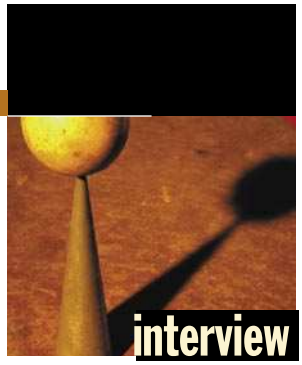

A Conversation with Douglas W. Jones and Peter G. Neumann

Does technology help or hinder election integrity?

Douglas W. Jones and Peter G. Neumann have long been active participants in promoting integrity in the election process, with special emphasis on the dependable use of information technology, as well as on the weak-link nature of the entire process, from beginning to end.

Elections form the fundamental basis of all democracies. In light of many past problems with the integrity of election processes around the world, ongoing efforts have sought to increase the use of computers and communications in elections to help automate the process. Unfortunately, many existing computer-related processes are poorly conceived and implemented, introducing new problems related to such issues as voter confidentiality and privacy, computer system integrity, accountability and resolution of irregularities, ease of administration by election officials, and ease of use by voters—with many special problems for those with various handicaps. Overall, the issues relating to computer security provide a representative cross-section of the difficulties inherent in attempting to develop and operate trustworthy systems for other applications. These issues, of course, have relevance internationally and are increasingly timely.

This interview attempts to capture some of the most basic problems and potential solutions. We explicitly recognize the end-to-end nature of the election process—from voter registration to voter authentication to ballot casting to vote counting and results distribution—in which each step has potential vulnerabilities. Here, however, we focus primarily on the issues related to the use of computer systems, although we begin with some questions that put the technological problems in the context of the overall election process.

PETER G. NEUMANN To what extent is election integrity a technological problem, in contrast with a socio-economical, political, or other kind of problem?

DOUGLAS W. JONES It is certainly all of those. Even pure hand-counted paper ballots raise technical questions. Whatever the technology, any attempt to scientifically investigate elections has unavoidable political implications. I have described this as "very political science"; work in this area is as politically loaded as work on evolution or stem cells. Merely claiming that research into election integrity is needed is seen by many politicians as challenging the legitimacy of their elections.

PGN Let's start with the big picture. What are the most important technological issues that we need to consider?

DWJ Much of the difficulty stems from the fundamental requirement that the election system be transparent. To borrow a phrase from Dan Wallach, associate director of ACCURATE (A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections), the system must convince the losers that they lost. Those who lost an election and their partisans typically have no required technical qualifications, so the entire system must be sufficiently open and comprehensible that nontechnical observers can believe the results. Furthermore, the losers in an election have no reason to believe any assertions made by a government that is run by the winners.

