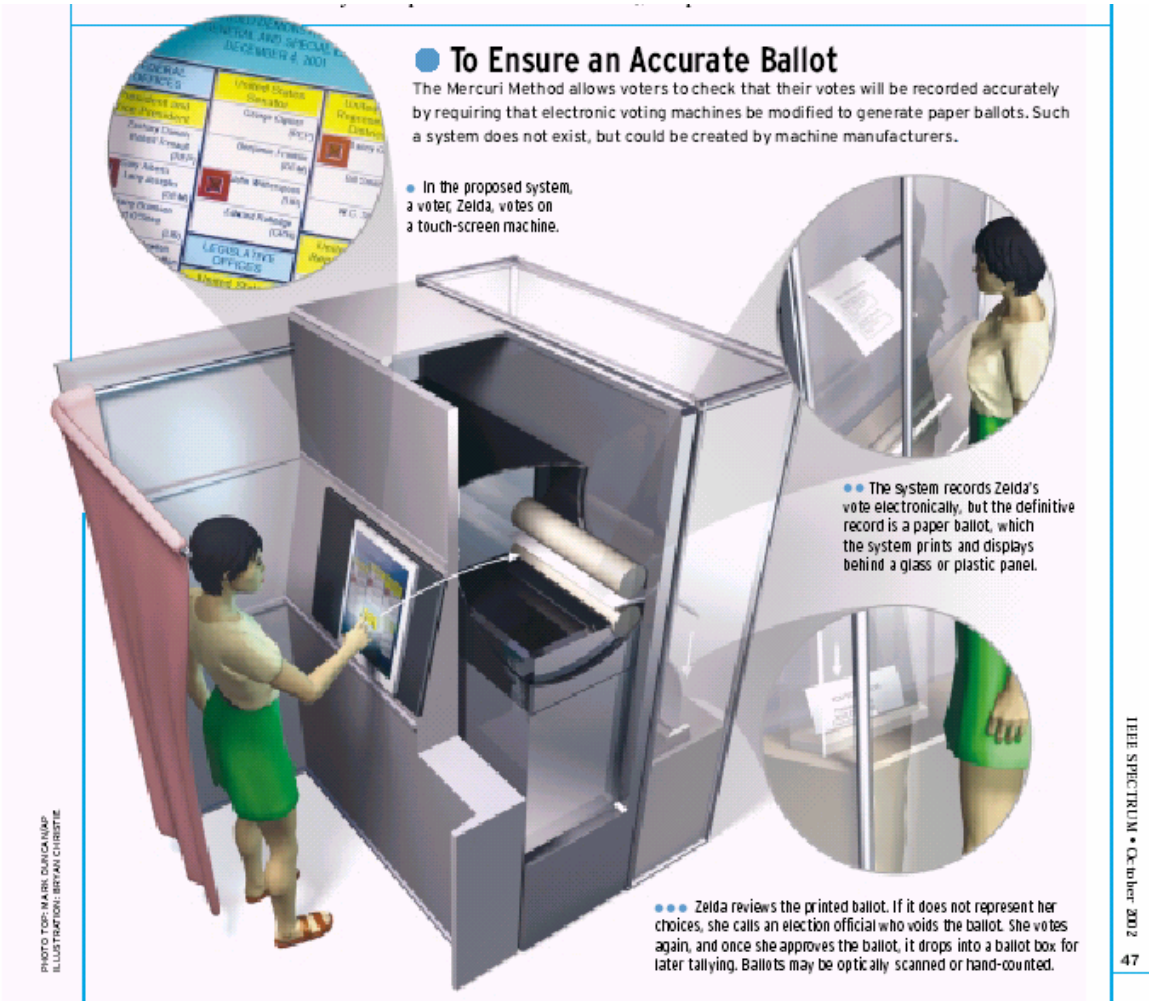

Figure 1. Illustration of The Mercuri Method
[Source:[MERC2002] Mercuri, Rebecca. "A Better Ballot Box." IEEE Spectrum, October 2002.]

Blind Signatures

Birgit Pfitzmann, in [PFIT1996] offers the following definition, "Blind signatures are signatures that the signer issues without being able to see the message she signs." It is reasonable to imagine that blind signatures may be used to allow a voter to cast a vote anonymously, while still providing authentication.

A protocol for voting with blind signatures is described in [SCHN1996]. This protocol can be summarized as follows.

