Reducing the Trusted Base

Douglas W. Jones

Department of Computer Science
University of Iowa
Iowa City, IA 52242
(jones@cs.uiowa.edu)

Presentation for
A Framework for Understanding Electronic Voting
Computer Science and Telecommunications Board
December 9, 2004
Washington DC

Last spring, Dan Wallach surprised many of us by saying that the purpose of an election is not to name the winner, it is to convince the losers that they lost [1]. In retrospect, this is fairly obvious; while it is generally easy to convince the announced winner that the election was conducted honestly and accurately, the losers are frequently skeptical. It is generally the announced losers who will look carefully at the results, seeking out any discrepancy and suspicious of all assurances of honesty provided by the experts and officials.

In most areas of computer science, we are willing to accept systems where a proof that the system is correct is more complex than the system itself. When we do not forgo such a proof entirely, we usually rely on trusted experts to certify that the system meets a sufficient standard of correctness, without formal statement of either the standard of correctness or the means used to demonstrate that the system is sufficient.

This approach cannot generally meet the goal of convincing the losers in an election that they lost because the losers in every election know that every participant in the electoral process have a vested interest in the outcome. If only a very small and privileged class is able to verify, to their own satisfaction, that they system operates correctly, then it is generally possible that all of the members of this class are partisans of the announced winner and therefore uninterested in disclosing weaknesses or dishonesty in the system.

This is quite different from the situation with most computer systems. Even conventional national security computing is operated under the assumption that the majority of those with access to the machinery are on our side, and that if the adversaries have penetrated the organization, the majority of those with access remain on our side. In this situation, unlike voting, it is prefectly reasonable to trust a small number of experts to check each other's work and certify that the systems operate correctly.

This line of reasoning leads to one conclusion: Whatever technology is used in elections must be sufficiently simple and sufficiently transparent that its correctness is apparent to large numbers of observers [2]. This does not demand abandonment of technology, but rather, careful use of technology.

In discussions with David Chaum and Jim Adler about the cryptosystems underlying their elections, I have suggested, on several occasions that what we need are cryptographic solutions where the proof of the correctness of the system is accessible to a bright high school math student. The classic exclusive-or cypher with a one-time pad is at this level of complexity, but I am not convinced that public key cryptography meets this test.

Asking that the demonstration of the integrity of an election system be sufficiently simple that a bright high school student can understand it does not require public disclosure of all of the entire system if the assertions to be proven are sufficiently well formulated and the mechanisms used to enforce these assertions are carefully selected. The following example illustrates this:

Consider the problem of demonstrating that the computer system running the election management system in the county elections office cannot be attacked through the Internet. Many voting system vendors have simply asserted that their systems cannot be attacked this way because they are not connected to the Internet. I have long argued that this defense is false because the export of data from the election management system to the Internet requires a connection, whether that connection is directly through copper wires or indirect, for example, through *sneakernet*, or hand-carried electronic media.

To demonstrate that attack from the public Internet is impossible, we must demonstrate that the data flow from the election management system to the Internet is strictly one-way. If hand carried electronic media are used to move data between the election management system and the Internet, we must therefore demonstrate that no write operations are performed at the Internet end of the transfer or that the media are physically erased before return to the election management system for rewriting. Proof of erasure for floppy disks can be trivial, if they are degaussed before return to the election management system, but short of this, it is not trivial to demonstrate, to an outside observer, that a disk drive is operating in a read-only mode. Assurances provided by write-protect tabs on disks, for example, are extremely weak.

If the network connection to the outside world is hard-wired, we the proof that it is unidirectional may be quite difficult. For example, for an ethernet connection or a USB connection, proof may be impossible without complete disclosure and examination of the entire protocol stack of the network connection, therefore ruling out the use of proprietary systems.

If we back away from these sophisticated network technologies and use a simple asynchronous data channel, the proof can be quite simple. All we must do is cut all of the wires but the outgoing transmit data wire and the signal ground wire. It is easy for an observer to see that only these two wires remain, but the observer must still assume that the interfaces at each end of the cable are actually standard asynchronous interfaces. We can go farther and move the outgoing data path into the optical domain with an LED on the source side and a CdS Photocell on the destination side. CdS photocells are preferred here because they do not resemble any devices that are capable of radiating data, while phototransistors are packaged identically to LEDs. The circuitry required for this is trivial, and it may be packaged in such a way that anyone with an elementary knowledge of electronics can verify that it is good only for one-way data transmission.

The above example ignores the problems posed by wireless network technology and by Internet over powerline technology. Clearly, we must block these paths as well, for example, by operating the election management a Faraday cage and by use of an appropriate UPS or filtering transformer.

What I believe we must seek is a decomposition of the election problem into technological components where the essential properties of those components are subject to the type of easily accessible proof I have outlined above. This does not require that the internal workings of the system be entirely revealed, but rather, it places each such component inside a shell of easy-to-audit checks. Another way of thinking about this is that we are attempting to reduce the size of the trusted base of software to the point that the trusted base can be entirely disclosed and where the logic of that trusted base is clear to a bright highschool student.

The normal election certification procedures followed by Miami-Dade County provide an excellent example of how entire large system components can be removed from the trusted base. In Miami-Dade county, after all of the data from the precincts has been entered into the election management system and after the election management system has computed the election totals, the final step before certifying the canvass is to compare the printouts of precinct election totals that were created at the precincts with the totals presented by the election management system [3].

This entirely removes the part of the election management system that computes vote totals from the trusted base, but it should be noted that it still serves a useful purpose -- hand processing of printed election totals is notorious for introducing clerical errors, and with this check, a crook would have to corrupt both the electronic and the paper records before the attack would have a chance of success.

Similarly, the push for a voter-verified paper audit trail can be interpreted as an effort to remove the software within the voting machine itself from the trusted base on which the voting system rests. If only a small fraction of the electorate, at random, takes the time to check the voter-verifiable paper ballots (or ballot images) printed by the voting system, the likelihood that inaccurate or dishonest voting software would survive in the polling place plummets [4].

The frequency of clerical errors in manual vote tallying is such that we must take very seriously the old maxim of secure computing that whatever a malicious attacker could do could also be done by a careless legitimate user. There are far more careless users than malicious attackers, so the low probability of a really severe careless error by any one careless user is offset by the sheer number of users. It is useful to remember that the number of election administrators required to run a general election in the United States is on the order one percent of the turnout!

Because of the high likelihood of accident, we must build voting systems where the integrity of the system does not rest on a single mechanism and the assumption that that mechanism is perfectly administered. Instead, we must adopt a policy of defense in depth, where multiple layers of defense protect against the failure of any given layer [5].

As an illustration of the use of defense in depth in voting systems, consider the problem of canvassing the election, consolidating the vote totals from the precinct into an overall vote total. In the above, one of the suggested procedures removes the election management system from the trusted base of software, while another suggestion protects the election management system against intrusion. If we did not care about defense in depth, one or the other of these defenses might be sufficient, but a defense in depth policy suggests that we should adopt both, if at all possible.

Finally, open standards for data representation offer another possible defense. If the data formats used for election data reporting and election setup are sufficiently transparent and are disclosed to the public, third-party tools can be developed that allow observers to independently verify election results. If we require the publication of all of the relevant election data, including the election configuration files and the raw precinct-level data, then we will extend, to election observers, considerable rights that they have not had since the dawn of computerized vote tabulation.

References

[1] Quoted by: Pete Slover, Some Texas counties are clinging to the chad, *Dallas Morning News*, March 8, 2004.

[2] A less general formulation of this rule is presented in: Douglas W. Jones, Auditing Elections, *Communications of the ACM*, *47*, 10 (October 2004) 46-50.
➠ http://portal.acm.org/citation.cfm?id=0922594.1022622

[3] See part 5 of: Douglas W. Jones, *Recommendations for the Conduct of Elections in Miami-Dade County using the ES&S iVotronic System*, report to Miami-Dade County revised June 7, 2004.
➠ http://www.cs.uiowa.edu/~jones/voting/miami.pdf

[4] Adler, J. Confidence: What it is and how to achieve it. *NIST Symposium on Building Trust and Confidence in Voting Systems*. December 2003, Gaithersburg, MD.
➠ www.votehere.net/papers/NIST_121003.pdf

[5] *Defense in Depth*, National Security Agency Security Recommendation Guides, number 1.
➠ http://nsa2.www.conxion.com/support/