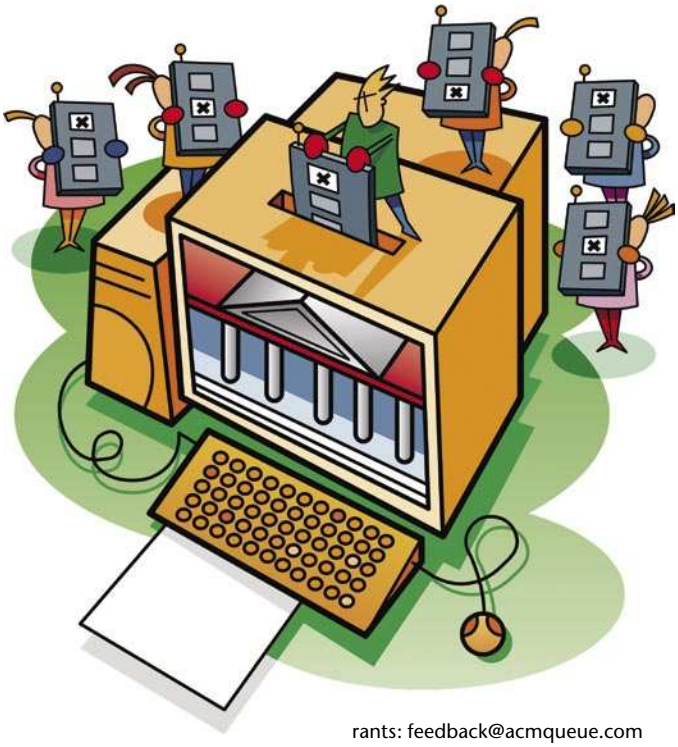

Generally, systems become more trustworthy when redundancy is used to protect against internal errors, accidents, and attacks, but the mere presence of redundancy offers no guarantees unless the placement and transmission of the copies are carefully planned and there are clear procedures for identifying and resolving discrepancies between the copies.

Vote counting is an accounting function, and as such, the accounting standards for voting ought to be comparable to those for money. We ought to design auditability into our voting systems. So long as we are forced to trust the software of voting systems, we need a tightly controlled software development process, so that all code is attributable both to the specifications it fulfills and to the programmers who wrote it. We also need secure authentication of hardware, software, configuration files, and results to defend against both fraud and accident. In addition to self-authentication, we need trustworthy paths for transport of all system components.

Finally, we need to learn how to evaluate whether our systems actually meet design requirements. Certification by anonymous bureaucrats or government contractors operating behind closed doors cannot be expected to convince the loser in an election. Can you imagine what would happen if the NSA (National Security Agency) certified the security of voting systems? The conspiracy theorists would have a field day. Even so, NSA certification would be far stronger than the current approach to voting-system certification.

PGN What about other issues, such as preventing insider misuse?

DWJ One of the problems in public discussions of voting-system integrity is that the different participants tend to point to different threats. Election-system vendors and election officials generally focus on effective defense against outside attackers, usually characterized as hackers. Meanwhile, many public interest groups have focused on the possibility of election officials corrupting the results.

Over the past two centuries of American history, we have seen ample evidence of attacks from both directions, but it appears that the vast majority of election fraud has involved corrupt jurisdictions protecting their positions. This pattern has appeared repeatedly in big-city political machines, as well as in rural areas.

PGN What can be done to address all of the seemingly nontechnological problems? Might technological approaches such as registration databases help or hinder?

DWJ Both, of course. It is clear that technology can be used as fairy dust. Crypto fairy dust has a long record in the world of elections. The most famous case involves the voting machines made by Global Election Systems (later Diebold), where all the machines had the same DES encryption key hard-coded into their software. Use of cryptographic authentication of object code falls into the same category if you expect the voting machine to self-authenticate by printing out the authentication code of the software it's running. You advance beyond the fairy-dust-stage only if you allow an external observer to examine the contents of memory and independently compute the authentication code.

At the same time, it is clear that the technology exists to overcome many of the defects of current electronic voting systems. The pyramid-of-trust model used to authenticate casino game firmware is very well tested (and patented, see U.S. Patent 6,149,522). We know how to use both symmetric-key and public-key cryptosystems to authenticate all data flow both to and from the voting machine. We know how to automate the key-management problem for a flock of voting machines so that the cryptographic keys are reset to new random values for each machine for each election without placing a huge burden on election administrators.

The problem is, there are other things we cannot do. The proof that there is no effective software to detect all malware has a structure very similar to that of Turing's proof that the halting problem cannot be solved. Assume you have a malware detector that never misidentifies well-behaved code as malware and never identifies malware as being well behaved. Create an application that incorporates your malware detector. If that application detects malware, it behaves entirely honestly. If it does not detect malware, it becomes malware itself. Now, submit your application to itself. The malware detector is obviously wrong, so our initial assumption that it worked was wrong.

We cannot reliably inspect a program to determine that it does not contain hidden functionality, and it is impossible to guarantee detection of all hidden functionality by black-box testing. The prevalence of "Easter eggs" in commercial software demonstrates this. Whereas some Easter eggs may be intentional tools used to detect illegal copying, others are clearly examples of unauthorized functionality that has slipped through the quality-control tests at the vendor.