> Every voter produces, say, ten sets of messages. Each set contains a valid vote for each possible outcome of the election. For example, if the election decides a simple yes or no question, each set contains two messages: the messages "yes" and "no." Each message includes a random identification number that is large enough to avoid any duplication with other identification numbers generated by other voters.
>
> The voter blinds each message and sends them to a centralized tabulating entity, along with the blinding factor for each message.
>
> The centralized tabulating entity verifies that the voter has not submitted blinded votes previously. Then it chooses nine of the ten sets, and examines them using the blinding factors to see that they are properly done. The centralized tabulating entity individually signs each message in the tenth set (the one that was not examined), and sends these signed messages to the voter. The centralized tabulating entity stores the name of the voter in a database.
>
> The voter un-blinds the set of signed messages. Now the voter has a set of votes, such that each vote is signed by the centralized tabulating entity.
>
> The voter chooses one of the votes, encrypts it with the centralized tabulating entity's public key, and sends in the vote.
>
> The centralized tabulating entity decrypts the vote with its private key and checks the signatures. It also checks for a duplicate identification number. Then it saves the identification number and tallies the votes. It publishes the results of the election, and every identification number with its associated vote.

A voter cannot generate valid votes without the involvement of the centralized tabulating entity because he does not know the private key that belongs to the centralized tabulating entity. A voter also can not intercept and change other people's votes because they are encrypted with the centralized tabulating entity's public key. If a voter tries to send in the same vote twice, the centralized tabulating entity will notice the duplicate identification number and discard the second vote. If a voter tries to get multiple votes signed by the centralized tabulating entity, it is probable that this will be discovered when the centralized tabulating entity examines nine of the ten sets of messages.

Voters can check that their votes were counted correctly by reviewing the published list of identification numbers and their associated votes. However, if a voter discovers that the centralized tabulating entity changed his vote, he has no way to prove it.

The centralized tabulating entity can not figure out how each voter cast his vote, assuming that the centralized tabulating entity did not examine the tenth set and observe the identification numbers on the votes before they are cast.

If the voter sends in his vote using a mechanism that is not entirely anonymous, the centralized tabulating entity can determine how the voter cast his vote. Also, the centralized tabulating entity could generate a large number of signed valid votes by itself, and thus alter the result of the election.

It is also important that randomly generating identification numbers is implemented well. Let's say that each identification number must be 64 bits in length. This gives a set of 2^64, or about 10^19, possible identification numbers. This is a large set of numbers, and the probability that two voters randomly choose the same number is small. If each identification number must only be 8 bits in length, there would only be 2^8 or 256 possible identification numbers. This is not a very large set of numbers, and the odds of a collision are much greater. Also, it is important that the numbers are generated in a truly random way. If it is possible for a group of voters to all generate the same numbers based on some quirk of their pseudo-random number generators, it is possible that some valid votes may be discarded by the centralized tabulating entity.

The voter has to generate quite a few messages to implement this protocol. The voter also has the task of blinding and un-blinding messages. Presumably, both tasks can be automated to reduce the burden on the voter. Additionally, the voter must wait for a response from the centralized tabulating entity. What happens if the centralized tabulating entity does not respond, or does not respond in time for the election?

An Anonymous Electronic Voting Protocol for Voting Over the Internet

Indrajit Ray, Indrakshi Ray, and Natarajan Narasimhamurthi published a paper [RaRaNa] that describes a cryptographic protocol for voting via the Internet. Although I am not convinced that Internet voting is viable for national elections, I believe the basic ideas presented in the paper are worth discussing.

Here is a brief summary of the protocol.

> Each voter has a public key and a private key. Before the election, the voter registers with a voter registration authority. This voter registration authority compiles a list of registered voters and issues a certificate for each registered voter that contains the voter's identity and public key.
>
> After a voter authenticates himself over the Internet with the certificate, a ballot distribution authority provides a signed blank ballot and a signed digest of the voter certificate. The blank ballot includes a serial number and a signed (by the ballot distribution authority) digest of the serial number.
>
> The voter generates a unique voter mark by performing a one-way permutation of the serial number. The voter receives a blind signature on the voter mark from the certifying authority. The blinded voter mark has the voter's signature on it, to authenticate the voter at the certifying authority.
>
> Once the voter receives the signed voter mark, the voter removes all identifying information and fills in the ballot. The completed ballot and the voter mark signed by the certifying authority are sent to the vote compiler via a protocol similar to anonymous ftp.
>
> When the voting period ends, the vote compiler publishes all the cast ballots in a public place and announces the results. The certifying authority publishes the voter marks in a public place. The ballot distribution authority publishes the number of blank ballots distributed and the serial numbers of those ballots.
>
> The voter can verify that his ballot has been counted. A voter can detect and prove fraud by anonymously submitting a copy of his ballot signed by the certifying authority.

