The case for networked remote voting precincts

Daniel R. Sandler Dan S. Wallach
Rice University
{dsandler,dwallach}@cs.rice.edu

Abstract

Voting in national elections from the comfort of one's home computer may never be practical or secure, but we argue that remote network voting can be both practical and secure. Provisional and postal absentee ballots, which trade some amount of anonymity for the ability to determine the eligibility of a distant voter, serve as a template for how electronic remote voting might proceed. We propose the "remote voting center": a government-operated facility located in embassies, consulates, and other remote areas where voters might normally need to vote by mail. Each remote voting center would maintain one or more electronic voting systems and a registration system. A voter presents identification to the registrar on site and is then directed to cast a ballot in a private electronic voting booth. The cast ballot is encrypted and forwarded to the registration system, where it is wrapped with the voter's identifying information. This double enclosure is signed by the voting center and posted publicly where it can be examined and canvassed by officials in the voter's home precinct. If and when the ballot is accepted, it can be combined with existing tallies using standard cryptographic techniques to preserve the voter's anonymity. The resulting system has privacy properties comparable to provisional voting in a local polling place, and represents an improvement over postal voting by offering the voter privacy in a supervised voting center.

1 Introduction

Never. This, the answer given by e-voting security researchers when asked when we will be able to vote in national elections over the Internet, is unsatisfying to many because of the tremendous convenience it would seem to afford (see, e.g., Alvarez and Hall [2]). The number of endeavors, from personal entertainment to securities trading, that have been profitably brought online would imply that the Internet can improve any task that does not absolutely require one to be physically present.

Voting, unfortunately, requires absolute trust in two factors that cannot be adequately controlled in the residential Internet scenario: *environment* and *equipment.* The voter's PC may be compromised; the voter may be coerced. It is not the only such task; academic testing, for example, requires a testing environment free of distraction, collusion, and unauthorized assistance.

Participating in national elections from the comfort of one's home computer may never be practical or secure, but we argue that *remote* voting can be both. Voters in many jurisdictions are currently permitted to cast provisional ballots in situations where their eligibility to vote is in doubt; the voter's identification is submitted along with the sealed ballot for consideration by elections officials. Postal voting (or "vote-by-mail") functions similarly. Both of these schemes trade some amount of anonymity for the ability to determine the eligibility of a prospective voter.

These techniques inspire our vision of practical electronic remote voting. We note that a DRE that encrypts individual ballots provides the sealed ballot described above; when digitally signed along with plaintext attesting the identity of the voter, it becomes an electronic replica of the conventional provisional ballot, albeit one that can travel faster and more safely than a postal envelope.

In this paper we propose that voters far from their home precincts visit a "remote voting center," a facility maintained and supervised by government officials (perhaps in foreign embassies or in controlled areas on military bases and ships). The remote voting center consists of one or more electronic voting booths and a registration system. Voters present their personal identification, and are then directed to cast a ballot in a private electronic voting booth with the proper local ballot (provided in advance by the election director of the voter's home precinct). The cast ballot is encrypted and returned to the registration system, which then in turn wraps the ballot ciphertext in the voter's identifying information. This might include a scanned signature or ID card or even a digital photograph of the voter taken at the time of voting. This double enclosure is then digitally signed by the voting center and posted on a public "bulletin board" where it may be examined and canvassed by the voter's home election officials. Once the election officials have determined that the ballot was cast properly (e.g., the voter's identification matched up with records on file and the proper ballot definition was used), then they can approve the still-encrypted ballot for inclusion in the final tally.

We continue by reviewing the procedures currently in place for postal and provisional balloting (Section 2), giving special attention to the security guarantees made to the voter for these (Section 3). Subsequently, we sketch remote voting using the VoteBox e-voting platform (Section 4) and show how it relates to other proposed Internet voting schemes (Section 5). We conclude with a review of our proposal (Section 6).

2 Provisional and postal voting

Postal voting is used widely in the U.S. and is growing in popularity. The state of Oregon, for example, votes exclusively by mail. Many states offer "no fault" postal voting; voters may declare their desire to vote by mail without requiring any reason. In California, voters may declare their desire to vote exclusively by mail, and need never again cast ballots in person.