Finally, of course, no highly technical measure is going to be entirely convincing to a naive observer. We therefore need to emphasize understandable measures. In the area of cryptography, for example, the work being done by people like David Chaum to develop cryptographic election protocols that can be explained to those of us

without Ph.D.s is promising, but it seems to me that they have a long way to go. I see no hope of reducing software correctness proofs for paperless voting machines to a form that average voters should trust.

PGN By the way, it was Doug Jones who in 1997 discovered the DES key hard-coded into the software and reported it. Ironically, despite many changes since then in the software, that very same key is reportedly still in the software. (Only recently was the election administrator given the option to change the key manually!)

Moving on, what issues have we not yet addressed?

DWJ The centralization of voter registration records provides interesting opportunities. The anarchic system of voter registration record keeping, which until this year was common nationally, was awful. In many jurisdictions, a large fraction of the records were long out of date, and no two counties seemed to keep their records in the same format. Now, HAVA (Help America Vote Act of 2002) has forced migration to statewide registration databases. This promises to improve things, but it also means that mismanagement, when it occurs, will have statewide consequences. One common problem is already appearing in many states—namely, that of matching driver's license records with voter records. As it turns out, many of us haven't been entirely careful to spell our names identically in our different official government records, and some states have opted to go with very bad name-matching software in an effort to clean up the new statewide registration databases that are mandated by HAVA.

One lesson has struck home repeatedly. The consequences of human error and fraud are frequently indistinguishable. An investigation of voting-machine event logs in Miami-Dade County in a minor municipal election in 2003 showed that some vote records came from machines that had not been deployed at the polling place. When I was asked to investigate this, my first hypothesis was that someone had fraudulently used a machine in the county warehouse to record the votes that should have come from the precinct. The problem with this hypothesis was that the election was so minor and the outcome so expected that fraud didn't make any sense. Occam's Razor was wrong: Working with the vendor and the county, we eventually determined that the misrecorded serial numbers were the result of an unlikely sounding coincidence involving a fairly common administrative error (failing to plug in the machine) combined with a programming error (i++ instead of ++i during the recording of the low-battery message in the event log) combined with a software design error (failing to use the redundant memory format properly).

PGN What about operational risks?