It occurs to me that an eligible voter could use this system to obtain a blank ballot and sell it. If an illegitimate voter can obtain a blank ballot and the private key of an eligible voter, that illegitimate voter can cast a vote. The system seems to assume that voters will keep their private keys secret, but what incentive is there to do so?

It is even easier for an illegitimate voter to assume the identity of an eligible voter. In this case, the identity of the eligible voter is any information the voter must provide to the voter registration authority. I see nothing to prevent someone anywhere in the world from pretending to be me, especially if I cooperate. It is a little more difficult, with more risk involved, for someone to walk into a polling location and pretend to be me. For example, it may be difficult for Indrajit Ray, one of the authors of this paper, to walk into a polling location and pretend to be Lelia Barlow.

Further, if the eligible voter has his identity stolen, that eligible voter will be prevented from casting his vote since the system should ensure that no eligible voter can vote more than once.

The voter may retain all information during these transactions. Therefore, there is a link between the voter and the cast vote. This may also invite vote buying.

Let's assume now that voters are honest, that they have the means to ensure the secrecy of their private keys, and that the identities of eligible voters cannot be stolen. This protocol is still a lot of work for the voter. (a) The voter has to obtain or produce a public/private key pair, and understand which to use in the protocol. I'll argue that this is beyond the scope of the average voter. (b) The voter has to register with a registration authority, and possibly wait for the registration authority to issue a certificate. If this transaction occurs via the Internet, what happens if the registration authority does not issue a certificate in time for the election, or if the certificate is intercepted and not allowed to reach the voter in time for the election? What if the certificate is corrupted during transmission? (c) The voter has to generate his own unique mark by applying a transformation to the serial number of the blank ballot. (d) The voter has to send his unique mark away to receive a blind signature, and wait for the result. In this sense the certifying authority could also be a bottleneck. What if the certifying authority refuses to apply a blind signature to the voter mark in time for the election, preventing an eligible voter from casting a vote with this system. What if this is intercepted, or corrupted? (e) The voter has the responsibility of removing any identifying information from the ballot. This, again, may prove difficult for the average voter to accomplish without the help of software, and software may or may not be trusted to do a good job at this task. (f) Finally, the voter gets to cast a vote. Hopefully, everything has worked to this point, and the vote can get through to the vote compiler. If not, the voter may have no recourse during the voting interval. (I wouldn't want to be working tech support on election day!) (g) Although it is not explicitly a requirement, the voter should verify that his ballot was counted correctly and he should report any errors.

The conclusion of the paper states:

> "We are aware of three shortcomings: Since the voter can identify his ballot, we can not prevent vote buying. Secondly, if several voters, after obtaining CA's [the Certifying Authority] signature, decide not to cast their ballot, then the three agents can cast fraudulent ballots. This fraud cannot be detected if the number of fraudulent ballots is less than the number of signed ballots that were not cast. If such a fraud is detected, then proving it will require the cooperation of the voters who did not cast their signed ballots. Finally, we allow a cast ballot to be traced back to an IP address (not to a voter). By using public voting kiosk we can avoid the IP address to be linked to a voter."

I disagree with the statement, "Since the voter can identify his ballot, we can not prevent vote buying." As previously described, there appear to be other, more serious issues with this protocol (as it is described) which make it vulnerable to vote buying. If the voter could identify his ballot, without there being a direct link between himself and the identifier, vote buying may not be a significant concern.

That said, I like the fact that the voter can verify, in an anonymous manner, that his intended choices are included in the final tally. The voter also can anonymously contest the results if his intended choices are not included in the final tally. The reader can imagine that this protocol could be applied to electronic voting systems with voter-verifiable paper ballots (ala the Mercuri Method) at a polling location that are protected by some degree of physical security. Additionally, election officials could be summoned if there is any problem.

Secret-Ballot Receipts and Transparent Integrity: Better and Less-Costly Electronic Voting at Polling Places

David Chaum published a paper [CHAUM] that describes an idea for voting via electronic voting machines at a polling location.