A month in advance of the proper election date, ballots are mailed to these voters, giving the voter time to cast the ballot or request an alternative ballot if the original is lost or spoiled. Completed ballots are placed in an opaque return envelope. The back of this envelope has designated areas for the voter to inscribe her personal identifying information, including her signature. A paper flap (or, in some cases, another enclosing envelope) conceals this personal information while the ballot is in transit.

When envelopes arrive in the mail at the elections office, they are counted and stored. Each envelope's signature and personal information is verified by hand against available registration data to determine whether the ballot inside should be counted. If an envelope is rejected, election officials may then attempt to contact the voter to offer then an additional opportunity to cast a vote, assuming the election is still ongoing. Ballots that are determined to be legitimate can are then removed from their envelopes and stored as any other ballot might be stored. These ballots can then be tabulated using the same optical scan machinery that can be used for paper ballots cast in traditional precincts.

Provisional voting, required as part of the 2002 Help America Vote Act, is semantically very similar to postal voting. Provisional voting occurs when a voter arrives at what she believes to be the proper precinct only to discover that she is absent from the registry of voters for that precinct. At that point, the voter may conclude that there was an error and declare the desire to cast a vote, regardless.

At this point, procedures vary from voting system to voting system. One solution, used in a number of locations, is that the voter is handed a paper ballot along with an envelope. The paper ballot is filled out, as normal. The envelope, much like a postal voting envelope, contains information about the voter's identity along with why he or she claims the right to cast a vote in this particular precinct. Some DRE voting systems offer similar functionality, tagging provisional votes with an identifier of some kind that corresponds to paper records describing the voter's situation.

Provisional votes are generally not tallied until a recount occurs or if the number of provisional votes is large enough to impact the outcome of the election. At this stage, election officials hold a public hearing to individually discuss each provisional voter and determine whether his or her vote will be counted. Once the envelope has been validated, the inner ballot can be removed and tabulated.

3 Security and privacy of remote voting

3.1 Conventional approaches

Voter anonymity is necessarily harder to safeguard when the voter's name, address, and signature accompany each ballot. Present-day provisional and postal voting attempt to preserve privacy through a combination of technology (ballots are enclosed in opaque envelopes) and procedure (envelopes are only opened if eligible, and once validated, a ballot is separated from its envelope).

Postal voting, however, suffers from several obvious problems. The postal mail channel is slow and not sufficiently reliable, particularly when delivering mail overseas. Furthermore, there are a wide variety of opportunities for election fraud with postal voting, ranging from outright bribery and coercion (i.e., selling unvoted ballots) to attacks upon and within the postal system (e.g., postal workers destroying or tampering with ballots). While some voters may detect that their ballots failed to arrive at their destination, it would be difficult to automatically detect and correct such errors. In cases where the postal delivery channel is too slow or too lossy, multiple round-trips with the voter are likely infeasible in the time allotted for voting.

Provisional voting, when performed inside a properly supervised voting location, is more robust against bribery and coercion, since the vote will have been cast in the privacy of a voting booth. Likewise, there are fewer concerns about loss or damage to votes while in transit. Nonetheless, as with postal voting, the ties between the voter's identity and ballot allow subsequent opportunities for fraud, whether the provisional vote is cast on paper or with current-generation DRE systems. We necessarily trust that the election officials will properly manage the process to preserve voters' privacy.

Both postal and provisional voting share the property that a variety of attacks can be *detected* even when they cannot necessarily be *corrected*. Voters can detect whether their ballots were received and whether they were tabulated. They cannot learn whether their ballots were tabulated accurately.

3.2 Goals for a networked replacement

If we are to create a remote voting system that replaces postal mail with Internet transmission, we must retain comparable security and privacy semantics to postal or provisional voting. While the ballot should be accompanied by information identifying the voter so that only eligible remote votes are counted, we must retain the opaque envelope: the voter's choices must be concealed until eligibility is determined, and then separated from the voter's identity before tabulation.

We also wish to preserve the ability to detect problems, even if we cannot necessarily correct them right away. For example, as Internet hosts are indisputably vulnerable to denial of service attacks, we must preserve the ability to cast a ballot regardless of whether or not the election authority can be reached from the polling place. That is, despite the supposition of a network connection, we cannot use "online" methods that require constant uptime of that connection. Finally, we should improve on postal voting by providing a voting environment that resists voter coercion and fraud.

