

Inside Risks

The Risks of Election Believability (or Lack Thereof)

With 90% of the 2020 U.S. general election ballot contents verifiable by paper, why do only 65% of voters trust the results?

DESPITE OR PERHAPS because of COVID-19 health concerns, a record 155 million ballots were cast for President in the 2020 general election, with both the winner and runner-up each individually getting more votes than any candidate in U.S. history. Yet, according to post-election polling,¹¹ only two in three voters felt confident the election was free and fair. Even the voter-verified paper ballot⁹ for direct-recording electronic voting machines (which enables hand-counting via an audit or 100% tally) does not in and of itself create sufficient credibility in the election results.

Trustworthiness in elections is inherently a total-system problem (as considered more generally²). Every part of the overall process (for example, voter registration, ballot layouts, casting and counting, audits, and recounts) provides potential points of compromise. Problems may result from human errors, intentional manipulations (such as ballot tampering, creative disinformation, and insider fraud), imbalanced redistricting (for example, local

and state gerrymandering), Electoral College issues, unlimited funding and targeted advertising (for example, Citizens United, Cambridge Analytica), delays at the polls due to malfunctioning equipment, and even the effects of environmental conditions (such as

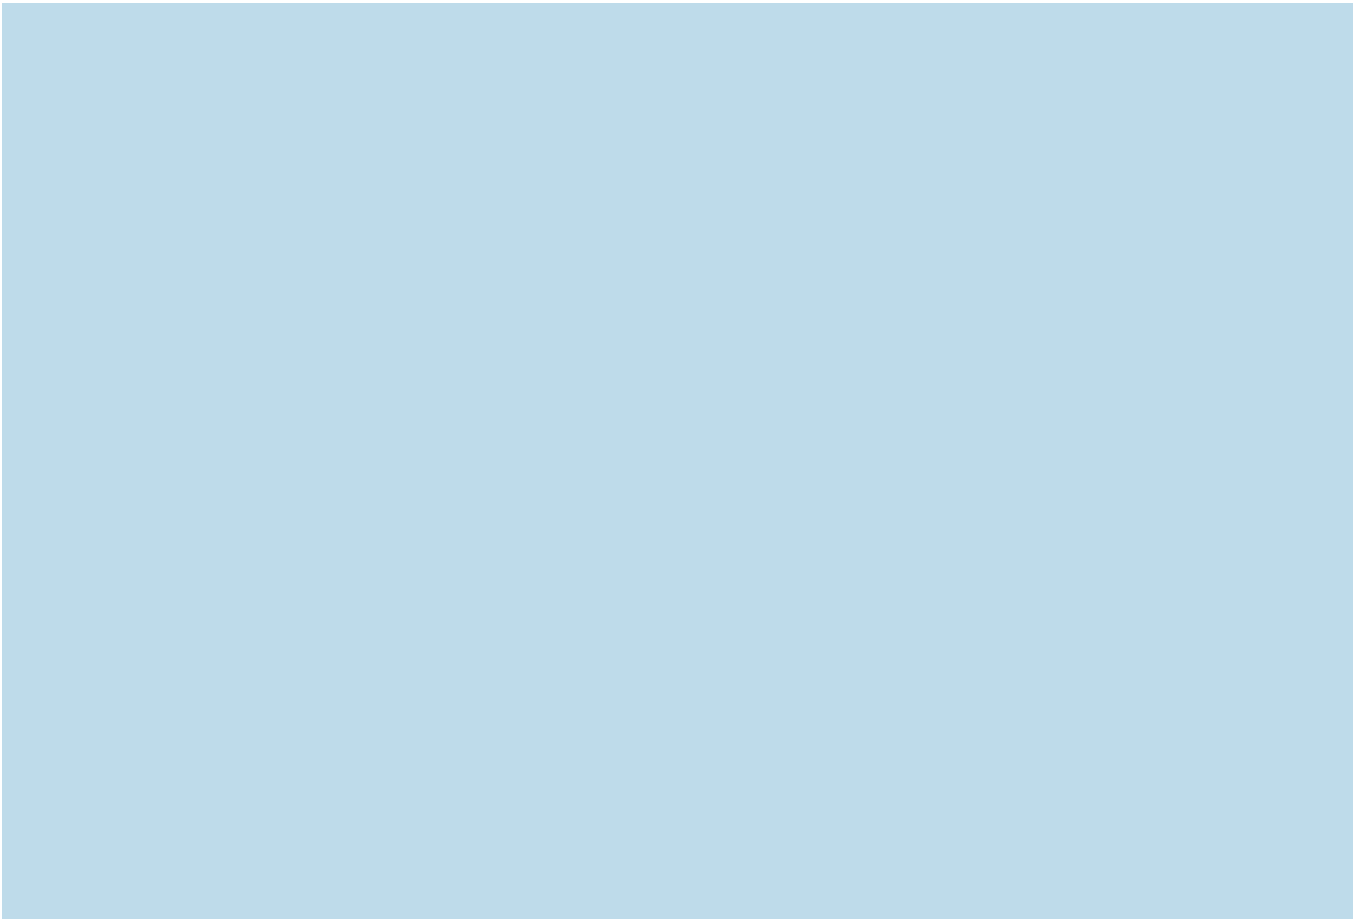
For the last four decades, computer scientists have been among the most outspoken advocates for trustworthiness, security, and reliability in election systems.

