OUT OF BAND

Douglas Jones on Today's Voting Machines

Hal Berghel, University of Nevada, Las Vegas

We catch up with computer scientist and voting machine guru Douglas Jones to get a deeper understanding of current challenges in electronic voting technology.

Douglas Jones, a professor in the Computer Science Department at the University of Iowa, has been involved in voting technology research since 1995 and was a principal investigator for the National Science Foundation (NSF)-funded ACCURATE project (A Center for Correct, Usable, Reliable, Auditable, and Transparent Elections; accurate-voting.org). His recent book with coauthor Barbara Simon, *Broken Ballots: Will Your Vote Count?* (CSLI Publications, 2012), is the seminal work in the area of current voting technology and is highly recommended to anyone who believes in fair elections.[1] Much of Jones's professional work is available on his website (www.cs.uiowa.edu/~jones). The "interview" that follows resulted from our email exchanges during July and August 2016.

ELECTION MANAGEMENT SYSTEMS

HAL BERGHEL: You, Aviel Rubin, Bruce Schneier, and many other prominent computer scientists have been highly critical of DRE [direct-recording electronic] voting machine vendors for refusing to build DRE equipment around robust security models. Please provide us with a 2016 status update on the security of these machines.

DOUGLAS JONES: Most of the DRE voting machines being sold today are based on designs from the 1990s. That is to say, there's been little change in DRE voting technology in the past 15 years. Software upgrades over this interval have improved the GUI design significantly, as well as fixed some security flaws, but this has largely been an incremental process. Finally, it's important to note that all of the major DRE voting system vendors have added voter-verifiable paper-trail mechanisms.

In contrast, there's a new generation of optical mark scanners on the market. Whereas the scanners of the 1990s used either discrete component sensors or 100-pixel-per-inch monochrome contact image sensors, the new scanners use high-resolution and, in many cases, color-image sensors originally developed for desktop scanners. Memory has become inexpensive enough that these scanners typically capture full images of each ballot instead of merely a summary of the votes cast.

Another major development involves accessible voting devices for voters with disabilities. DRE machines

0018-9162/16/$33.00 © 2016 IEEE

now have serious competition in this arena in the form of touchscreen ballot-marking devices [BMDs] that allow disabled voters to mark a paper ballot for input into a ballot scanner. All voting systems based on ballot scanners are now marketed in conjunction with BMDs.

The greatest liability faced by today's voting system vendors lies not in the vote capture technology, whether DRE or scanner based, but in the election management systems [EMSs] used to configure the vote capture system and accumulate precinct totals. These frequently include legacy support for the full range of voting systems sold by the corporate predecessors of the current vendors. If some county is still using a system, continuing support is required, and it's more expensive to strip out code for a system no longer in use than to retain it. As a result, the code in these EMSs tends to grow larger and more brittle with each passing year.

BERGHEL: Electronic voting machines have been widely discussed, but I've seen very little discussion of EMSs. What are the greater security vulnerabilities in these systems? Have computer scientists ever analyzed any? If so, what did they find? Is there any reason to assume that the back end is secure enough to prevent fraud from election insiders, network attacks, and so on?

JONES: A typical EMS contains a database that holds all the machine settings required to configure the voting system to meet local election laws, plus the mapping from precincts to election districts, the offices up for election in each district, and the candidates for those offices. Before the election, the EMS automatically generates configuration files from this database for each DRE machine or ballot scanner, and after the election, the EMS consolidates the totals from each machine to produce jurisdiction-wide results.

In most cases, configuration files are written to removable media such as compact flash cards for transfer to voting machinery, and official election results are returned on the same media. There's immense pressure from the news media for rapid reporting of unofficial results, so most EMS vendors offer modem banks so that voting machinery in the precinct can report by modem after the polls close. Similar pressures ask election offices to report election results on the Web, so there's frequently a data path from the EMS to the jurisdiction's webserver.

Because of its central role in preparing for an election and aggregating the returns, a compromised election management system is very dangerous.

Because of its central role in both preparing for an election and aggregating the returns, a compromised EMS is very dangerous. It has the potential to misconfigure all the voting equipment in the jurisdiction, and it can potentially alter the election results after the polls close.

Election officials frequently respond to allegations of software vulnerabilities by reassuring the public that voting equipment isn't connected to the Internet. For the machinery in a precinct, this is generally true. However, almost all of the machinery in a precinct can be equipped with modems, and the EMS can have a modem bank. Reporting returns to the Internet can be air-gapped with hand-carried media, but in the past many counties have had network connections from the EMS to a webserver.