4 RemoteBox: connecting remote precincts over the net

Our design for electronic remote voting builds on VoteBox [13], our current electronic voting platform designed for use in a single polling place. We begin this section by describing how VoteBox works before extending it to encompass our proposed remote voting model.

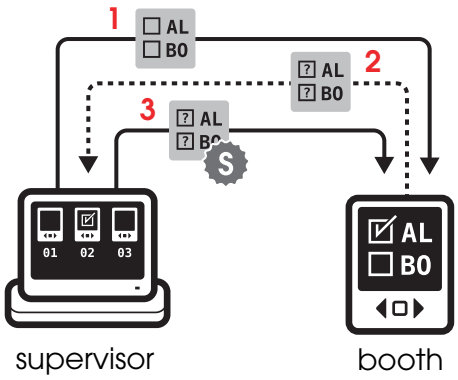

Figure 1: Casting ballots with VoteBox. For each voter, the supervisor console broadcasts a blank ballot (1), authorizing a particular VoteBox to accept the voter's input. The voter makes selections and casts the ballot, which is encrypted and broadcast by the booth (2). The supervisor acknowledges receipt of the encrypted ballot (3), allowing the booth to inform the voter that it is safe to depart. (All messages are received and logged by every device on the network, including other voting booths.)

4.1 VoteBox

VoteBox is a Java-based electronic voting system, developed to serve as a platform for broad e-voting research. Borrowing a technique due to Yee et al. [16], its GUI is mostly pre-rendered, resulting in a smaller runtime software stack and facilitating the code auditing process. It uses a local broadcast network called Auditorium [14] to provide fault-tolerance and tamper-resistant audit logs that show proof of the order of election-day events. A precinct whose VoteBoxes are linked via Auditorium is governed by a single *supervisor console*, under control of the on-site election administrator; the console simplifies the process of introducing each voter to a machine and improves the ability of poll workers to monitor and maintain voting systems in use.

Voting occurs in this system via a sequence of broadcast, hash-chained Auditorium messages. The protocol is described in more detail in prior work [14]; Figure 1 shows the essential steps involved in casting a single ballot. The final precinct tally is made by combining the logs from all machines in the polling place, removing duplicate entries, and decrypting the votes.

4.2 Remote electronic voting

We extend this model to provide a remote voting environment with the security properties of current provisional voting. We create the "RemoteBox" remote polling place from a VoteBox precinct by adding:

- A **remote polling place**, maintained and monitored by trusted / non-partisan government officials. We envision such a facility in embassies, consulates, and military bases: anywhere a large population of remote voters may be served.
- A database of **eligible remote voters**, mapping name and home precinct information to the correct blank ballot design for that voter. (Jurisdictions wishing to allow their voters to cast remote electronic ballots must furnish this information in advance.)
- A requirement that the voter **present identification** on election day: a government-issued ID or voter registration card, plus an interactive authenticator like a handwritten signature. Just as with a postal or provisional ballot, the voter must be identified so that he or she may be given the correct ballot, and so that election officials can decide whether the voter's cast ballot should be counted.

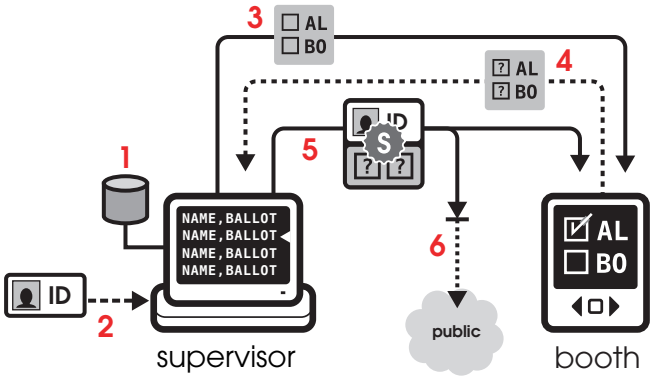

Figure 2: Voting with RemoteBox. A database (1) is furnished in advance, containing a blank ballot design every voter allowed to vote at this location. On election day, the voter presents identification (2) which is used to select the correct ballot for voting. As in a conventional polling place, the blank ballot is sent to a VoteBox (3) for voter input, and returned in the form of an encrypted cast ballot (4). The supervisor combines the result with the voter's identification and signs it (5), broadcasting it to the polling place for storage as well as on a one-way channel to a public medium (6).