weather, power outages, and pandemics). Education is also a vital consideration, especially in the ability of voters to cope with changing conditions, new instructions, politically biased reporting, and a host of other incidental or even disingenuous factors.

As technologists, we feel it is important to assert that while recognizing the plethora of other potential risks that can affect election outcomes, a good starting place for reform must focus on the systems used to enable the casting, counting, and reporting of votes. Here, we highlight the fundamental importance of incorporating transparency and trust methods from other confidence-building processes into these specialized computational systems, including some novel approaches that have not previously been applied to elections. While we recognize much of this column is highly U.S.-centric, many of the issues raised also translate to election concerns around the globe.

Hack the Vote

For the last four decades, computer scientists have been among the most



outspoken advocates for trustworthiness, security, and reliability in election systems.” Early conferences, such as those sponsored by the Electronic Frontier Foundation and Computer Professionals for Social Responsibility, provided opportunities for discussion of election issues and potential solutions. Peter Neumann’s ACM Risks Forum has logged several hundred reports of election and voting-related human and computer errors, evident fraud, and other problems, from the 1980s to date. Additionally, 11 previous *Communications Inside Risks* columns have been devoted specifically to this topic, with several other articles pointing out relevance to election integrity.

One of the earliest researchers to raise serious concerns related to hacking was Roy Saltman. His 1988 treatise¹⁴ for the National Bureau of Standards (now NIST) was cited frequently in testimonies about the Florida 2000 Presidential election. Since 1975, Saltman’s position has been that reduced public confidence in the election process can be related to the lack of assurances that software modifications have not occurred, as well as to the vulnerabilities inherent in

commercial-off-the-shelf (COTS) products used in voting systems. These and many other issues he had flagged continue to fail to be sufficiently addressed by the Voluntary Voting System Guidelines (VVSG) established by the U.S. Election Assistance Commission (EAC). Actually, the Federal Voting System Guidelines are not mandatory, and are not applied nationwide—due to the U.S. Constitution’s preservation of states’ rights.

Some states may leave the choice of voting method and system up to their counties. Even within a single county, the voting and tabulation systems may differ—depending on whether the ballot is cast or counted at a polling location, at a voting center, at the election office, or via in-person or remote accessibility. Thus, there is a total lack of uniformity among the states and also within many states. For example, the November 2020 election in California used 21 different voting systems or versions from seven different vendors: seven products from Dominion, five from ES&S, five from Hart, and one each from Democracy Live, Interactive, Runbeck, and VSAP.¹⁵ Ensuring all of these systems

are operating in accordance with whatever guidelines the state has decided to impose for each particular election is a daunting logistical task.

After the 2000 election, the IEEE (well-respected for its 802.11 family of communications standards) launched a project intended to establish a performance standard for the evaluation of voting equipment (P1583). Unfortunately, the multiple-year effort bogged down when “attempts to insert adequate security into the standard [were] thwarted by vendors attempting to protect legacy systems, software, and proprietary trade-secret products that produce no independent method for auditing the election.”¹² These issues (and others related to Mean Time Between Failures, accuracy thresholds, and COTS) were vigorously argued between election integrity advocates and system vendors, eventually leading to an unresolvable stalemate—resulting in the non-issuance of the draft as an accepted IEEE standard.

The EAC’s VVSG 1.0 (2005), 2.0 (2015) and 2.1 (2021) each pertain to new products; no guidance is provided

as to obsolescence or with respect to vulnerabilities later discovered in systems previously certified. Older systems are grandfathered and continue to be used. Also, due to the length of time needed to update and recertify voting systems under the new VVSG, purchases of election equipment during 2021, 2022, and likely even into 2023, will have been designed under the prior guidelines, and will not adequately address more recent concerns. Communities should be advised to wait until VVSG 2.1-certified products become available.

As an example of such built-in obsolescence, in the early and mid-2000s (despite strong efforts by knowledgeable citizens to educate county and state officials about the dangers of paperless electronic voting systems), most New Jersey counties replaced their legacy lever machines with AVC Advantage DREs that were certified only to the then-obsolete FEC 1990 standards. Anomalous situations were reported early on, with observations of vote flipping, where a press for one candidate selects another instead. In the Super Tuesday Presidential Primary of February 2008, some 37 voting machines in eight New Jersey counties showed the Republican candidates to have received more votes than the number of voters who had signed in at the polls. This tabulation error was eventually attributed to a software bug. Princeton University Professor Andrew Appel demonstrated in open court that he could pick the lock on these machines, replace the ROM containing the software, and relock the door in less than seven minutes.¹ To date, most of New Jersey continues to use these vulnerable and unauditable systems.

DEFCON25, held in 2017, featured the first-ever Voting Machine Hacking Village “to highlight cyber-vulnerabilities in U.S. election infrastructure—including voting machines, voter registration databases, and election office networks.”³ Lessons learned indicated: the systems could be hacked even with limited time, information, and resources; foreign-made parts introduce supply-chain concerns; the exercise was not merely a stunt, and demonstrated that a diverse community of stakeholders should be engaged; and “affirmed what election security advocates have been arguing for years:

There is urgent need for election officials to implement measures to secure U.S. election infrastructure.”

The 2018 and 2019 DEFCON Voting Villages continued to report similar vulnerabilities, also finding that equipment was sometimes shipped with security features turned off, previously studied equipment showed new vulnerabilities, ballot-marking devices posed new systemic risks, remote attacks were possible even with air-gapped equipment, some hacks could occur in two minutes (less than the time it takes to vote), and earlier hacked equipment models were not remediated by their vendors, even though they had been informed about the known risks.

Election Forensics

Forensics is the process by which evidence is examined and described for presentation in a legal setting (for example, at a trial, hearing, or mediation), in order to allow for adjudication or resolution of a dispute. Some key aspects of forensics include these: each side is allowed independent access to the evidence; the evidence has been preserved in a traceable and pristine fashion; and the forensic review occurs using standardized tools and approaches that can be replicated.

What is immediately evident in comparing forensics to elections is this: typically, only the losing candidates are permitted to challenge the results, and often they are not given an opportunity to directly examine the evidence; some of the evidence may not

One would not inspect a few vehicles for emissions and then provide a pass for all others of the same model and type, but this is what is done with election equipment.

be properly secured (for example, it may sit at voting locations for days before being transported to warehouse facilities that may also be insecure); and while there may be methods in place for performing recounts, there is a complete lack of specified procedures and protocols for conducting a thorough forensic review of any aspects of elections. In fact, critical features of such examinations are typically prohibited by restrictive trade-secret agreements forced on the municipalities by the election equipment vendors.

There is the general assumption that election administrators are unbiased, and that the voting and vote-counting processes have been honest. Most do abide by the rules of the offices in which they serve, but notably, some do not. Still, all are given powers that, unlike in a legal trial, enable them to prevent (intentionally or by invocation of laws and procedures) a full triage of the equipment, software, ballots, and other evidence by the forensic examiners for any or all candidates in the election (whether winners or losers). This was illustrated in Georgia’s 2020 election, due to the need to prepare the voting equipment quickly to allow early voting for the state’s run-off Senate races. Such preparations would necessarily reset the equipment, which eliminates the possibility of a forensic investigation. Imagine a murder scene where no one is allowed to examine the dead body, while the murderer has access to the evidence and is allowed to eradicate all traces of it—before the forensic examiners arrive. That is the situation we have been dealing within election investigations, and it must change.