The defense against an outside attack therefore depends on procedural defenses such as printing the official precinct totals and writing them to removable media before connecting [to the modem] to upload unofficial totals, and doing a cold start and restore from backup on the EMS after turning off the modem bank before processing the official results. Just as paper ballots from randomly selected precincts can be hand counted to detect miscounts in ballot scanners, the paper records of precinct totals can be reconciled against the totals reported by the EMS. Numerous jurisdictions have done this routinely for decades, but it appears that there are many that still don't take these precautions.

In general, EMSs haven't been subject to the scrutiny that DRE voting machines have faced. In part, this is because they aren't as widely available. When jurisdictions replace vote tabulators and DRE machines, the old ones have sometimes been sold at government surplus auctions, where they become available to researchers. EMSs generally run on commodity computers, and when these go to surplus, their disks are routinely scrubbed.

In addition, the major focus of researchers has been on those parts of the election system that aren't software independent. MIT cryptographer Ronald Rivest and NIST researcher John Wack coined the term "software independent" to refer to voting systems in which we don't need to rely on the correctness of the software to assure ourselves that the results are correct. Paperless DRE voting systems are

FURTHER READING

For those interested in further information about today's digital election systems, the definitive book on the subject is *Broken Ballots: Will Your Vote Count?* (CLSI Publications, 2012) by Jones and Barbara Simon.

An introduction to this topic was also presented in last month's Out of Band column (vol. 49, no. 9, 2016, pp. 104–109), and an overview of the various categories of election fraud (versus imaginary voting fraud) can be found in the January 2016 column ("Digital Politics 2016," vol. 49, no. 1, 2016, pp. 75–79).

purely software dependent, whereas paper-based systems are subject to hand recounts and audits that can, in principle, defend against malicious or erroneous software. With the procedural defenses outlined above, we can defend against a faulty or corrupt EMS. These procedures have been used for decades in some jurisdictions and are required by law in some states, but their use is far from universal.

CHALLENGES OF OPEN SOURCE

BERGHEL: If there were any application of computing that cries out for high-confidence code, it's the voting machines that determine our nation's future. This is precisely the sort of application in which open source code excels. However, DRE voting machine equipment is proprietary: neither open source nor high-confidence. How did we get to the point that the public finds this acceptable?

JONES: In the first place, DRE voting systems predate the open source software development model. The first DRE voting machine sold commercially was the VideoVoter, first deployed in 1975 by a predecessor of Election Systems and Software. By 1990, the DRE marketplace was vibrant, with several vendors offering a range of machines, and it wasn't until the 1990s that research began to demonstrate that open source software was, on the whole, more robust and secure than competing proprietary software.

There's a second problem with open source software, and that is that it might not be the right model. In 2003, I helped found the Open Voting Consortium [OVC] in hopes that it would create a framework for open source voting system development. The OVC still exists, but to this day, we don't have a consensus on how an open source voting system development framework should function. The problem is, you can't just invite everyone to contribute code; you need tight controls over what goes into the final product. This applies to all security-critical code. At this point, I'm convinced that what we need isn't open source voting code, but a disclosed-source model. That is, vendors should rely on copyright and patent law, not trade secrets, to protect their intellectual property rights. The problem with this is that any vendor that relies on trade secrets can copy its competitor's code with impunity, so how do we manage the transition to a disclosed-source model?

Researchers interested in studying current voting systems face several legal barriers. It's not clear that it's legal to reverse-engineer software or to experimentally test it for the purpose of assessing software security, even if this evaluation is critical to the public interest. Recent stories about the legal barriers to this have focused on the Volkswagen emissions control scandal, but it's clear that the same questions are relevant in the election domain.

BERGHEL: Independent Testing Authorities [ITAs] and the Voting System Testing Laboratories [VSTLs] that replaced them are approved by the government to certify that voting systems meet the federal Voting System Standards and the more recent Voluntary Voting System Guidelines [VVSG; www.eac.gov/assets/1/Documents/VVSG.1.1.VOL.1.FINAL.pdf]. However, these organizations are paid by the manufacturers seeking the certification, and negative results aren't reported to the public. This appears to go beyond conflict of interest all the way to creating a moral hazard. What should be done to ensure legitimate certification [note that the Diebold AccuVote TS system that was easily hacked was certified by an ITA]?

JONES: The ITA and VSTL models closely parallel the product-testing and -certification models used in a wide range of industries. Manufacturers of electrical products pay for UL testing. Medical apparatus manufacturers pay for the testing needed to get FDA [US Food and Drug Administration] approval. Manufacturers of airplanes pay the cost of airworthiness certification. So long as products are developed and manufactured by for-profit private companies, it makes good sense that they should pay the price of bringing the products to market.