- The notion of a **provisional electronic ballot**, which is a signed enclosure certifying the identity of the voter and her encrypted vote. This is an analogue of the conventional provisional ballot envelope (identifying the voter outside and sealing her choices inside) described in Section 2.
- A one-way channel to a **public medium**, for posting provisional electronic ballots. This could be an online channel such as an Internet link (perhaps on the other side of a data diode [11]) or an offline one such as a CD-ROM burner.

We show how these components fit together on election day in Figure 2.

4.3 Ballot definitions

An obvious complexity in this system lies in managing the *ballot definitions*, which will vary widely from county to county and state to state. If there were a single, standardized, national voting system, particularly based on pre-rendered ballots, then we could imagine these ballot definitions being collected by a centralized organization within the federal government. State- and locality-specific issues (e.g., Texas requires a "straight ticket" voting option while California forbids it) would be encoded in the ballot definitions, requiring the remote voting machines to be sufficiently generic to accommodate any voter from any jurisdiction. We could also envision current DRE systems' ballot preparation tools being augmented to output a standardized description of the ballot which could then be processed independently.

4.4 Cryptographic and pragmatic details

While this paper is not focused on the details of how cryptographic voting mechanisms work, we will provide a brief explanation of how "standard" cryptographic techniques can be applied to protect voter anonymity and provide for end-to-end verifiability of votes using this system.

Key management. Regardless of the cryptosystem that we might use, VoteBox requires that every piece of voting equipment (booths and supervisors) to have its own local key material for digital signatures. Moreover, each jurisdiction's election administration office (county clerk, etc.) must have individual public keys such that ballots cast remotely may be encrypted for their eyes only. We note that this problem is similar in scope to the issues surrounding ballot definitions (described in the prior section). Again, assuming the existence of a centralized organization within the federal government, this key material could be collected

and redistributed in advance of elections. Ballot definitions could likewise be centrally collected and disseminated. Each ballot definition would include the appropriate public key to use when encrypting votes cast with that ballot.

Cryptosystems. We may choose a cryptosystem to provide additional useful properties, such as end-to-end verification. For example, *homomorphic* cryptosystems [5] have the property that the (encrypted) sum of encrypted values may be found without decrypting them; that is, $E(a) \circ E(b) = E(a + b)$ for some appropriate homomorphic operation $\circ$ and encryption function E. We can combine this technique with threshold encryption [8, 3] to protect individual votes from prying eyes; the key material necessary for decryption may be distributed among a number of different election trustees who will only consent to decrypt the encrypted totals, rather than the individual votes. Alternatively, mixnets [6, 12] can allow the same set of trustees to each shuffle and reencrypt the ballots while proving that no ballots were lost or corrupted, ultimately achieving the same properties as homomorphic schemes. Furthermore, any schemes that allow voting machines to be "challenged" as part of their operation (e.g., a scheme due to Benaloh [3, 4]) or any schemes that allow voters to take home some sort of evidence that can later be compared against an electronic bulletin board are as compatible with this scheme as they are with any other DRE system.

Bulletin boards. A standard feature of many cryptographic voting protocols is the concept of a bulletin board where ballots are posted for all the world to see. We propose this as a mechanism for disseminating the results from remote voting precincts back to their proper home for tabulation. With proper key management and ballot definition distribution, performed in advance of the election, local election officials should easily be able to identify ballots on the bulletin board which are intended for their local consumption. These ballots would be encrypted with local election officials' public keys and signed with the keys of the remote voting system. The entire bulletin board from each remote precinct could then be signed by the remote precinct itself, protecting the bulletin board against tampering.

Networks. Ultimately, the bulletin board ballots must be transmitted from remote polling places to election officials. As a real-time feed is unnecessary (and possibly infeasible for some remote locations), ballots may be batched and sent infrequently, perhaps at the end of each day.

This gives us some flexibility in how exactly to transmit ballot data. For example, in order to isolate the remote precinct from the Internet, the supervisor console might burn a CD-ROM. This could then be transmitted via an overnight courier or hand-carried to a computer connected to any sort of network, whether public or private. All the remote results could be aggregated (but not tabulated) by the same centralized federal agency that coordinated the distribution of cryptographic keys and ballot definitions.