Voting System Testing

In addition to maintaining the VVSG, the EAC also arranges for certification and testing (paid for by the equipment vendors and with results shielded by trade secrecy) for a small number of sample machines of each version and style, intended to provide an assurance of compliance with the guidelines. Analogously, one would not inspect a few vehicles for emissions and then provide a pass for all others of the same model and type, but this is what is done with election equipment. Lacking individual acceptance testing for each voting system, there is no way to know

whether the procured units all actually conform to the VVSG.

The only U.S. state that ever claimed to perform such scrutiny was Georgia, which for many years had contracted with Kennesaw State University's Center for Election Systems to provide services that involved checking each voting system on procurement, before delivery to the counties. This contract was terminated by Secretary of State Brian Kemp in October 2017, following an investigation that shockingly revealed that critical vulnerabilities to the systems were known by the Center to have existed prior to the 2016 general election.⁵ This information was never properly reported, and the evidence was deliberately deleted—leaving the accuracy of the election results in doubt.

With such a diversity of complex systems, amplified by the necessity of correct recognition of hundreds of different ballot layouts with thousands of candidates, it is difficult—indeed impractical—for state and county officials to fully verify that proper functionality existed before, during, and after each election. Instead, they typically rely on the use of sample ballot sets during pre-election setup. But it has been demonstrated that sample ballots could act to circumvent or even install malware in vulnerable systems. Nor are officials typically provided with the in-depth knowledge needed to establish believability that these products are in compliance with the state's voting system standards. On the other hand, those states with a single type of voting system, or products from only one manufacturer, may produce a monoculture risk such that an attack could potentially affect their entire election's results.

In recent flurries of legislation, we are now seeing that states can and do establish their own rules for how elections are conducted. Variations may include the dates and times for voting, the manner in which ballots are laid out (some states use the party of the current governor to determine which candidates appear first—above or to the left of others), the types of voting equipment that will be used, absentee and mail-in ballot rules, and so on. One complex problem is that there has been no reconciliation of the various state laws pertaining to what ballot mark (X, dot, check, circle, arrow, and so forth) constitutes a legal

One of the myths of having trustworthy computer systems for elections is that everything depends only on the quality of the software.

vote, so many voting systems fail to recognize legally valid ballot choices. This can be remediated with hand counting of the full set of ballots (as long as the people doing the counting have been properly instructed), but this is customarily not performed. Partial audits may not pick up enough of these misreported votes to make a difference in the election outcome.

Audits vs. Recounts

Some states (including Florida, Georgia, and Michigan) have instituted policies and laws that actually prohibit the hand counting of paper ballots and allow only rescanning. Many other states prescribe only a partial statistical audit. The only reason that Georgia was able to conduct a 100% manual recount (in violation of its state law) for the 2020 presidential race was that wording in the state's Risk-Limiting Audit (RLA) legislation allowed for such, if the disparity between the candidates was close enough that nearly a full count would have had to be performed for results to be above the threshold for sufficient assurance of correctness. Had Georgia's law been worded differently, a tedious process of randomization for ballot selection would have had to occur, possibly not providing results in the time needed to certify the election, and with less believability than the multiple full counts that actually were performed throughout the state.

As of 2017, approximately half of the U.S. states required some form of post-election audit, typically only of 1% of the ballots cast. The RLA method has been gaining in popularity, because the number of ballots counted depends on the margin of votes between the leading

candidates. If there's a wide difference, counting can stop sooner. Actually, the formula that determines when to stop counting is fairly sophisticated,¹³ such that most election officials (beyond a rare few with deep knowledge of statistics and probability) would not be able to conduct an RLA without computer assistance. So, while the method may seem to be transparent, the calculations and their correctness are not necessarily obvious or comprehensible.

Some have suggested referring to the RLA as a Recount-Limiting Audit, since a primary intention of this method is to speed up validation of the election results by preventing full recounts. In addition, since the number of ballots to audit is determined by the initially computer-generated vote tallies (which have not yet been certified), there is a believability problem with regard to whether the RLA will be sufficient to reveal ballot tabulation anomalies.

Another issue is that since the selection of ballots to hand-count can be based on precinct groups, RLA is inherently better at detecting localized problems (such as with a particular voting machine or scanner) than it is with dispersed issues (a few votes added or subtracted here and there). Localized problems are usually more immediately evident in the vote tallies anyway. Dispersed problems are harder to detect, and are also less likely to be caught via RLAs. For example, major cities often show consistent voting patterns for particular party candidates in large numbers, year after year. A hacker might siphon off some votes from the winning candidate in these cities, adding them to the runner-up, or even to a third-party candidate, without detection. This would not affect the outcome of local races, but could be sufficient to alter the results of statewide races (such as for president, senator, and governor) without detection. Note that audit discrepancies may provide little or no clues as to how the tally anomalies occurred, so a forensic review of the voting systems should be (but typically is not) performed.

Given these flaws and concerns for audits, and including the time that it takes to conduct them, a better alternative might be to publicly count paper ballots on election night (as is done in the U.K. and Canada). These proceedings should be live-streamed and recorded for later scrutiny. One of N races

(where a single candidate is selected by the voter out of N choices) is especially easy to perform using the bin-sort method well known to computer scientists.

CISA Oversight

For the 2020 election cycle, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) took a proactive role in trying to thwart election disinformation via its [cisa.gov/rumorcontrol](https://www.cisa.gov/rumorcontrol) Website. The agency's October 20 pre-election statement asserted "We remain confident that no foreign cyber actor can change your vote, and we still believe that it would be incredibly difficult for them to change the outcome of an election at the national level." Noted is that this seemingly positive press release does not encompass the full gamut of election shenanigans, including those that are capable of being performed by U.S. citizens.

Following the 2020 general election, on November 12th, CISA issued a 10-person joint statement that included leaders from the Election Assistance Commission, the National Association of Secretaries of State, the National Association of State Election Directors, and numerous representatives from election service

companies (Unisyn, Hart InterCivic, ES&S, ERIC, and DemocracyWorks). The statement included the assertion: *There is no evidence that any voting system deleted or lost votes, changed votes, or was in any way compromised.*⁶

Four days later, Professor Matt Blaze issued a letter signed by 59 computer scientists^a (many of whom are well known to the election integrity community), which echoed the CISA joint statement asserting that, "To our collective knowledge, no credible evidence has been put forth that supports a conclusion that the 2020 election outcome in any state has been altered through technical compromise."

What is curious about both the CISA and joint computer scientists' statements is that they were premature. The secretaries of states would not be certifying their election results for many weeks. Most of the nation's voting systems were still on lockdown for pending

election challenges and recounts, and could not be physically examined. Some vote counting was still ongoing. Absolutely none of the 10 people on the CISA Joint Statement, or the 59 people on the computer scientists' letter, had performed any post-election forensics or triage that would support these conclusions. In fact, many of the scientists (including Andrew Appel, Richard DeMillo, Alex Halderman, Harri Hursti, and Philip Stark) had, a few months prior to the November election, provided testimony in a Georgia U.S. District Court matter on behalf of plaintiffs against Brad Raffensperger, et al., objecting to the use of the electronic Ballot Marking Devices on the grounds that they might not be sufficiently accurate and that the scanners may incorrectly report results.

The closing pages of U.S. District Judge Amy Totenberg's ruling in that case^b included the following pertinent statements: "The stealth vote alteration or operational interference risks

a Scientists say no credible evidence of computer fraud in the 2020 election outcome, but policymakers must work with experts to improve confidence (Nov. 16, 2020); <https://bit.ly/3mZeG2u>

b Totenberg, Honorable Amy, Opinion and Order in *Curling, et al. v. Raffensperger, et al.*, Civil Action No. 1:17-cf-2989-AT (Oct. 11, 2020); <https://bit.ly/2QxtysN>

Transforming research and practice

Introducing three new journals from Cambridge University Press:

- Exploring the impact of AI and Data Science on different domains;
- Bridging the gap between research and practice;
- Peer-reviewed and overseen by international editorial teams;
- Open Access with support for unfunded authors.



Data-Centric Engineering explores the transformative potential of AI and Data Science on all engineering disciplines, publishing original research and translational papers. Supported by the Lloyd's Register Foundation.



Data & Policy promotes deeper understanding of policy-data interactions by publishing research, translation and commentary. In association with the Data for Policy Conference and supported by the Alan Turing Institute, UCL and ONS.



Environmental Data Science is dedicated to the use of AI and Data Science to deepen our understanding of environmental processes, including climate change. Publishing applications, methods and data papers.

Three new Open Access journals

- ▶ Enhance the impact of your work through Open Access publication
- ▶ Authors supported by Read & Publish agreements and partnerships
- ▶ Publish with a global, not-for-profit university press

Learn more: cambridge.org/datajournals



CAMBRIDGE
UNIVERSITY PRESS

posed by malware that can be effectively invisible to detection, whether intentionally seeded or not, are high once implanted, if equipment and software systems are not properly protected, implemented, and audited. ... Given the masking nature of malware and the current systems described here, if the State and Dominion simply stand by and say ‘we have never seen it’ the future does not bode well.”

Not heeding this warning, the CISA and computer scientists’ statements effectively said “we have never seen it” without conducting any actual investigation to determine if a cybersecurity breach affecting the election results had happened or not, in Georgia or anywhere else in the country. This is not reassuring to the voters. Far better to have said “we don’t know” or “we are investigating” than to prematurely issue statements intended to convey that everything was copacetic.

Ironically, we should note that after the election concluded, it was later learned that CISA and other security-related government agencies had been breached with the SolarWinds attack, suffering an as yet fully unknown extent of damages and loss of information to foreign agents. This unprecedented hack remained undetected for about nine months, spanning the time of the 2020 election.

If those charged with protecting elections cannot defend their own assets, their claims of “no compromise” of the election systems must be considered with a skeptical eye. On the other hand, research-based claims that voting systems and election results have been or could be compromised, may begin to be suppressed. This is now being put to the test by the defamation lawsuits seeking \$4.3B in restitution by the voting system vendors Dominion and Smartmatic. Whether free speech and freedom of the press will prevail with regard to exposure of election security concerns is left to be seen.

Paper Ballots, Anonymity, and Cryptography

While much of the controversy in the 2020 presidential election focused on absentee vs. mail-in ballots, structurally these paper ballots are the same, and may even be identical to the scanned paper ballots used at the polling loca-

Elections must be constructed and conducted such that everyone can, with extremely high confidence, rationally believe the results reflect the will of the voters.

tions. Their differentiation is legal in nature—an absentee ballot is issued to a registered voter who has requested one because they will not be able to visit a polling place on Election Day (or during early voting) for a legitimate reason. A mail-in ballot is one that is issued by the locality or state without having been specifically requested by the voter. In 2000, Oregon became the first state to eliminate precinct voting, replacing it with all mail-in ballots. More recently, and especially due to postal delays, local drop-boxes have become a convenient way of depositing absentee or mail-in ballots, but some states are now trying to roll back their use.

As for the paper ballots themselves, very little has evolved to make them more trustworthy over the past quarter-century. By comparison, paper money and checks have changed dramatically during this time. One can obtain 10,000 checks that have 20 security features (including a foil hologram, prismatic multicolor background, micro-printing, heat-sensitive icons, watermarks, invisible fluorescent fibers, red/blue visible fibers, toner adhesion, and chemical sensitivity) for a little less than 13 cents each. Still, the printing companies and voting system vendors have yet to implement any of these types of document authentication methods for paper ballots. Certainly the paper stock that contains the authentication would need to be strictly inventoried and controlled, to prevent spoofed ballots from being subversively created. However, there has been no demand for such ballot fea-

tures, not even by the individuals and groups who have been complaining about the possibility of dead people voting (which turns out to be exceedingly rare).

With the greater availability of non-precinct voting comes the increased risk of vote selling or coercion. But it may also be the case that as people have become more willing to share personal details via social media, and as photographing and publishing things (such as your filled-out ballot) has also become commonplace, there may be less concern about the secrecy or privacy of one’s voting choices. A major finding of Rebecca Mercuri’s Ph.D. dissertation (partly synopsized in Mercuri¹⁰) established that “the need for anonymity precludes the use of transaction logging for providing access assurances” in direct-recording electronic voting systems. In other words, it is not possible in fully anonymous voting to ensure individuals (such as voters or precinct workers) have not tampered with the voting system during the election in order to alter the vote totals, without the availability of voter-verified paper ballots to use in performing a cross-check. This is not just a theoretical speculation but rather is based on NP-completeness proofs.

With the elimination of full anonymity (such as happens in a stock shareholder election—one casts ballots that are tracked to the owner of particular shares)—the voter can be contacted to verify that they did cast their votes as recorded. Actually, the U.K. does use ballot numbers to track votes, and under very constrained circumstances can require a voter to later validate that their ballot was cast as intended.

Various cryptographers have devised methods for generating encrypted ballots. Some of these schemes (notably Chaum’s⁴) enable the voter to decode their votes or track them in order to confirm that not only was the ballot received for counting, but also that their vote choices have been entered into the tallies correctly. Unfortunately, as with Risk-Limiting Audits, these crypto algorithms are too complex for most people to understand, which limits the believability of correctness of the implementing software (which may also be insecure) to an elite intelligent few. Some have considered

the use of blockchain as a voting method, but this would require extensive reliance on trustees and software to properly maintain the cryptographic records. Quantum voting has also been suggested, but this technology is not yet mature—and likely to be overkill. These techniques could someday show promise for elections but are not yet well-enough understood by the general public to ensure believable results.

Opensource Software and Open-Architecture Hardware

One of the myths of having trustworthy computer systems for elections—and indeed for trustworthy applications in general—is that everything depends only on the quality of the software. This myth is finally being debunked by exploits that take direct advantage of already existing hardware risks. Examples include speculative execution vulnerabilities introduced by the Spectre exploits,⁷ and the Thunderclap vulnerabilities⁸ whereby a USB-C stick could take over most systems with or without IOMMUs. Neither of those cases can be resolved only in software. In addition, hardware supply-chain compromises have long been a concern, most recently exemplified by the hacked water-treatment facility in Florida.

Ideally, total-system hardware-software architectures should be deployed to eliminate vulnerabilities as well as simplify the programming process that is presently riddled with potential flaws. An example of such an architecture is provided by the emerging Capability Hardware Enhanced (CHERI) RISC Instructions hardware instruction-set architecture and its operating systems.^c Existing versions that could be used to develop voting systems include the opensourced CHERI-RISC-V hardware ISA and Arm's prototype Morello board of the CHERI spec integrated into their Version 8 hardware, each with appropriate opensourced software.

A combination of CHERI's least-privilege access controls, fine-grained and course-grained compartmentalization, with highly principled design,

has the potential of dramatically increasing the trustworthiness of the computer aspects of future elections. In addition, formal proofs that critical security properties are satisfied by the CHERI instruction-set architecture could further increase the believability of that claim.

Still, the admonition by Ken Thompson in his classic 1984 ACM A.M. Turing Award Lecture—Reflections on Trusting Trust—holds true today: “You can’t trust code that you did not totally create yourself. (Especially code from companies that employ people like me.)” Indeed, many issues continue to persist (including some noted in this column) that can impact the believability of election results from voting systems whose vendors deliberately prevent open software and hardware reviews.

But neither COTS nor opensource hardware and software are inherently immune to exploits. Consequently, election integrity must also involve protections against insider misuse, including accountability in reliably recording all changes in hardware, software, data, and system configurations. This is a non-trivial problem that posting source code on GitHub cannot solve entirely. The very real possibility of exposure (via an open review) of a major flaw or breach shortly before, during, or after an election, stands to wreak havoc on the believability of the vote totals, but this risk is not reduced by refusing to make the code open to review.

Conclusion

Elections must be constructed and conducted such that everyone (all of the winning and losing candidates, as well as those who have supported them) can, with extremely high confidence, rationally believe the results reflect the will of the voters. Technology is only one part of the end-to-end voting process. Repeatability and transparency can provide critical assurances that enhance trust—but voter-verified paper ballots and open software and hardware do not ensure correctness if scrutiny is thwarted or lacking. Risk-Limiting Audits and cryptographic methods, while perhaps mathematically sound, are not believable if the general public lacks the intellectual sophistication to understand how they work, or the capa-

bility to independently demonstrate correctness proofs for themselves. In short, the risks of seeking election believability are that the cure can be worse than the disease. As computer scientists, we must bear responsibility for warning about election vulnerabilities and proposing solutions, while also being careful not to make unfounded statements of assurance or promote voting and auditing methods that are incomprehensible by the preponderance of the electorate. **□**

References

1. Appel, A. et al. The New Jersey Voting-machine Lawsuit and the AVC Advantage DRE Voting Machine, EVT/WOTE'09 (Aug. 2009); <https://bit.ly/32tWIRS>
2. Bellovin, S.M. and Neumann, P.G. The big picture: A systems-oriented view of trustworthiness. *Commun. ACM* 61, 11 (Nov. 2018); <https://bit.ly/3mYVpOR>
3. Blaze, M. et al. Voting Machine Hacking Village Report on Cyber Vulnerabilities in U.S. Election Equipment, Databases, and Infrastructure, (Sept. 2017); <https://bit.ly/3suT6UH>
4. Chaum, D. Secret-ballot receipts: True voter-verifiable elections. *IEEE Security & Privacy*. (Jan. 2004); <https://bit.ly/3eh7nzm>
5. Gordon, G., Condon, C. and Dunlap, S. Georgia election officials knew system had 'critical vulnerabilities' before 2016 vote. McClatchy DC Associated Press (Aug. 6, 2018); <https://bit.ly/3mZGQKN>
6. Joint Statement from Elections Infrastructure Government Coordinating Council and The Election Infrastructure Sector Coordinating Executive Committees (Nov. 12, 2020); <https://bit.ly/3v3I9Lo>
7. Kocher, P. et al. Spectre Attacks: Exploiting Speculative Execution. In *Proceedings of the 2019 IEEE Symposium on Security and Privacy* (May 2019); <https://bit.ly/32qH2o>
8. Markkettos, A.T. et al. Thunderclap: Exploring Vulnerabilities in Operating-System (IOMMU) Protection via DMA from Untrustedworthy Peripherals, Network and Distributed System Security Symposium. NDSS 2019 (Feb. 24–27, 2019).
9. Mercuri, R. A better ballot box. *IEEE Spectrum* (Oct. 1, 2002); <https://bit.ly/3xao8oz>
10. Mercuri, R. Uncommon criteria. *Commun. ACM* 45, 1 (Jan. 2002); <https://bit.ly/3gjSjtP>
11. Monmouth University Poll. More Americans happy about Trump loss than Biden win. Press Release (Nov. 18, 2020); <https://bit.ly/3swuuej>
12. Rein, L. The IEEE P1583 Voting Machine Standard. *IEEE Internet Computing* 8, 1 (Jan.–Feb. 2004); <https://bit.ly/3edEwM9>
13. Risk-Limiting Audits Working Group, Risk-Limiting Post-Election Audits: Why and How (Oct. 2012); <https://bit.ly/3srOqRe>
14. Saltman, R. Accuracy, Integrity, and Security in Computerized Vote-Tallying. NBS/NIST (Aug. 1, 1988); <https://bit.ly/3sx8y2z>
15. Weber, S.N. California Secretary of State, Voting Technologies In Use By County, as of 15 October 2020; <https://bit.ly/32tPfg0>

Rebecca T. Mercuri (mercuri@acm.org) has been researching, writing, and testifying about election system integrity since the late 1980s. She also provides digital investigations and expert testimony on a variety of criminal and civil matters through her company, Notable Software, Inc.

Peter G. Neumann (neumann@csl.sri.com) is Chief Scientist of the SRI International Computer Science Lab, and has moderated the ACM Risks Forum since its beginning in 1985.

The authors are grateful to Steven M. Bellovin, E. John Sebes, Vanessa Teague, and Brent Turner for their constructive comments during the development of this column.

Copyright held by authors.

^c See the CHERI website for published papers and reports, including the hardware instruction-set architecture report, the compartmentalization paper, the Thunderclap paper, and so on: <https://bit.ly/2RN5Fb>