From *SIAM News*, Volume 37, Number 3, April 2004

Works in Progress: Trustworthy Cryptographic Voting Systems

By Sara Robinson

After a statistical tie in the 2000 presidential election exposed flaws in U.S. voting systems, Congress funneled money to states for mandated upgrades of voting technology. The most popular upgrade—electronic voting machines—comes with problems of its own, however.

The fundamental problem, researchers say, is that without an independent record of voters' actions, a complex computer program cannot be trusted to count votes correctly. "Paperless voting requires software and hardware to be perfect," says David Dill, a professor of computer science at Stanford University. "It must never lose or change votes."

The secret-ballot requirement makes it particularly easy to compromise a voting system. Each voter wants to know that his vote was counted exactly as cast, but given a record of his actions, a voter could prove to a third party how he voted, and be paid for his vote, or he could be coerced into voting a certain way.

As discussed in the March issue of *SIAM News*, computer scientists have developed new voting methods that bypass the need to trust sophisticated software. The simplest is the voter-verified paper audit trail: Voters would check paper printouts of their filled-in ballots for accuracy before depositing them in the ballot box. The paper records would be counted or sampled, removing the need to trust the machine.

The "frog" system, proposed by Ronald L. Rivest of MIT, Shuki Bruck of Caltech, and David Jefferson, now at Lawrence Livermore National Laboratory, gets around the problem by separating out and securing "vote casting," the most sensitive step of the voting process. In their system, a voter would fill in her ballot on any old voting machine, which would store the ballot on a memory card or print it on paper in bar-code form. The secure "vote-caster," which would have a dedicated operating system and simple, open-source hardware and software, would read and display the ballot, and, if the voter is satisfied that her choices are correctly represented, cast the vote.

Neither of these methods ensures that each ballot was counted as cast, however. The paper audit trail relies on the checks and procedures now used for paper ballots, while the frog method relies on experts for careful checking of the vote-caster's software and hardware.

Two new cryptographic voting systems, the culmination of two decades of cryptographic voting research, address this issue. In these systems—developed independently by Andrew Neff, a mathematician and chief scientist at VoteHere, a Bellevue, Washington-based start-up, and David Chaum, a Los Angeles-based independent cryptographer and privacy expert—votes would be counted publicly, with cryptographic techniques used to secure privacy. Voters could use their receipts to verify that their votes were counted as cast, but not to prove to a third party how they had voted. Both schemes are provably reliable, even in the presence of compromised machines and dishonest election officials.

Josh Benaloh, a cryptographer at Microsoft Research who is familiar with both methods, is confident that their security and privacy claims are mathematically verifiable. These cryptographic systems are a good solution to the voting problem, he says, adding: "Anything we can do to reduce the trust we need to place in computers and people is a good thing."

Cryptographic Voting Systems

Each of the systems described in this article is accomplished in two steps, which can be handled independently. The first step produces an encrypted, filled-in ballot on an untrusted voting machine. The second step is a method by which untrusted election officials take a publicly posted list of encrypted, filled-in ballots, each associated to a particular voter, and decrypt them, removing voter-identifying information, without being able to change the contents of the ballots.

Both systems rely on key participants to enforce security or privacy:

The Voters. Some voters must check that the encrypted receipts on which their votes are recorded were printed correctly and that the receipts appear in the virtual ballot box posted on the Internet. After the votes have been decrypted, voters can also verify the counts.

The Election Trustees. Responsible for ensuring ballot secrecy, they decrypt and count the votes in a public forum, where all their actions can be monitored. A flexible fraction of the trustees (Neff's scheme) or all the trustees (Chaum's scheme) would have to cheat in order to compromise the privacy of an election.

Interested Third Parties and Observers. Any observer or voter can check to make sure that the voting machines are trustworthy and that an encrypted receipt is valid. Chaum sees this role being played by the League of Women Voters, the American Civil Liberties Union, or other such organizations. In Neff's scheme, voters or workers at the polling place would perform this function.

Mathematically sophisticated voters and observers can also verify the proof that no ballot was changed in the decryption process.

Step One: Generating Encrypted Ballots on an Untrusted Machine

The first task of these voting systems is to provide each voter with an encrypted receipt that correctly represents his ballot choices.

The potential adversary is an untrustworthy voting machine that might fool the voter into thinking all is well while secretly changing his vote. Both Neff and Chaum handle this problem by way of procedures that randomly check the output of the voting machine. Because the machine does not know when its output will be checked, there is a high probability of detecting any attempts to cheat.