If this agency should sustain days-long attacks on its Internet connection, then this fact would certainly be visible to the public. All of the election results could then be disseminated through slower means (copying/mailing CDs, etc.). All that matters is that the various cryptographic signatures are properly verified. These may be verified both by local election officials and by the remote voting center's officials.

Various attacks. A voter with access to multiple remote voting centers (or, perhaps, a coalition of attackers using the stolen identity of one valid voter) could use the system described thus far to cast one vote per voting center. This would not necessarily be detected during the voting day. Nonetheless, each encrypted vote would be contained in a public envelope with the voter's identifying information present. Election authorities could certainly detect multiple votes having been cast, exactly as they can in the case of postal or provisional voting. It then becomes a policy problem to determine which vote should be counted and whether a crime has been committed. Alternatively, voters could be required to declare, in advance, which remote voting center they intend to use. When a voter shows up at the proper remote voting center, his or her

name is present and the vote proceeds normally. At other remote voting centers, the voter would be absent from the database and would either need to vote provisionally or would be turned away.

5 Related work

The U.S. military planned to deploy in 2004 an Internet-based electronic voting system called the Secure Electronic Registration and Voting Experiment (SERVE). They convened a panel of experts to evaluate it. A subset of them wrote a report describing all of the problems with voting over the Internet, such as easily compromised client platforms [10]. The military canceled the program, replacing it with a fairly simple fax-based scheme that is arguably less secure than SERVE [9].

In the U.S., several "primary" elections have been conducted over the Internet, including the recent "Democrats Abroad" primary election. Standard web browsers on standard client computers were used, and no particular measures were taken (or really could have been taken) to prevent voter bribery and coercion, much less deal with viruses or worms that might try to compromise the browser's behavior. In fact, the Democrats Abroad's primary did not have a secret ballot. In a radio interview[1], the administrator of the election said that the votes were actually public. The official disseminated results[2] only present country-by-country subtotals, so it's unclear exactly how much privacy is granted to Democrats Abroad's voters.

Internet voting has been used, perhaps more successfully, in national elections in Estonia [15]. The user authentication builds on a national ID card which contains a smart-card chip. Prospective voters insert the card into their computer, with a suitable adapter, and it allows them to authenticate to a government web site over an SSL-encrypted channel where they may cast a vote. Voters may vote as many times as they like, with the final one actually being tallied. The ability to cast multiple votes provides some limited resistance against bribery and coercion attacks. The use of SSL provides resistance against network man-in-the-middle attacks. Nothing in the Estonian voting architecture provides any protection against compromised client platforms.

Among commercial DRE voting systems, several vendors allow the use of modems to transmit election results (insecurely). While some states ban the use of these modems, others allow them under the guise of "unofficial" early election results. While this ignores the risk that an attacker may be able to compromise the tabulation system by calling it up on the telephone, these states are assuming that the records stored in the DRE systems themselves will survive the interregnum between the end of the election and their return to the voting warehouse, after which electronic results can be extracted directly from the voting machines. A similar property works for the bulletin boards in our scheme. Our bulletin boards can be disseminated in any way that data can be transmitted.

Two recent research systems confront the general Internet voting problem. Civitas [7] is an ambitious cryptographic voting system designed to allow Internet-based voting on a large scale. It suffers some limitations that preclude its straightforward deployment in nationwide elections, notably the requirement that each voter be issued a long-term cryptographic key pair for the purpose of acquiring per-election voter credentials. Moreover, an explicit design goal of the Civitas work is allowing *unsupervised* Internet voting. The authors admit that this requires trust in the end user's computer, and they respond to this by suggesting that voters seek out a voting terminal that they trust (e.g., one maintained by a political party or social organization). We note that this proviso causes a *practical* Civitas deployment to look quite a bit like the remote polling places suggested in this paper. Helios [1] is a Web-based system that sacrifices coercion resistance for a verifiable and minimally complex crypto-voting system that can be used from a voter's home computer. It employs a simplified variant of Benaloh's ballot challenge [3] with a single trusted server maintaining a

[1] http://weekendamerica.publicradio.org/display/web/2008/01/25/demsabroad/

[2] http://www.democratsabroad.org/sites/default/files/DA%20Global%20Primary%20Results%20FINAL%20REVISED.pdf

bulletin board for cast ballots. Helios is intended for "low-coercion elections," but if used exclusively in remote polling places it could be suitable for high-stakes national contests as well.