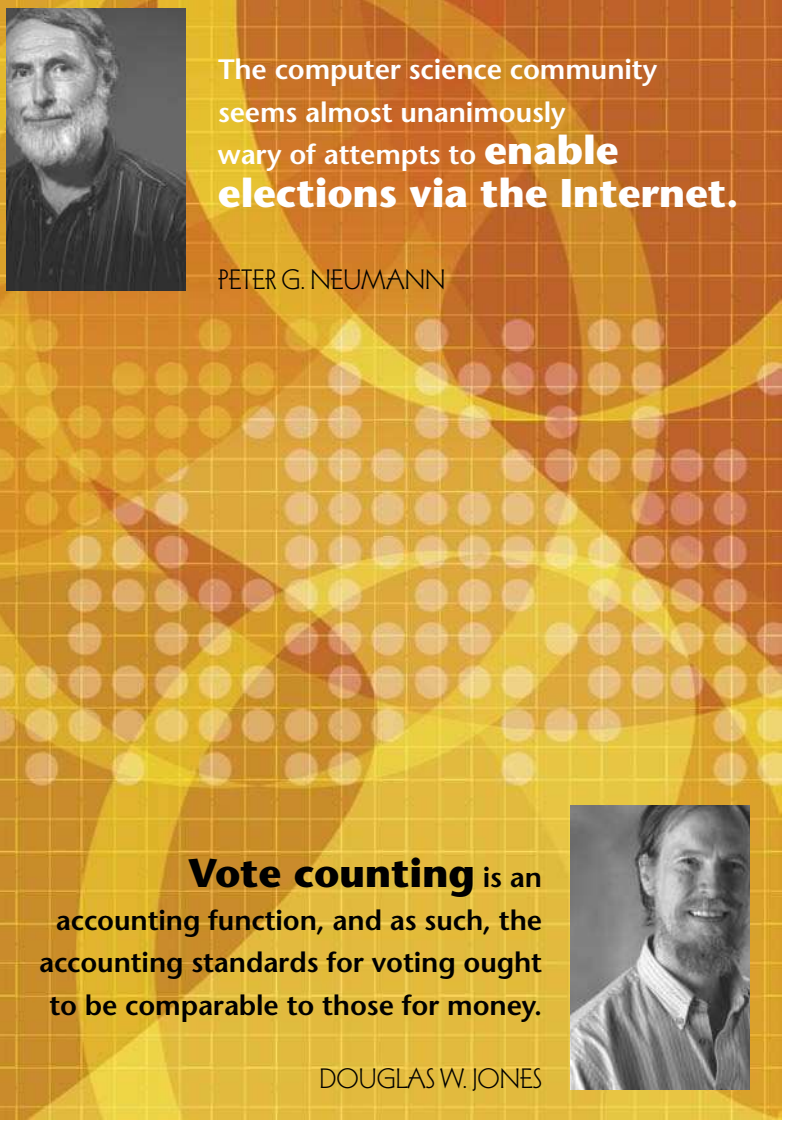

Photo courtesy of University Relations–Kirk Murray

DWJ Voting-system administrators walk a tightrope, balancing the cost of obtaining in-house expertise with dependence on vendor technical support. In-house expertise is expensive, sufficiently so that many small jurisdictions will never be able to afford to operate without outside support. Over the past century, it is clear that voting-system vendors have generally made more money selling election support services than selling the machines themselves. Service contracts may involve as little as preventive maintenance or as much as complete election support—relieving the county of the need for any local technical expertise.

The intermittent nature of election work makes it almost impossible for anyone to dedicate full-time positions for all of the different technical jobs involved in elections. Much of the election workforce must, therefore, be made up of temporary employees. Voting-system vendors transfer large numbers of workers from other roles in the company into the field as election consultants for the duration of each election season; in jurisdictions that rely on local expertise, it is common to draft computer programmers and technicians from other government departments. In addition, both the vendors and the voting jurisdictions depend on hiring and training temps.

Whether dealing with temporary employees or contractors, it is clear that election officials face some serious problems. At what point do you draw the line between functions that are performed only by trusted permanent employees of the jurisdiction and functions that you allow contractors or temporary employees to perform? Do you allow them to handle ballots? Do you allow them to do on-site preventive maintenance? If something breaks, what repairs do you permit them to perform?

PGN You mention making the process more open. What procedural steps might help to that end, particularly those that could overcome potential insider attacks?
DWJ One of the biggest problems with the current regulatory framework is that it contains poor provisions for feedback. When failures occur in our election systems, they are reported in an ad hoc way, and it is not at all clear that those who are involved with formulating or applying our voting-system standards have ready access to information about system failures. We need to close the loop, so that voting-system failures are routinely investigated and the results of those investigations are made available to the public, to certifying agencies, and to those in charge of further development of the voting-system standards.

The definitions we use have some startling effects on openness. For example, consider the definitions of the words *programming* and *software*. We in computer science take the meanings of these words for granted, but they have come to have different meanings in the context of elections. Voting-system software is generally the proprietary intellectual property of the voting-system vendor and not subject to public inspection. This may itself be a problem, but the creeping definition of the word *software* has led some jurisdictions to object to public disclosure of the contents of election configuration files. These files do not contain (or at least ought not to contain) anything proprietary to the vendor; in theory, their only content is descriptions of the ballots and voting rules for the jurisdiction. Unfortunately, creating the configuration files to set up a voting system for an election has come to be described as *programming* the election, and these files have come, therefore, to be described as software.