The problem with the current situation is that, in these other industries, there are strong feedback loops in the regulatory system. Defects in electrical products lead to insurance claims, and UL is the creation of the insurance industry. Medical professionals have strong incentives to report failures and side effects to the FDA, and every incident in the aviation industry is reported to the FAA [US Federal Aviation Administration]. Regulators in these fields respond very rapidly to reports of problems.

In contrast, local election offices have strong incentives not to report problems. Public disclosure of failures in voting systems reduces voter confidence in the integrity of our democracy. Currently, the Election Assistance Commission [EAC] requires voting system vendors to report all problems with voting systems certified to meet the EAC's VVSG, but the VVSG update process is extremely slow and the threshold of what constitutes a reportable problem appears to be rather high.

One positive change we have seen in the past decade is the move by the EAC to routinely post VSTL reports on their website. This is a major change from the era of confidential ITA reports that were rarely available to the public.

BERGHEL: Any serious student of human factors understands how important ballot design is to ballot effectiveness (for example, to avoid unintentional undervoting and accidental vote flipping, voter confusion, banner blindness, and so on), and yet there seems to be no attempt to set standards for ballot layout in the 2015 VVSG [see Section 3 of the VVSG: Usability, Accessibility, and Privacy Requirements]. Am I missing something or is this a glaring failure of the EAC?

JONES: The voting system guidelines are written with an understanding that state laws largely dictate the details of the presentation of the ballot. State laws have frequently *required* horrible presentations, and the federal government is largely powerless to intervene unless you can show discriminatory consequences under federal civil rights or disability rights laws.

There is a glaring failure here, but the root of the problem is congressional. The Help America Vote Act of 2002 [HAVA] that established the EAC contains this text: "The error rate of the voting system in counting ballots (determined by taking into account only those errors which are attributable to the voting system and not attributable to an act of the voter) shall comply with … [VVSG Section 301 (a)(5)]." That is to say, human factors are explicitly excluded from any discussion of the accuracy requirements.

Section 3 of the 2005 VVSG tries hard to address usability within the scope permitted by HAVA and the range of state requirements, but the emphasis is on accessibility. It's likely that more can be done under the current legal framework, but it will probably take a change to this framework to properly address the issue.

BERGHEL: The 2000 US presidential election in Florida illustrated the dangers of having political partisans serve as chief election officials. What are your thoughts on how we might depoliticize the office of chief election official in the US?

JONES: I distrust suggestions that you can simply require that election administration be depoliticized. The problem is how to do this. In a democracy, it verges on irresponsible for a person not to have political opinions. I would much rather know the politics of the people running our elections than have them hide their politics. So, the problem isn't how to depoliticize elections, it's how to manage the fact that people are inherently political.

In states with good civil service systems, it's possible to erect a fairly solid firewall between the elected and partisan appointees and the actual administration of elections. The other alternative is to rely on mutual distrust, requiring that representatives of both parties be involved in all critical decisions. This works reasonably well in a balanced two-party democracy, but it becomes unwieldy as the number of parties grows; and, because it relies on mutual distrust, it breaks down badly where there are partisan coalitions or when one party is significantly more powerful than any others.

There are strong incentives not to report problems—public disclosure of failures reduces voter confidence in the integrity of our democracy.

BERGHEL: Let's discuss the two models of election secrecy for a moment. The British model holds that the ability to recover the individual voter's preference is a state secret. What you call the "absolute secrecy model," which is the default in the US, holds that no information can be retained that would allow any observer to determine a particular voter's preferences. Computer scientist Michael Shamos faults VVPAT [voter-verified paper audit trail] systems as egregious violations of the voters' right to a secret ballot. Does Shamos's observation speak in favor of eliminating VVPAT systems altogether, or to moving to the British model of election secrecy? Is there a middle ground?

JONES: The generation of VVPAT systems that were introduced after the 2000 US presidential election used continuous rolls of thermal-printer paper to record a paper trail. Shamos is correct that these prevent absolute ballot secrecy. There's also ample evidence that the number of voters who read the VVPAT on these machines is small enough that they're not very good at achieving their stated purpose.

There are two answers to the middle-ground question: first, a pair of scissors. Ideally, the VVPAT could be snipped after each voter's record is printed inside the voting machine. Many modern receipt printers can do this. Alternatively, before any person is allowed to look closely at the VVPAT contents during an audit, it could be snipped into segments by hand to achieve anonymity.

Second, we can create cryptographic links between voter and ballot. A number of proposals for end-to-end [E2E] cryptographically verifiable elections do this with multiple key custodians. The key custodians must cooperate to decrypt the ballots, but voter privacy is assured so long as just one key custodian does not join in a conspiracy to violate that privacy. At this point, there aren't any E2E systems that would meet the requirements for a public general election using DRE or Internet voting, but several are in widespread use in less critical contexts.

TECH EXPERTISE IN ELECTIONS

BERGHEL: You mentioned in your book that Iowa statute requires that at least one of the Board of Examiners for Voting Machines and Electronic Voting Systems "… shall have been trained in computer programming and operations." [Note that Jones once held that position.] This requirement seems beyond eminently sensible. How might other state legislatures be incentivized to create similar laws?

JONES: Some of them already do, but this isn't necessarily a successful requirement. In Iowa, when they asked for volunteers from the tech sector to serve on the Board of Examiners, I was the only volunteer. When I told Shamos this story, he said that was exactly how he got on the Pennsylvania Board—in his case, there were three openings and exactly three volunteers.

When I volunteered to serve as an examiner for Iowa's voting machines, I significantly overestimated the technical competence of the vendors, and I seriously misestimated where the problems would be. I expected interesting cryptography and interesting embedded systems. I didn't expect to see system failures that were dominated by human factors and amateurish software development methodologies.

In most states, voting system examination is essentially a volunteer job with a token reimbursement that might have been significant a century ago. It took me years to reach the point where I felt confident in my criticism of the process and the marketplace. Not many people who have the technical expertise can make this commitment.

Several states hire outside consultants to evaluate voting systems. This model would make sense if there was a pool of outside consultants who were both well informed about the current state of voting systems and free of entanglements with the voting system industry. Unfortunately, such a pool is hard to identify.

> Elected officials at all levels are reluctant to face any questions about the election system that put them in office.

BERGHEL: On a personal note, several computer scientists and election officials have experienced firsthand the wrath of electronic voting equipment manufacturers, ITA executives, and the leadership of influential special interest constituencies for speaking out about insecure voting systems. In fact, attempts to censor or silence both you and Rubin were directed to the presidents of your respective universities, and at least one election official in Utah was forced to resign for allowing Diebold equipment to be inspected by computer security experts. Of course, truth is always disadvantaged when it confronts power, but elections are so important that it would seem that a special case should be made to protect experts, officials, and whistleblowers. What are your thoughts?

JONES: In both my case and Rubin's case, our institutions did an excellent job of responding to the attacks. Working in academia has its advantages.

It is much harder to protect voting system administrators who raise unwelcome questions about the systems they're using. Elected officials at all levels are reluctant to face any questions about the election system that put them in office. When there are suggestions that the voting system is flawed, common defenses include shifting the focus. For example, politicians love to talk about [protecting against] voter fraud, while most election fraud has been instigated by the struggle of ruling parties to preserve their status in the face of voter discontent.

Computer scientist Dan Wallach at Rice University pointed out that those who lose elections are the ones who ask the hard questions, while the winners generally prefer that their victory go unquestioned. Short of broad-based public outcry and blatant misconduct, election officials willing to expose voting systems to close scrutiny by outside investigators will invariably place their jobs on the line.

BERGHEL: Your book quotes Rivest: "Coming up with 'best practices for Internet voting' is like coming up with 'best practices for drunk driving.' You really don't want to go there." Let's close with your current thoughts about Internet voting.

JONES: Internet voting faces two huge problems: Internet security and human factors.

Questions of Internet security have received more attention in recent years. There's an almost constant drumbeat of reports about government databases that have fallen to malicious hacking, and there's no reason to believe that voter databases, election configuration databases, or election result databases are immune to this threat.

Proposals for E2E cryptographic, voter-verifiable elections are interesting in this context. If voters could compute elliptical polynomials in their heads, these cryptosystems might actually solve the security problems, but real people can't do this. As a result, the cryptography must be done on the voter's computer, and done by software that, ultimately, the voter cannot be sure of. So long as voters' personal computers are vulnerable to malware, there's no guarantee that the vote reported to the EMS is the same as what the voter intended.

And then there's the problem of human factors. All Internet voting systems are, at heart, DRE voting systems where the Internet replaces the memory cartridge used to communicate with the EMS. I've run experiments on DRE interfaces at the University of Iowa, and David Byrne has run even more comprehensive experiments at Rice University that show significant error rates when people vote on DRE voting systems. What becomes rapidly obvious is that we're very good at designing user interfaces for routine use, but most voters only vote once every few years. All of our assumptions about how people learn user interfaces and how people develop expertise fly out the window in this context. Voting systems must be accessible to the most technologically unsophisticated without any training. This sets an extremely high bar, and we're not there yet.

HAL BERGHEL is an IEEE and ACM Fellow and a professor of computer science at the University of Nevada, Las Vegas. Contact him at hlb@computer.org.