6 Conclusion

We have argued that the remote polling place is a model for networked remote voting that brings the benefits of DRE voting (convenience, speed, fault-tolerance) to provisional and postal voting. The security and privacy guarantees of these conventional remote voting methods are met or exceeded by this approach. We showed how an existing system design, VoteBox, can be straightforwardly extended to accommodate this voting model by enclosing anonymized, encrypted ballots in a public wrapper identifying the voter. A similar transformation (comprising a remote polling place and double-enclosure provisional ballots) should be applicable to any DRE-style voting system, provided that it is engineered (or re-engineered) with the necessary properties from VoteBox: it must accommodate a potentially large number of ballot designs (perhaps by loading them on-the-fly per voter), and it must provide the essential "opaque envelope" by encrypting each individual cast ballot.

References

[1] B. Adida. Helios: Web-based open-audit voting. In *Proceedings of the 17th USENIX Security Symposium (Security '08)*, San Jose, CA, July 2008. *To appear.*

[2] R. M. Alvarez and T. E. Hall. *Electronic Elections: The Perils and Promises of Digital Democracy*. Princeton University Press, Princeton, New Jersey, 2008.

[3] J. Benaloh. Simple verifiable elections. In *Proceedings of the USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '06)*, Vancouver, B.C., Canada, June 2006.

[4] J. Benaloh. Ballot casting assurance via voter-initiated poll station auditing. In *Proceedings of the 2nd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)*, Boston, MA, Aug. 2007.

[5] J. D. C. Benaloh. *Verifiable Secret-Ballot Elections*. PhD thesis, Yale University Department of Computer Science, 1987.

[6] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2), Feb. 1981.

[7] M. R. Clarkson, S. Chong, and A. C. Myers. Civitas: A secure voting system. In *IEEE Symposium on Security and Privacy*, Oakland, CA, May 2008.

[8] Y. Desmedt and Y. Frankel. Threshold cryptosystems. In *CRYPTO '89: Proceedings of the 9th Annual International Cryptology Conference on Advances in Cryptology*, pages 307–315, Santa Barbara, CA, July 1989.

[9] D. Jefferson, A. D. Rubin, and B. Simons. A comment on the May 2007 DoD report on voting technologies for UOCAVA citizens, June 2007. http://www.servesecurityreport.org/SERVE_Jr_v5.3.pdf.

[10] D. Jefferson, A. D. Rubin, B. Simons, and D. A. Wagner. A security analysis of the secure electronic registration and voting experiment (SERVE), Jan. 2004. http://www.servesecurityreport.org/.

[11] D. W. Jones and T. C. Bowersox. Secure data export and auditing using data diodes. In *Proceedings of the USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '06)*, Vancouver, B.C., Canada, Aug. 2006.

[12] K. Sako and J. Kilian. Secure voting using partially compatible homomorphisms. In *CRYPTO '94: Proceedings of the 14th Annual International Cryptology Conference on Advances in Cryptology*, pages 411–424, Santa Barbara, CA, 1994.

[13] D. R. Sandler, K. Derr, and D. S. Wallach. VoteBox: A tamper-evident, verifiable electronic voting system. In *Proceedings of the 17th USENIX Security Symposium (Security '08)*, San Jose, CA, July 2008. *To appear.*

[14] D. R. Sandler and D. S. Wallach. Casting votes in the Auditorium. In *Proceedings of the 2nd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)*, Boston, MA, Aug. 2007.

[15] A. H. Trechsel, G. Schwerdt, F. Breuer, R. M. Alvarez, and T. E. Hall. Internet voting in the March 2007 parliamentary electiosn in Estonia, July 2007. http://www.vote.caltech.edu/reports/EvotingEstonia2007.pdf.

[16] K.-P. Yee, D. Wagner, M. Hearst, and S. M. Bellovin. Prerendered user interfaces for higher-assurance electronic voting. In *Proceedings of the USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '06)*, Vancouver, B.C., Canada, 2006.