Another problem revolves around the interpretation of public records. All of the records of an election are gener-

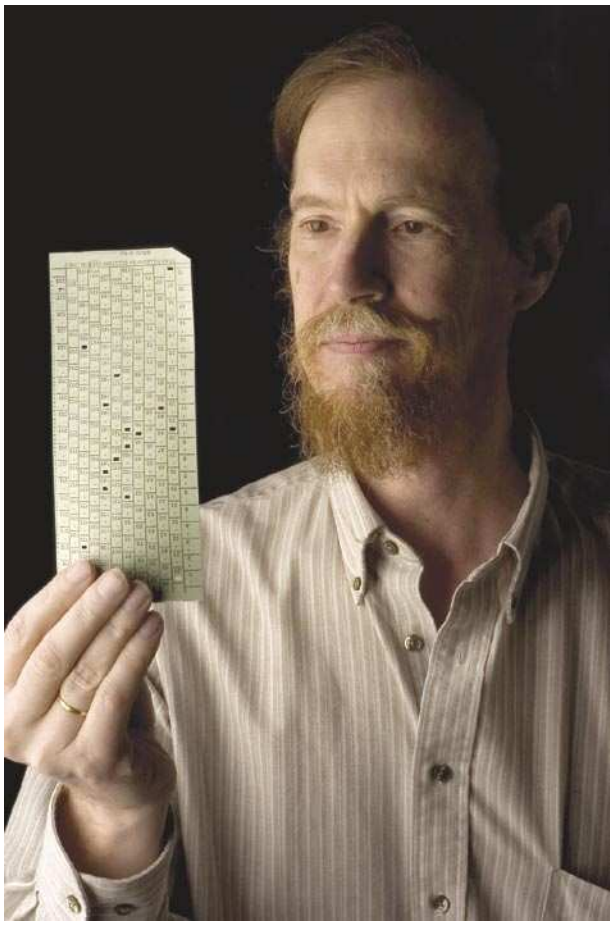

Photo courtesy of University Relations–Kirk Murray

ally public, but many of these records are stored in proprietary data formats. Even the ink-on-paper records printed by election-management systems are frequently loaded with cryptic abbreviations. Unfortunately, the documentation necessary to interpret these records is generally considered confidential. This is absurd. It is as if the public records were being made available only in encrypted form. As far as I am concerned, the documentation necessary to interpret any public records must itself be public.