Chaum's method. Chaum's method can be realized with a simplified version of a scheme for "visual cryptography" developed in 1994 by Moni Naor and Adi Shamir, computer scientists at the Weizmann Institute of Science. The scheme uses pixel patterns on layers of translucent paper to spell out a printed message. When the layers are aligned, one on top of another, the opaque and semi-opaque areas combine to make the message legible, although it's impossible to determine the message from either layer alone. For readers familiar with the term, the scheme produces a visual "one-time pad," where one layer, generated pseudorandomly, is the key.

Chaum envisions a voter signing in at the polling place and selecting her choices on a voting machine. Her completed ballot is then displayed on a two-layer printed page, each layer of which is a random binary image. So that both layers contain equivalent information, each page is a mixture of the two pages given by the visual cryptography scheme, obtained by swapping every other pixel. Hence, at the end, each layer contains half of the randomly generated "key." The other pixels contain the vote, which can be decrypted if the other half of the randomly generated key can be regenerated.

Some voters must check that the encrypted receipts on which their votes are recorded were printed correctly and that the receipts appear in the virtual ballot box posted on the Internet. After the votes have been decrypted, voters can also verify the counts.

The bottom of the page (printed on both layers) contains three numbers: a serial number, and the seeds, in encrypted form, used to generate the random pixels in each half of the voter's ballot. The voter needs to visually check that her ballot accurately represents her choices and that the three numbers are the same on both layers. She then chooses one of the layers; the machine digitally signs the serial number and the information in the entire layer, printing the signatures at the bottom and releasing both layers to the voter. The voter can now check again that the two halves together form a readable ballot that correctly represents her choices. The voter must present the half she didn't choose to a voting official, who checks that it is the correct half and destroys it in a shredder. The machine memory's copy of this other layer must also be destroyed in order to protect the voter's privacy.

Outside the polling place, interested third parties can use the digital signatures and encrypted seeds to ensure that the receipts are valid voting records and that the machine did not falsify the receipt layer. Having the voter choose a ballot half ensures that the machine prints both layers correctly. Each time it does not, there is a 50% chance that the bad layer will be chosen and detected. With the odds of getting caught increasing with each attempt to cheat, it becomes infeasible for the machine to cheat on a significant scale.

Once the polls close, digital images of the receipts are posted on the Internet, and each voter can use the serial number on her receipt to make sure it was included in the virtual public "ballot box."

Neff's method. The method for producing encrypted ballots in Neff's scheme differs from Chaum's both in technical details and in the voter's experience.

Neff's system provides each voter with a voting token that contains the appropriate blank ballot for that voter's party and district. The voter inserts his token into a touch-screen voting machine and views his ballot choices. Next to each option on the ballot is a number that Neff calls a "verification code." After the voter has made his selections, the voting machine displays a ballot summary, showing all the choices with their corresponding codes, along with the races the voter skipped.

At the same time, the machine produces a ballot receipt (in either paper or electronic form) that lists the codes corresponding to the voter's choices, without indicating what the codes represent. Before casting his ballot, the voter must check that the numbers on his receipt correspond to the numbers of his ballot choices as viewed on the screen.

The assignment of verification codes to ballot choices—which Neff calls a "codebook"—is a random permutation that is generated separately for each voter. The codebooks—committed to before the election—are generated by a parameter representing a secret shared by k of the n election trustees.

To ensure that the machines do not display false codebooks, election officials and voters continually test them on election day. Later, audits of the values displayed during the tests are done to check that the test codebooks were correctly generated. Because the machine cannot distinguish between a test display and a display that will end in a cast ballot, any attempt by the machine to cheat will be caught with high probability, given a sufficient number of such tests.

As with Chaum's scheme, the encrypted receipts have serial numbers and are digitally signed by the voting machine so that third parties can verify their legitimacy. After the election, the receipts are posted on the Internet, where voters can check for them.

Step Two: Mixing and Counting

Step one in both methods results in a list of encrypted receipts that are posted on the Internet and that, with high probability, are correct representations of voters' ballot choices. The next task, performed by the trustees, is to decrypt the ballots, strip off the serial numbers, and shuffle the decrypted ballots. The goal is to provide a list of decrypted ballots that cannot be traced to individual voters. This list can then be counted publicly on the Internet, without fear of betraying voters' privacy.

The tricky part is to set up the decryption and shuffling process so that privacy is preserved yet outside observers can ensure that dishonest trustees are not changing votes. The shuffling step, called a "mix," uses methods developed in two decades of cryptography research and is worth highlighting as a separate achievement. A mix typically involves an interactive proof, a way to prove to a third party that something has certain properties without revealing what it is.

In Search of a Perfect Mix

Chaum introduced the concept of mixing in 1981, primarily as a method for issuing untraceable mail, but also as a potential component of an electronic voting system. With his method, each ballot (or other item) is enclosed in layers of encryption, as in a series of sealed envelopes. The trustees perform the decryption in stages: Each trustee, using a secret known only to him, opens and removes the outermost envelope. He shuffles the inner envelopes before passing them on to the next trustee. Because there is no universal way to verify that votes were not changed, a dispute-resolution mechanism is necessary for voters who claim that their votes were not proper-ly included in the final tally; voters shed their privacy when they lodge protests.

The goal is to provide a list of decrypted ballots that cannot be traced to individual voters. This list can then be counted publicly on the Internet, without fear of betraying voters' privacy.

During the 1980s and early 1990s, cryptographers improved on Chaum's scheme. The biggest advance, Benaloh says, came in 1993, when Choonsik Park, Kazutomo Itoh, and Kaoru Kurosawa produced a method for shuffling and renaming encrypted ballots. The following year, Kazue Sako and Joe Kilian simplified and improved the method, adding a way to prove that no ballots were changed during the mixing process. The new technique decoupled the encryption from the shuffling, simplifying the process and making it possible for other features to be added to the encryption. Description of the shuffling method and proof of correctness is beyond the scope of this article.

The researchers' scheme for renaming encrypted ballots borrows from an encryption method called ElGamal encryption. It begins with a large prime p so that $p - 1 = qr$, where q is a large prime that is co-prime to r. Let H be the order q subgroup of Z_p^*. For any g and h in H, the ElGamal encryption of an element m is given by a pair $(x,y) = (g^a, mh^a)$, where a is chosen uniformly from $[1,q]$. In this form, an element (x,y) can be renamed by taking a random c in $[1,q]$ and computing (xg^c, yh^c), which is the same as (g^{a+c}, mh^{a+c}) and thus also an encryption of m. The drawback of the Park–Itoh–Kurosawa method is its high computational cost, Benaloh says.

In 2002, M. Jakobsson, A. Juels, and Rivest devised a fast and remarkably simple method for verifiably shuffling and renaming encrypted ballots. The researchers show how a list of encrypted elements can be permuted in such a way that no element can be traced to its predecessor and such that, although the elements are renamed, the permutation can be verified by a compu-tationally simple interactive proof. The idea is to do two shuffles (two permutations composed).

To prove that the permutation was performed correctly, the shuffler provides an interactive proof to an observer as follows: The observer first chooses half of the ballots from the intermediate stage, before the second shuffle. For those ballots, the shuffler reveals their predecessors under the first permutation. For the other half of the ballots, the shuffler reveals their forward images under the second permutation. If the ElGamal renaming procedure is used, each predecessor or successor link revealed is accompanied by a proof that the two linked elements are equivalent, i.e., a c is provided such that (g^c, h^c) is equal to the component-wise quotient. Because the method will detect an output ballot that does not correspond to an input ballot with probability at least 1/2, any attempt to alter k ballots will succeed with probability at most 2^{-k}. Thus, any significant attempt to change ballots will almost certainly be detected.

Chaum's mixing method uses a combination of his original layered encryption idea and the verification method of Jakobsson et al. Neff's method, like the method described by Park et al., decouples the encryption from the shuffling. He relates the two encrypted lists of k ballots (pre- and post-shuffle) to the roots of two degree k polynomials over a finite field determined by the cryptographic key size. Using an interactive proof, he then shows that the two lists have a high probability of being permutations of one another by demonstrating that the polynomials take the same value at a randomly chosen point.

A Future for Cryptographic Voting Systems?

Both systems will soon be realized and tested. VoteHere plans a software release implementing Neff's method sometime this spring. Poorvi Vora, a professor of computer science at George Washington University, and colleagues Jonathan Stanton and Rahul Simha are directing a student implementation of Chaum's system and hope to use it for a student election.

The most serious issue for these cryptographic systems, however, is whether they can win the trust of voters who have very little background in mathematics or cryptography. For most voters, several systems experts point out, cryptographers and cryptographic approaches must seem just as untrustworthy as proprietary computer programs.

At the very least, building trust in such methods will require a concerted effort by the researchers that support them. "We can work hard at building better voting machines, but this is really only a means toward the ultimate goal of building trust in the election results," Rivest says.

Sara Robinson is a freelance writer based in Pasadena, California.

Readers who missed Sara Robinson's introductory article on electronic voting, "What's So Special About Voting?" in the March issue of SIAM News *can find it on the Web: http://www.siam.org/siamnews/03-04/e-voting.htm.*