I will attempt to summarize Chaum's protocol here:

> After a voter makes his selections, the electronic voting machine prints a receipt that shows all of these selections in a human-readable form. If the voter agrees with the printed receipt, he is asked by the machine whether he wishes to take the top or bottom "layer" of the receipt. After the voter has indicated a choice, the machine prints out the end of the receipt, and releases the receipt to the voter. One layer of the receipt is labeled "Voter keeps this privacy-protected receipt layer" and the other layer of the receipt is labeled, "Voter must surrender this layer to poll worker."
>
> The voter takes both layers of the receipt out of the booth. A poll worker destroys the layer that is labeled "Voter must surrender this layer to poll worker." The voter takes the other layer with him. The layer kept by the voter contains a digital signature in barcode format that can be used to determine the authenticity of the receipt. For example, this barcode could be scanned by a volunteer organization when the voter leaves the polling location. Meanwhile, the electronic voting machine saves a copy of the layer of the receipt labeled "Voter keeps this privacy-protected receipt layer."
>
> After the polls close, this electronic version of the receipt is posted on an official website. The voter can privately view the website, and see his receipt listed by serial number. The voter could check that it has been posted correctly by printing it and overlaying it with the layer of his receipt, for example.
>
> The set of receipts from all voters is posted on the website, as well as a set of human-readable images that contains exactly the same number of items. However, the receipts and human-readable images are not presented in the same order. A chain of intermediate steps, performed by election trustees, link the two sets and ensure correspondence between the sets. The trustees make use of cryptographic techniques to prepare the human-readable images and tally the vote.

The layered receipts are constructed using the principle of a one-time pad. The layers are not constructed at random, but should be generated as "indistinguishable" from random except by a set of trustees in the tabulation process. It is, therefore, essential that the voter select a layer after the contents of the receipt have been printed to prevent "cheating" by electronic voting machinery. Because the machine does not know which layer the voter will select, it has a fifty percent chance of creating an "invalid" layer that is not detected by the voter.

When the layers are overlaid, they produce either a part-transparent image or an opaque image. (See Figure 2, below.)

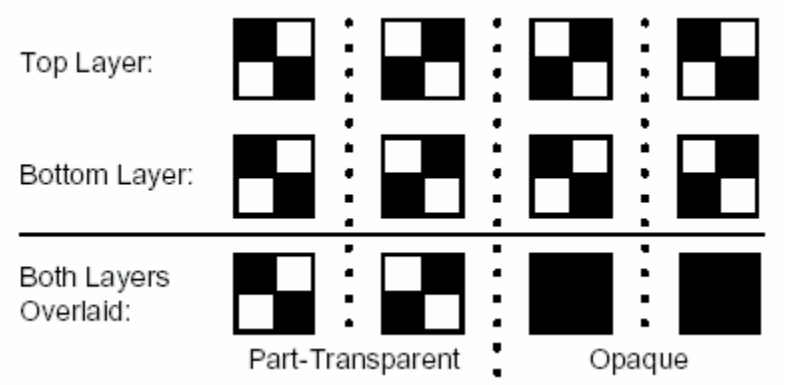

Figure 2. Two symbols, separate and overlaid.
Source: [CHAUM]

With only one layer, an observer is unable to determine the resultant image when both layers are overlaid. A human-readable image is constructed, as described in the figure below.

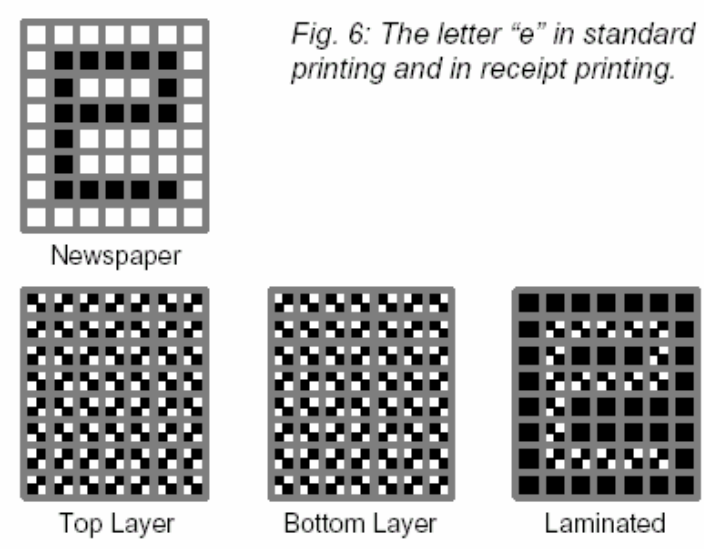

Figure 3. Construction of a human-readable image
Source: [CHAUM]

The tabulation process begins with the set of electronic receipts and produces a final tally based on the corresponding human-readable images. Chaum's paper states, "Then the first trustee can produce the first "intermediate batch" from the receipt batch. After that, a trustee forms the second intermediate batch from the first intermediate batch, and so forth, until the last trustee forms the tally batch from the last intermediate batch." This process is further described by the following figure:

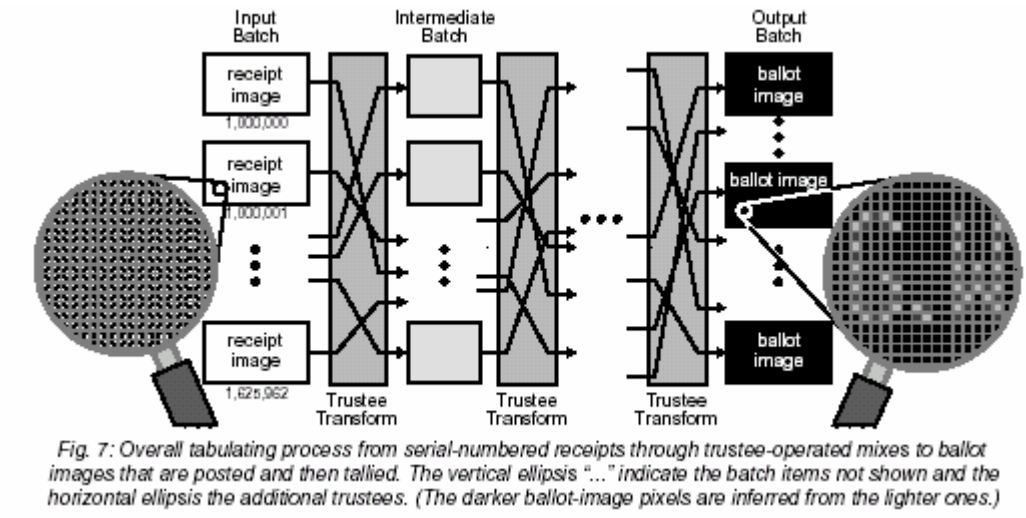

Fig. 7: Overall tabulating process from serial-numbered receipts through trustee-operated mixes to ballot images that are posted and then tallied. The vertical ellipsis "..." indicate the batch items not shown and the horizontal ellipsis the additional trustees. (The darker ballot-image pixels are inferred from the lighter ones.)

Figure 4. Tabulation process.
Source: [CHAUM]

Rebecca Mercuri, in [MERC2002], writes,

> "David Chaum, a Palo Alto, Calif., cryptologist who, 20 years ago, invented electronic cash, has a technique that provides the best of all possible worlds: a computer-generated, voter-verified physical ballot that also gives the voter a receipt that can be used to determine that his or her vote was tabulated correctly, without revealing its contents. One drawback of Chaum's method is that to demonstrate that the votes are tallied correctly requires a lot of math. As a result, it is difficult to explain to election officials, poll workers, and voters how it establishes the correctness of the balloting and tabulation process. But it gives a glimpse of the type of voter-verifiable systems that may be used for future elections."

After reading Chaum's paper, I am inclined to agree. In a technical sense, his ideas are interesting and (as far as I have been able to determine) technically sophisticated. In particular, I think the idea (first proposed by Naor and Shamir in 1995) to use the concept of a one-time pad to generate multi-layered receipts is clever. However, I remain unconvinced that his methodology would work well in a real-world election.

Voters can match their receipt to the electronic copy of the receipt on the official website. Does this tell them that their vote was counted correctly? Perhaps the electronic voting system will simply post the electronic copy of the receipt on the official website, and then construct a different tally.

Chaum's paper refers to a "simple, open-source program downloadable from any of multiple suppliers (which can also check the consistency of each entry in the receipt batch)." Therefore, if the voter wants to obtain this software, he can run it to check that there exists a one-to-one correspondence between the set of electronic receipts and the set of human-readable images. Honestly, I can't see the average voter going to this kind of trouble. Presumably, volunteer organizations would assume this role.

Even if the voter goes to this effort, will it make him confident that the system didn't cheat? I think it is safe to assume that the average voter does not possess the necessary mathematical background to understand the process of transforming a set of electronic receipts into a set of human-readable images. Nor does the average voter have the ability to examine source code. The average voter will still have to trust someone else to tell him that everything worked OK, and that his privacy was protected.

Further, I do not understand why it is important to protect the privacy of the voted ballot within the tabulation system. Would it not be possible for the electronic voting machine to unscrupulously record an image of the voter's choices without the voter's knowledge anyway? If we assume that the voter uses the

electronic voting machine without identifying himself to it (i.e., "Name: John Doe, Address: 1234 A Street, Corvallis, OR.") why does the voted ballot need to be kept secret from all but the combined group of trustees? Is this not the same as dropping a voted paper ballot into a sealed ballot box? Certainly, a voted paper ballot is not encrypted or obfuscated in any way. Any obfuscation would only be a hindrance to computing the tally.

In addition, I think human users of a system that incorporates this method are likely to be confused by the process. There may also be vulnerabilities that could be exploited by an unscrupulous individual. The user may not understand why he has to select either the top layer or the bottom layer of the printed receipt. Perhaps the user could rip the receipt out of the machine before it is finished printing. The user may leave the receipt in the machine or it may fall to the floor unnoticed and un-separated. The election official may destroy the wrong layer, or may fail to destroy either layer. Is there a way for a malicious individual to tamper with the serial number on the printed receipt, rendering it invalid? What is the policy that governs voters "trying again" if there is a problem with their printed ballot? Could someone forge printed receipts and use them to bog down a process of disputing the system?

Conclusions

So far, in my opinion, I have not read a description of a technique that could be implemented to create a "perfect" electronic voting system. Such a system would, among other things, provide acceptable levels of accuracy, privacy, and usability. A variety of protocols have been created, some of which contain interesting ideas. Certain ideas, if implemented well, could potentially improve the process of conducting elections, and improve voter confidence. Surely, there is room for improvement.

I am concerned about the length of time it is taking for the security (or, rather, the insecurity) of electronic voting systems to become an issue. Protocols designed to improve the security of computerized voting have been discussed for decades, in some cases.

In his book, Applied Cryptography [SCHN1996], Bruce Schneier describes a variety of protocols for what he calls computerized voting. His claim is that, "Computerized voting will never be used for general elections unless there is a protocol that both maintains individual privacy and prevents cheating." Unfortunately, this statement has not been proven true.

A source of irritation for me is the perception that "Security can be added later – we need to get the system working first." As my former colleague, George Cox, was fond of saying, "Security should be an adjective, not a noun." The goal should be creating of a secure system, not applying security techniques to an existing system like a patch on a flat tire. If anything is done about the insecurity of these current electronic voting systems at all, it is likely that a lot of money and effort must be expended to revise systems that are already in use.

Experience has shown that it typically requires a major system failure before security is considered an important enough problem to fix. Will the 2004 presidential election provide the incentive to pay attention to the security of our election systems?

References:

[BARL2003] Barlow, Lelia. "An Introduction to Electronic Voting." November 2003. http://islab.oregonstate.edu/koc/ece399/f03/explo/barlow.pdf

[CHAUM] Chaum, David. "Secret-Ballot Receipts and Transparent Integrity: Better and less-costly electronic voting at polling places." (Date of publication unknown) http://www.vreceipt.com/article.pdf

[MERC2002] Mercuri, Rebecca. "A Better Ballot Box." IEEE Spectrum, October 2002. http://www.notablesoftware.com/Papers/1002evot.pdf

[MERC2003] Mercuri, Rebecca. Website on "Electronic Voting" last updated September 1, 2003. http://www.notablesoftware.com/evote.html

[PFIT1996] Pfitzmann, Birgit. *Digital Signature Schemes: General Framework and Fail-Stop Signatures*; Springer-Verlag, Berlin 1996.

[PITT2003] Pitt, William Rivers. "Electronic Voting: What You Need to Know," Interview with Rebecca Mercuri, Barbara Simons, and David Dill, October 20, 2003. http://truthout.org/docs_03/102003A.shtml

[RaRaNa] Ray, Indrajit and Ray, Indrakshi and Narasimhamurthi, Natarajan. "An Anonymous Electronic Voting Protocol for Voting Over the Internet." (Date of publication unknown.) http://citeseer.nj.nec.com/471417.html

[SCHN1996] Schneier, Bruce. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*; John Wiley & Sons, New York 1996.