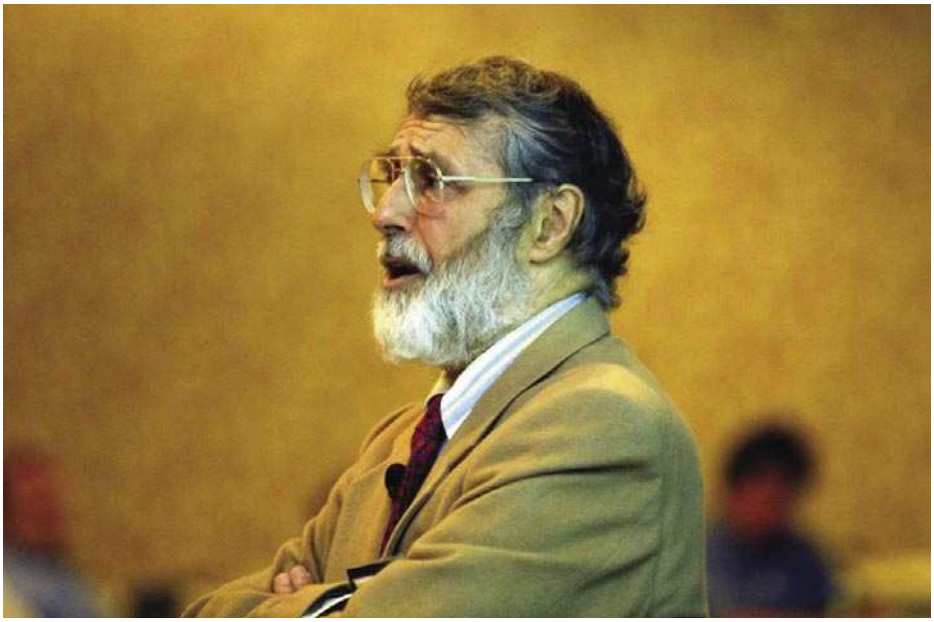

PGN The computer science community seems almost unanimously wary of attempts to enable elections via the Internet—for example, the SERVE (Secure Electronic Registration and Voting Experiment) report. To what extent could better computer-system security and sound uses of cryptography satisfy the necessary requirements?

DWJ A good part of the problem is that most Internet voting proposals involve absentee voting systems that happen to use the Internet instead of postal ballots. Absentee voting, where voters fill out their ballots outside the controlled context of polling places, allows voters to easily prove, to a crook, that they filled out their ballots as required. This is why I have long recommended increased use of satellite polling places for early voting instead of unrestricted postal voting.

Internet voting from the home has many additional problems. Browser neutrality is one. Several current schemes operate correctly only under Microsoft Internet Explorer. I don't believe that such limits are appropriate for any official government use of the Internet.

With Internet voting from the home, we have no control over what spyware may be present on the voter's machine. Even if the voting application is distributed as a bootable system image, we cannot prevent it from being run under an emulator. Emulators that accurately simulate the realtime behavior of the emulated system are entirely possible—I've written one!

On the other hand, I see no reason that the Internet cannot be used for any functions that currently use wireless systems or other public networks such as the switched telephone network. Use of the Internet for reporting unofficial results from the precinct to the central election headquarters would seem to be no less safe than the wireless schemes used in some jurisdictions today. We know how to authenticate such transmissions, although I am unaware of vendors that are doing this correctly, and we know how to assure that they do not corrupt the official results: Just print the official results to paper and extract machine-readable media from the voting machine before inserting the communications card or cable in the machine. Many jurisdictions already get this right.

PGN "No less safe than the wireless schemes?" Wow, that is damning with faint praise. How unsafe are the wireless schemes used today, which are typically remotely spoofable and certainly misusable by insiders?

DWJ I have no great confidence in wireless transmission, but it has been used for several years in some jurisdictions. In urban areas, for example, it is popular to set up polling places in building lobbies. It turns out that the lobby of a typical high-rise building is the one place in the building that frequently has no phone lines, so if the election administrators want a fast electronic report of the results from each polling place, the obvious way to do it is using wireless technology.

There are security advantages to both hand-delivering results from the polling place and electronic delivery, even wireless electronic delivery. Copies delivered by hand are necessarily delivered slowly, allowing an

attacker the time needed to construct a counterfeit. This applies equally to paper and electronic media. Immediate electronic delivery over a public wired or wireless network denies the attacker the time to carry out an elaborate attack, assuming that the keys needed to authenticate the data have not been compromised.

PGN In many countries, voters still vote with a single piece of paper and have thus far resisted the use of high-tech approaches. Other countries seem to be pursuing various levels of computerized voting. Do you see any trends around the world for or against the potential "technologization" of elections?

DWJ It is important to understand that the U.S. system of elections places more contests on a single ballot than any other country in the world. In Clay County, Iowa, a relatively small rural county, the November 2000 ballot held 20 distinct contests, ranging from President of the United States to County Agricultural Extension Council, as well as 11 judicial retention questions and one referendum. Hand-counting ballots of this complexity is far more difficult than hand-counting ballots in a parliamentary democracy where there is only one race on the ballot, for member of parliament. This complexity is why the United States began applying machinery to vote counting a century ago.

Electronic voting technology has been used in some countries to deal with other complexities. For example, Belgium uses a hybrid party list scheme, where voters may vote for either a party list or for individual candidates. The complexity of Belgian rules for voting and ballot counting, even when there is only one race on the ballot, invites the use of computers; as a result, electronic voting has been used extensively in Belgium.

The canton of Geneva, Switzerland, has allowed Internet voting for several years in order to allow the large number of Geneva citizens who work overseas to cast ballots. As in the United States, the Swiss typically hold as many as four elections a year, and they routinely place multiple issues on the ballot for some of their elections (although rarely as many as 10).

India has what must be the world's simplest paperless electronic voting machines, used nationwide. Of all the electronic systems in use today, it comes the closest to being simple enough that it might actually be possible to prove its integrity.

Brazil also uses a paperless electronic voting system nationwide, one that is far more complex than the Indian system, in part because the ballot access rules used in Brazil routinely allow for hundreds of candidates running for a single office.

Russia, in contrast, has opted to use optical mark-sense voting, although apparently, its rules forbid hand-counting of the ballots, making the system as potentially dangerous as a purely electronic system.

In November 2005, I participated in the Organization for Security and Cooperation in Europe election-observing mission in Kazakhstan. The e-voting system developed in Kazakhstan includes some very interesting innovations involving stateless voting machines and use of smartcard technology. Unfortunately, these innova-

Resources

For further background, the following should be useful:

Special issue on voting systems, *Communications of the ACM* (October 2004).
Jefferson, D., Rubin, A. D., Simons, B., Wagner, D. 2004 Analyzing Internet voting security. *Communications of the ACM* (October).
ACE Focus On; http://focus.aceproject.org/e-voting/countries. This is a useful but not definitive source for e-voting around the world.
Major scandal evolves around e-voting machines in Russia; http://english.pravda.ru/russia/politics/22-12-2004/7529-0.
Organization for Security and Cooperation in Europe; http://www.osce.org/item/18133.html.
Rubin, A. 2006. *Brave New Ballot*. Random House.
Web site for A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections (ACCURATE); http://www.accurate-voting.org.
Web site for Doug Jones; http://www.cs.uiowa.edu/~jones/voting.
Web site for Peter Neumann; http://www.csl.sri.com/neumann.

tions were balanced by some severe defects. The one that concerns me the most is that the voting-system standards themselves were a state secret. As in the United States, Kazakh voting machines must be approved by an independent testing authority, and the detailed reports by the authority on the system are delivered only to the voting-system developers, with no intention that these ever become public documents.

PGN Would you care to close with some personal comments on your experiences in tilting at this particular set of windmills?

DWJ When the Iowa Secretary of State called for computer scientists to volunteer to serve on the Iowa Board of Examiners for Voting Machines and Electronic Voting Systems, I was afraid that nobody would volunteer, so I did. As it turned out, I was right. Nobody else from Iowa's high-technology sector volunteered. When I resigned after a decade on the board, it was very difficult to find anyone from the high-tech sector willing to serve as my replacement.

I had served as a voting examiner for only a short time when I began finding serious problems in the voting systems that were being brought forward for certification in Iowa. Through the end of the 1990s, I remained quiet about these problems, simply voting against certification of defective systems and quietly telling the vendors about the defects they needed to fix.

Aside from a few postings in the ACM Risks Forum, I said little in public until spring 2000, after I had read the California Internet Voting Task Force Report of January 2000. At that time, I was chair of the board of examiners, and as of the 2000 general election, I was the only state election official making critical comments about voting technology.

This changed my life. Before the election, most of my time was spent on realtime embedded systems. In the years that followed that election, I have had little time for anything other than voting technology. Since then, I've learned to laugh when someone calls me a Luddite or a conspiracy theorist. The word *activist* is more troublesome. As a computer professional, I feel an obligation to advocate responsible computer use and to oppose irresponsible computer use. At the same time, being seen as an advocate endangers my credibility as an academic. Q

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums