

THE FRIDAY COVER

How to Hack an Election in 7 Minutes

With Russia already meddling in 2016, a ragtag group of obsessive tech experts is warning that stealing the ultimate prize—victory on Nov. 8—would be child’s play.

By BEN WOFFORD | August 05, 2016



Andrew Appel and a Sequoia AVC Advantage voting machine. | Alex Halderman

When Princeton professor Andrew Appel decided to hack into a voting machine, he didn’t try to mimic the Russian attackers who hacked into the Democratic National Committee’s database last month. He didn’t write malicious code, or linger near a polling place where the machines can go unguarded for days.

Instead, he bought one online.

With a few cursory clicks of a mouse, Appel parted with \$82 and became the owner of an ungainly metallic giant called the Sequoia AVC Advantage, one of the oldest and vulnerable, electronic voting machines in the United States (among other places it’s deployed in Louisiana, New Jersey,

Virginia and Pennsylvania). No sooner did a team of bewildered deliverymen roll the 250-pound device into a conference room near Appel's cramped, third-floor office than the professor set to work. He summoned a graduate student named Alex Halderman, who could pick the machine's lock in seven seconds. Clutching a screwdriver, he deftly wedged out the four ROM chips—they weren't soldered into the circuit board, as sense might dictate—making it simple to replace them with one of his own: A version of modified firmware that could throw off the machine's results, subtly altering the tally of votes, never to betray a hint to the voter. The attack was concluded in minutes. To mark the achievement, his student snapped a photo of Appel—oblong features, messy black locks and a salt-and-pepper beard—grinning for the camera, fists still on the circuit board, as if to look directly into the eyes of the American taxpayer: *Don't look at me—you're the one who paid for this thing.*

Appel's mischief might be called an occupational asset: He is part of a diligent corps of so-called cyber-academics—professors who have spent the past decade serving their country by relentlessly hacking it. Electronic voting machines—particularly a design called Direct Recording Electronic, or DRE's—took off in 2002, in the wake of *Bush v. Gore*. For the ensuing 15 years, Appel and his colleagues have deployed every manner of stunt to convince the public that the system is pervasively unsecure and vulnerable.

Beginning in the late '90s, Appel and his colleague, Ed Felten, a pioneer in computer engineering now serving in the White House Office of Science and Technology Policy, marshaled their Princeton students together at the Center for Information Technology Policy (where Felten is still director). There, they relentlessly hacked one voting machine after another, transforming the center into a kind of Hall of Fame for tech mediocrity: reprogramming one popular machine to play Pac-Man; infecting popular models with self-duplicating malware; discovering keys to voting machine locks that could be ordered on eBay. Eventually, the work of the professors and Ph.D. students grew into a singular conviction: It was only a matter of time, they feared, before a national election—an irresistible target—would invite an attempt at a coordinated cyberattack.

The revelation this month that a cyberattack on the DNC is the handiwork of Russian state security personnel has set off alarm bells across the country: Some officials have suggested that 2016 could see more serious efforts to interfere directly with the American election. The DNC hack, in a way, has compelled the public to ask the precise question the Princeton group hoped they'd have asked earlier, back when they were turning voting machines into arcade games: *If motivated programmers could pull a stunt like this, couldn't they tinker with the results in November through the machines we use to vote?*

This week, the notion has been transformed from an implausible plotline in a Philip K. Dick novel into a deadly serious threat, outlined in detail by a raft of government security officials. "This isn't a

crazy hypothetical anymore,” says Dan Wallach, one of the Felten-Appel alums and now a computer science professor at Rice. “Once you bring nation states’ cyber activity into the game?” He snorts with pity. “These machines, they barely work in a *friendly* environment.”

ADVERTISING

inRead invented by Teads

The powers that be seem duly convinced. Homeland Security Secretary Jeh Johnson recently conceded the “longer-term investments we need to make in the cybersecurity of our election process.” A statement by 31 security luminaries at the Aspen Institute issued a public statement: “Our electoral process could be a target for reckless foreign governments and terrorist groups.” Declared Wired: “America’s Electronic Voting Machines Are Scarily Easy Targets.”

For the Princeton group, it’s precisely the alarm it has been trying to sound for most of the new millennium. “Look, we could see 15 years ago that this would be perfectly possible,” Appel tells me, speaking in subdued, clipped tones. “It’s well within the capabilities of a country as sophisticated as

Russia.” He pauses for a moment, as if to consider this. “Actually, it’s well within the capabilities of much less well-funded and sophisticated attackers.”

In the uproar over the DNC, observers have been quick to point out the obvious: There is no singular national body that regulates the security or even execution of what happens on Election Day, and there never has been. It’s a process regulated state by state. Technical standards for voting are devised by the National Institute of Standards and Technology and the Election Assistance Commission—which was formed after the disputed 2000 presidential election that hinged on faulty ballots—but the guidelines are voluntary. (For three years the EAC limped on without confirmed commissioners—an EAC commissioner stepped down in 2005, calling its work a “charade”). Policy on voting is decided by each state and, in some cases, each county—a system illustrated vividly by the trench warfare of voter ID laws that pockmark the country. In total, more than 8,000 jurisdictions of varying size and authority administer the country’s elections, almost entirely at the hands of an army of middle-age volunteers. Some would say such a system cries out for security standards.

If such standards come to fruition, it will be the Princeton group—the young Ph.D.’s who have since moved on to appointments and professorships around the country—and their contemporaries in the computer science world who suddenly matter.

The Princeton group has a simple message: That the machines that Americans use at the polls are less secure than the iPhones they use to navigate their way there. They’ve seen the skeletons of code inside electronic voting’s digital closet, and they’ve mastered the equipment’s vulnerabilities perhaps better than anyone (a contention the voting machine companies contest, of course). They insist the elections could be vulnerable at myriad strike points, among them the software that aggregates the precinct vote totals, and the voter registration rolls that are increasingly digitized. But the threat, the cyber experts say, starts with the machines that tally the votes and crucially keep a record of them—or, in some cases, don't.

Since their peak around 2007, voting districts have begun to rely less on the digital voting machines—a step in the right direction, as states bolt for the door on what the programmers describe as a bungled, \$4 billion experiment. Instead, rushing to install paper backups, sell off the machines and replace them with optical scanners—in some cases, ban them permanently for posterity. But the big picture, like everything in this insular world, is complicated. As the number of machines dwindle—occasioned by aging equipment, vintage-era software that now lacks tech support, years without new study by the computer scientists, and a public sense that the risk has passed—the opportunities for interference may temporarily spike. Hundreds of digital-only precincts still remain, a significant portion of them in swing states that will decide the presidency in November. And, as the Princeton group warns, they become less secure with each passing year.

In American politics, an onlooker might observe that hacking an election has been less of a threat than a tradition. Ballot stuffing famously plagued statewide and some federal elections well into the 20th century. Huey Long was famously caught rigging the vote in 1932. Sixteen years later, 1948 saw the infamous “Lyndon Landslide,” in which Johnson mysteriously overcame a 20,000 vote deficit in his first Senate race, a miracle that Robert Caro reports was the almost certain result of vote rigging. But even an unrigged election can go haywire, as the nation learned in horror during the Florida recount in 2000, when a mind-numbingly manual process of counting the ballots left a mystery as to which boxes voters had punched—giving the nation the “hanging chad,” and weeks of uncertainty about who won the presidency.

In some ways, the country’s response was suggestive of the real crime committed in Florida: Not inaccuracy, but anxiety. Congress’s solution was to pass the Help America Vote Act in 2002, a nearly \$4 billion federal fund meant to incentivize states to upgrade their voting machines. It worked. All 50 states took the money. Requirements included upgrading voter registration methods and making polls disability-friendly, but Section 102 provided funds specifically allocated for replacing outdated voting machines; almost universally, “upgrade” meant a new, computerized touch-screen voting machine. By 2006, states had spent nearly \$250 million on new machines with Section 102 funds. In Pennsylvania, the funds purchased 20,597 new machines—around 19,900 of which were digital touchscreens. Some, like the Diebold TSX, Advanced WINvote, the ES&S iVotronic, and a variant of Appel’s AVC Advantage—the Sequoia Edge—would be the same models to come under scrutiny by cybersecurity experts and academics. Thousands of touchscreen DREs were similarly sold in state contracts. Between Election Day 2000 and the HAVA cutoff in 2006, the stock prices of the major companies soared.

The appeal of such machines seemed plain: Voting was crisp, instantaneous, logged digitally. To state officials—and, at first, voters—the free federal money seemed like a bargain. To computer scientists, it seemed like a disaster waiting to happen. Wallach remembers when he testified before the Houston City Council, urging members not to adopt the machines. “My testimony was: ‘Wow, these are a bad idea. They’re just computers, and we know how to tamper with computers. That’s what we do,’” Wallach recalls. “The county clerk, who has since retired, essentially said, ‘You don’t know anything about what you’re talking about. These machines are great!’ And then they bought them.”

Almost from the day they were taken out of the box, the touch-screen machines demonstrated problems (the same companies had a much better track record with Optical Scan machines). During the primaries in Florida in 2002, some machines in Miami-Dade malfunctioned and failed to turn on, resulting in hourslong lines that locked out untold numbers of voters—including then-

gubernatorial candidate Janet Reno. That year, faulty software (and an administrator oversight) on Sequoia models led to a fourth of votes initially omitted during early voting in Albuquerque’s Bernalillo County. In Fairfax County, Virginia, an investigation into a 2003 school board race found that a vote was subtracted for every 100 votes cast for one of the candidates on 10 machines. With margin sizes small enough to be noticed, local elections were vaulted into the forefront of these debates; Appel later found himself issuing expert testimony for a tiny election for the Democratic Executive Committee in Cumberland County, New Jersey, where a candidate lost by 24 votes. The margin was small enough that the losers sued, and called 28 voters as witnesses—who each swore they voted for them. The machine in use was a Sequoia AVC Advantage.

“Wow, these are a bad idea. They’re just computers, and we know how to tamper with computers.”

Cybersecurity researchers flocked to study the machines, but they say they were faced with an uncompromising adversary: the voting machine companies, which viewed the code of the machines as intellectual property. Until 2009, two companies, Diebold and ES&S, controlled the lion’s share of the voting machine market. The accreditation process is equally narrow: Since 1990, a voluntary federal accreditation process has certified voting technology, a system that has come under fire for its lack of transparency. The laboratories (“Independent Testing Authorities”) which conduct the certification reviews are typically paid by the manufacturers, and are usually required to sign nondisclosure agreements. In 2008, five labs were accredited; one was suspended that year for poor lab procedures, and another temporarily suspended for insufficient quality control.

State authorities can typically request these lab reports, as Kathy Rogers of ES&S reminded me in an email. (“For security reasons we did not make that code widely available to just *anyone and everyone* who simply wanted a copy for their own purposes. We truly have nothing to hide.”) But Appel, the Princeton group and others in cybersecurity have insisted that such measures—which they deem “security through obscurity”—pale to the types of rigorous testing that would result from releasing the code to the public or academics. One of the companies, Sequoia, later acquired by Dominion, once threatened Princeton’s Felten and Appel with legal action if they attempted to examine one of their models.

Election officials have sometimes complained that the lab reports they do receive lack vital detail, and information from the labs, bound by the NDAs, can be unforthcoming. In 2004, when the California Secretary of State Kevin Shelley—in charge of overseeing the state’s elections—asked one of the five laboratories for more information on the testing of machines, he was stonewalled, and told by a researcher, “We don’t discuss our voting machine work.” Because of a flood of machines

introduced to the market after HAVA, the 2002 accreditation standards are the ones that matter—the same process that approved touch-screen Diebold machines that had supervisor passcodes of “1111” in order to access the voting system. Shelley later banned Diebold TSX machines, calling Diebold’s conduct “deceitful.”

In 2003, an employee at Diebold mistakenly left 40,000 files containing code for the Diebold AccuVote TS, one of the most widely used machines on the market, on a publically viewable website. The computer scientists moved in, and one of the early and formative papers was published on the subject, co-authored by Wallach and led by Johns Hopkins’ Avi Rubin. Its findings were devastating: The machine’s smartcards could be jerry-rigged to vote more than once; poor cryptography left the voting records file easy to manipulate; and poor safeguards meant that a “malevolent developer”—an employee inside the company, perhaps—could reorder the ballot definition files, changing which candidates received votes. The encryption key, F2654hD4, could be found in the code essentially in plain view; all Diebold machines responded to it. (Rubin later remarked that he would flunk any undergrad who wrote such poor code.) “We read the code, and found really, really bad problems,” Wallach tells me, sitting at his Houston dining table. He catches himself. “Actually, let me change that,” he says. “We found *unacceptable* problems.” Diebold dismissed the report, responding that the code was obsolete, and the study’s findings thusly moot. But the 2003 report catalyzed a small movement: In CompSci departments across the country, vote hacking became a small, insular civic code of honor. Felten’s group at Princeton led the pack, producing some of the most important papers throughout the 2000s.

“We read the code, and found really, really bad problems,” Wallach tells me, sitting in his Houston dining table. He catches himself. “Actually, let me change that,” he says. “We found *unacceptable* problems.”

By the following year, professors in and around the Princeton group began the work of unwinding what they viewed as a 50-state debacle. Felten and Appel shared a taste for gallows humor and a flair for promotion. Felten took to blogging, and started a tradition: Each election, he snapped a photo standing alone with unguarded voting machines days before the election. In another study, the Sequoia AVC Edge was infected with malware that allowed it to do nothing but play Pac-Man; the students pulled off the feat without breaking the machines “tamper-proof” seals, and decorated the machine with Pac-Man logos. The team tore through topics including source code review of the larger Diebold voting system; advising election officials on security measures without new hardware; and designing malware for the Sequoia AVC Advantage that Appel had purchased, using a technique called a Return-Oriented Program. In less than a minute, they infected a Diebold

machine with self-duplicating code, spreading from machine to machine through an administrator card, and programmed it to swing an election for Benedict Arnold over George Washington.

The latter hack was the result of a curious and enigmatic email, when Felten received a message from an anonymous source, presumably with ties to the voting machine industry. Diebold's response to the Rubin and Wallach study was brittle and evasive; the source wanted to give Felten a Diebold TS machine—the same one whose code had leaked in the study. Studying the machine itself would offer an unmissable opportunity—Felten put his grad students, Feldman and Halderman, then 25 years old, in charge of the effort. One night in April 2006, Halderman drove to New York City, and double-parked his car, lights blinking, in front of a hotel just a few blocks from Times Square. Halderman jogged into an alleyway, where his source stood patiently, dressed in a charcoal colored trench coat and wielding a black canvas bag. After a few terse formalities, he handed Halderman the bag with the machine inside. Halderman never saw the man again. (“There's a lot of cloak and dagger in election security,” Halderman would tell me later.)

Throughout the summer of 2006, Feldman and Halderman set themselves to work in the basement of an academic building. Fearing retribution or a lawsuit, they didn't tell their colleagues in the department of their project. From noon until midnight, the two students met on the humid Princeton quad, and decamped to a claustrophobic, eggshell anteroom—enough space for a small table and two uncomfortable foldout chairs—and pored through reams of code and programming under the fluorescent lighting of the windowless room. At the center of the table was the subject of years of mystery: The squat, beige monitor of the Diebold TS. The authors would later describe the project as the first rigorous analysis of a physical touch-screen DRE—supposedly the kind of testing it would have received in one of the accredited labs.

When they were finished, they had another paper's worth of findings, and the most comprehensive understanding of how Diebold's machines worked. “We found the machine did not have any security mechanisms beyond what you'd find on a typical home PC,” Halderman told me. “It was very easy to hack.” Studying with Felten, Halderman had learned a key phrase—“Defense in Depth,” meant to describe a system with various rings of security. Halderman joked that the model should more aptly be called “Vulnerability in Depth,” so numerous were the entry points they discovered. Later, they found the key that opened the Diebold AccuVote TS was a standard corporate model, reproduced for minibars and other locks, available online. When their report revealed this detail, a commonplace reader found a picture of the key, filed down a blank from ACE Hardware and sent a copy to Feldman and Halderman as a souvenir (who then tested the key—it worked). That year, 10 percent of registered voters alone used the AccuVote TS to vote.

None of these breakthroughs were lost on states that had bought the machines, officials who were keeping an eye on academic reports. Felten would later write that the vulnerabilities in the Diebold

machine they tested likely could not be rectified without fully redesigning the machine; but the solution for state officials was simple. If they could include a paper trail—a voter-verified paper receipt that printed alongside the digital vote—the electronic tally could, in theory, be cross-tested for accuracy. In December 2003, Nevada became the first state to mandate that voter verified printouts be used with digital touch screens. A wave of states followed.

But the tipping point came in 2006, when a major congressional race between Vern Buchanan and Christine Jennings in Florida's 13th District imploded over the vote counts in Sarasota County—where 18,000 votes from paperless machines essentially went missing (technically deemed an “undervote”) in a race decided by less than 400 votes. Felten drew an immediate connection to the primary suspect: The ES&S iVotronic machine, one of the many ordered in Pennsylvania after they deployed their HAVA funds. Shortly after the debacle, Governor Charlie Crist announced a deadline for paper backups in every county in Florida that year; Maryland Governor Bob Erlich urged his state's voters to cast an absentee ballot rather than put their hands on a digital touch screen—practically an unprecedented measure. By 2007, the touch screens were so unpopular that two senators, Bill Nelson of Florida and Sheldon Whitehouse of Rhode Island, had introduced legislation banning digital touch screens in time for the 2012 election.

Precincts today that vote with an optical scan machine—another form of DRE that reads a bubble tally on a large card—tend not to have this problem; simply by filling it out, you've generated the receipt yourself. But that doesn't mean the results can't still be tampered with, and Felten's students began writing papers that advised election officials on defending their auditing procedures from attempted manipulation.

Each state bears the scars of its own story with digital touch screens—a parabola of havoc and mismanagement that has been the 15-year nightmare of state and local officials. The touch screens peaked in 2006, touching nearly 40 percent of registered voters; in 2016, most voters will use some combination of paper, optical scan or paper backup. In 2013, Maryland sped up its wind-down process, pushing through a transition to optical scans for use in the 2016 election. So did Virginia, which has rushed to phase out as many as possible in time for 2016—and later passed legislation to ban them permanently by 2020, just for good measure.

The Virginia ban was the quixotic crusade of one computer science expert in the private sector, Jeremy Epstein. In 2002, Epstein walked into the elections office in Fairfax, Virginia, to complain about the poor design of the touch screens—a WINVote model—and walked out with a mission to get them barred from the state. The machines were connected to Wi-Fi—vulnerable to “anyone who wanted to could hack them from the comfort of their car out in the parking lot,” Epstein told me. An investigation later revealed that the WINVote's encryption key was “abcde.” The machines were certified in 2003, running on a version of Windows from 2002, and hadn't received an update

since 2005.

Thirteen years later, Virginia announced its ban. “If these machines and elections weren’t hacked,” Epstein later told me, a credo he’s said for years, “it was only because no one tried.”

In 2001, the notion of foreign vote hacking felt like a far-fetched warning from a far-off time—it would be years, for instance, before North Korean agents would hack a company like Sony, or the Chinese would break into the federal government's personnel files. Citizen activists who had exposed the Diebold code leak and joined the counterreformation for paper ballots were concerned, but primarily about domestic hacking. Liberals tended to see the corporate voting machine companies as a threat to fair elections. Conservatives tended to see the incompetence of poorly designed machines as a threat to normalcy.

Today, Halderman reminds me, "the notion that a foreign state might try to interfere in American politics via some kind of cyber-attack is not far-fetched anymore.”

The Princeton group has no shortage of things that keep them up at night. Among possible targets, foreign hackers could attack the state and county computers that aggregate the precinct totals on election night—machines that are technically supposed to remain non-networked, but that Appel thinks are likely connected to the Internet, even accidentally, from time to time. They could attack digitized voter registration databases—an increasingly utilized tool, especially in Ohio, where their problems are mounting—erasing voters’ names from the polls (a measure that would either cause voters to walk away, or overload the provisional ballot system). They could infect software at the point of development, writing malicious ballot definition files that companies distribute, or do the same on a software patch. They could FedEx false software to a county clerk’s office and, with the right letterhead and convincing cover letter, get it installed. If a county clerk has the wrong laptop connected to the Internet at the wrong time, that could be a wide enough entry window for an attack.

“No county clerk anywhere in the United States has the ability to defend themselves against advanced persistent threats,” Wallach tells me, using the parlance of industry for highly motivated hackers who “lay low and stick around for a while.” Wallach painted an unseemly picture, in which a seasoned cyber warrior overseas squared off against a septuagenarian volunteer. “In the same way,” continues Wallach, “you would not expect your local police department to be able to repel a foreign military power.”

“ No county clerk anywhere in the United States has the ability to defend themselves against advanced persistent threats.”

In the academic research, hacks of the machines are far more pervasive; digitized voting registrations or tabulation software are not 10 years old and running on Windows 2000, unlike the machines. Still, they present risks of their own. “There are still plenty of computers involved” even without digital touch screens, says Appel. “Even with optical scan voting, it’s not just the voting machines themselves—it’s the desktop and laptop computers that election officials use to prepare the ballots, prepare the electronic files from the OpScan machines, panel voter registration, electronic poll books. And the computers that aggregate the results together from all of the optical scans.”

“If any of those get hacked, it could could significantly disrupt the election.”

The digital touch screens, even with voter verified paper trail, will still be pervasive this election; 28 states keep them in use to some degree, including Ohio and Florida, though increasingly in limited settings. Pam Smith, the director of Verified Voting—a group that tracks the use of voting equipment by precinct in granular detail—isn’t sure how many digital touch screens are left; no one I spoke with seemed to know. Nor is it clear where they’ll be deployed, a decision left up to county administrators. Smith confirms that after 2007, the number of states that adopted the machines plateaued, and has finally begun to shrink. The number of states using paperless touch screens—and nothing else—is five: South Carolina, Georgia, Louisiana, New Jersey and Delaware. But the number of states with a significant number of *counties* with the easily hacked machines is much larger, at 13, including Indiana, Virginia, and Pennsylvania. For hacking purposes, there’s little difference: In a close election, only a few precincts with paperless touch screens would be required to deflate vote totals, says Appel, even if the majority of counties are still in the Stone Age. Many of Felten’s mad-scientist experiments were designed to metastasize the nefarious code once it gained entry into a machine system.

The move away from electronic voting is a positive one, the professors say; the best option for election security are the optical scans. “Although the optical scan ballots are counted by the computer in the OpScan machine—which you can’t trust—you can trust the pile of ballots that accumulate in the ballot box, marked by users with their own hands,” Appel tells me. With the right auditing policies, “you can recount or do a statistical sample of the ballot boxes to make sure there aren’t cheating computers out there.”

State policymakers listened. In 2000, less than 30 percent of voters used the optical scanning

system. In 2012, 56 percent did. But in the interim, the touch-screen machines are still in place; their dwindling percentage of votes has not necessarily diminished the risk of an attack, the professors say. In some ways, it's heightened it—turning the issue of easy-to-tamper touch screens from a bell-curve problem to a hockey-stick graph, in which a small number of machines generate a high amount of risk. The machines that are left are often running on vintage Windows software from the late '90s or early 2000s, some of which has long surpassed its support date. “They’re probably about exactly as vulnerable as they were 10 years ago,” Appel tells me. “And they still get their program out of the same ROM.”

A study released by the Brennan Center last September, titled “Voting Machines at Risk” reached a similar conclusion. In 2016, 43 states will use machines that are at least 10 years old; 31 states suggested a serious need for new voting machines. Larry Norden, the report’s author, said everything from software support, replacement parts and screen calibration were at risk; he pointed me to a YouTube video of a precinct in West Virginia, where voters’ finger pressure on the screen selected an entirely different candidate, or caused the machine to go haywire (a symptom of the glue behind the screen loosening, Norden says). The HAVA money, says Wallach, was spent very quickly after 2002; “And it is not coming back,” he adds.

As late as 2011, a team at the Argonne National Laboratory of the Department of Energy revisited the Diebold TSX, five years after the Princeton group’s report. Its conclusion: With \$26 worth of parts and an eighth-grade understanding of computers, virtually anyone could tamper with it—a variant of the model that Feldman and Halderman procured in the Times Square alleyway. Five years later, cyber experts tell me that little has changed in voter cybersecurity. The Diebold TSX model is slated to be used in 20 states in 2016, including Pennsylvania, Ohio, Florida, Missouri and Colorado.

State officials recognize that digital touch screens are headed out the door—and the professors are quick to remind me of how government contracts work: When profit projections fall, upkeep suffers. “The level of security confidence when it comes to these voting machines is much lower than the sort of industry standard—the level of security you’d expect from top companies like Google, Facebook, Apple. I mean, your iPhone is probably much more secure than most of these voting machines,” says Ari Feldman, one of Felten’s acolytes and now a professor at the University of Chicago. “I think the level of technological competence of the people who work on these very popular commercial services and devices is just higher than those who these small voting machine manufacturers can attract.”

No one doubts that the companies take security seriously. But the approach to security shared by the manufacturers and election officials seem to hinge on the idea that hacking a school board vote would be just too boring for anyone talented enough to pull off. “You would be hard pressed to find

an example of our voting systems ever being hacked in a real election environment, as opposed to that of a hack attempt inside of a laboratory environment in which zero real world physical election processes are utilized,” writes Kathy Rogers, a spokesperson with ES&S, in an email, and correctly so—it’s never been proven that an election was deliberately hacked. “We feel *very* confident in the security of our voting systems—especially when you combine that security with the physical security, chain of custody, legal requirements and masses of pre-election testing.” She added, “We are not suffering from sleepless nights worrying about whether our voting systems might be hacked.”

A Virginia election official with decades of experience concurred, speaking to me on background. “I know that when some of the academics have hacked a machine, they’ve had unfettered access for an indefinite period of time,” the election official said, describing this as an unrealistic precondition. “But one of the security thresholds isn’t that it will be sitting in a public location here so anyone can have unfettered access for any in-depth period of time.” He demurred when I brought up Felten’s tradition of stalking the unguarded machines; he added, “Only people who have been authorized, sworn to uphold the process—they can have administrator access to these.

“It’s old school, I realize that,” he continued. “But it is the system in place.”

In the event of a state-sponsored attack—however unlikely—can old school match wits? The adversary, more than one member of the Princeton group pointed out, may be more practiced than we know: A June 2014 report linked Russian hackers to an attempt to alter the election outcomes in Ukraine, by targeting the computerized aggregation software—one of the attacks Appel fears.

How different is Kiev from Gary, Indiana? As is the case in cyberattacks—at least in the examples of Stuxnet and Sony—it’s never quite plausible, until it is. Hackers this year have targeted voter registration rolls in Illinois and possibly Arizona, another attack highlighted by the Princeton alums.

But most identified Pennsylvania as the greatest concern. There, according to Verified Voting 47 counties of 67 vote on digital voting machines without a written backup record if something were to go awry—a reality that is very much on the minds of state officials (legislation is working its way through the House to examine the issue of voting modernization.) In Pittsburgh and Philadelphia—two Democratic strongholds whose turnout typically decide the fate of the state’s outcome—around 900,000 voters will cast ballots entirely on paperless touchscreens DREs, if previous elections are any guide. Then, at least from the voters’ perspective, they will disappear into a sea of ones and zeroes.

Montgomery County, a crucial Democratic redoubt in the suburbs of Philadelphia—an area sometimes seen as having the potential to swing the entire state—is one such locality that uses a

paperless electronic machine, and only one machine, for all 425 precincts: Appel's Sequoia AVC Advantage.

"We are very, very confident in our machines," Val Arkoosh, the vice chair of the Montgomery County Board of Commissioners, tells me. She spoke with the staccato fervency and granular detail of someone who is thinking about this issue, and has been asked before. Yet when I asked her about Appel's hack and the Princeton group, next door across the Delaware River, she appeared not to have heard of it. She assured me their system is secure: "We program each of our machines individually—they're never connected to the Internet," and an internal hard drive "creates a permanent record each time that a vote is cast." At the end of the day, Arkoosh said, "the vote is transcribed on a thermal tape, the machines are closed to lock, the information is transferred to a standalone server that tallies the results." She describes the officials guarding the polling place, and adds for emphasis: "It would be extraordinarily difficult for someone to do something like that during the course of Election Day."

I asked Halderman to red-team Arkoosh's answer. "It's positive that they have procedures in place to cross-check that the counts produced by each machine match the tabulated results," Halderman wrote to me in an email. "However, none of that provides any defense against the kinds of attacks Andrew Appel wrote about, or the return-oriented programming attacks." He added, "An attacker with access to the administration system that's used to program the memory cartridges before the election could use ROP to distribute malicious code to all the machines."

"I can say that this is definitely a concern," says Kelly Green, the director of Voting Services in Montgomery County, who continued to describe efforts and conversations across Pennsylvania to improve the voting system. As a state issue, Green continued, "What I can tell you is, we've put it on the agenda."

What would be the political motivation for a state-sponsored attack? In the case of Russia hacking the Democrats, the conventional wisdom would appear that Moscow would like to see President Donald Trump strolling the Kremlin on a state visit. But the programmers also point out that other states may be leery. "China has a huge amount to lose. They would never dare do something like that," says Wallach, who recently finished up a term with the Air Force's science advisory board. Still, statistical threat assessment isn't about likelihoods, they insist; it's about anticipating unlikelihood.

The good news is that Wallach thinks we'd smell something fishy, and fairly fast: "If tampering happens, we will find it. But you need to have a 'then-what.' If you detect electronic tampering, then what?"

No one has a straight answer, except for a uniform agreement on one thing: chaos that would make 2000 look like child's play. (Trump aping about "rigged elections" before the vote is even underway has certainly not helped.) The programmers suggest we ought to allow, for the purposes of imagination, the prospect of a nationwide recount. Both sides would accuse the other of corruption and sponsoring the attack. And the political response to the country of origin would prove equally difficult—the White House is reported to be gauging how best to respond to the DNC attack, a question that poses no obvious answers. What does an Election Day cyberstrike warrant? Cruise missiles?

The easiest and ostensibly cheapest defense—attaching a voter-verified paper receipt to every digital touch screen—presents its own problem. It assumes states audit procedures are robust. According to Pam Smith at Verified Voting, over 20 states have auditing systems that are inadequate—not using sufficient sample sizes, or auditing under only certain parameters that could be outfoxed by a sophisticated attack—states that include Virginia, Indiana and Iowa. But relying on paper trails also assumes voters understand their importance. Many may simply discard the paper on the way out without giving it a glance, or leave it hanging in the machine printer.

Optical scanning machines are far and away the first choice of the programmers—as the Princeton group analogizes, they don't require receipts, they *are* the receipts—and states are increasingly ditching touch screens in favor of them. But the optical scans are still DRE models—we simply push paper, rather than push buttons. Jeremy Epstein, the Virginia computer scientist who led the charge against the WINVote system, points out that digital touch screens and optical scanning machines have something in common: "Whether it's an optical scanner or a DRE, the votes still get totaled on a memory card. And at the end of the election, you put that memory card into a central card system," Epstein tells me. "You could use it to infect the tabulator system, and once you infect the tabulator system, it could transmit on."

Then there are tech advancements that make the computer scientists shudder: To a person, they each warned me about the public's new delusion, one strikingly reminiscent of the aftermath of *Bush v. Gore*—Internet voting. As Halderman's work began to garner more attention, he sensed a new trend around the idea of voting online. With its lack of technical probity, an argument hanging entirely on convenience, and a stampede of purveyors from for-profit cyber companies, Halderman and others saw a facsimile of the voting machine companies they had sought to marginalize just years earlier. Yet elected officials found appeal in many of the same arguments. "In this world, we do so many things now online," Appel says, explaining the popularity of the idea. "You're banking online. You order coffee online. Somebody who's used to living so much of their life online will wonder why we're not voting online."

But Appel, and the others, share a categorical warning: "It would be a disaster," he tells me.

“Anyone could hack in. The Russians, the North Koreans, anyone who wishes.”

Like the voting machine companies, Internet voting services—mostly purveying their software in private or corporate elections—largely resist subjecting their work to public trial. That changed when, in 2010, the District of Columbia announced its intention to launch a citywide Internet voting platform, intended for overseas voters and a milestone for the concept. Just a month before the midterm elections in November, the District conducted a test drive. “It’s not every day, of course, that you’re invited to hack into government computers without going to jail,” Halderman says, muffling a giggle. “We didn’t want to let this opportunity, to have this be a realistic simulation of an attack, go to waste.”

On October 1, 2010, two employees in the Washington, D.C.-based Office of the Chief Technology Officer, stormed down a hallway and charged through the double-doors that opened into the basement-floor server room. Earlier that day, they had learned strange news: Someone had called into the hotline to report a bug on the board’s paperless ballot system. The program seemed to play obnoxious brass-band music each time subjects submitted their ballot. The names on the ballots had all been changed to villainous robots: Bender for State Board of Education (from *Futurama*); Hal 9000 for Council Chairman (from *2001: A Space Odyssey*). Then they learned that the hackers were likely watching them on the closed-circuit circuit feed, through the camera that was gazing down at them, right now.

Some 520 miles away, the scene played on a screen in the hacker’s cramped headquarters. A whiteboard behind the computer declared a series of instructions in brown and purple marker, each skewered with a squiggly strike-through, followed by a perfunctory checkmark: “Replace old ballots.” Check. “Steal temp ballots. Check. “Rig to replace new ballots.” Check. The hackers exchanged high-fives in adulation. And when the D.C. tech officers’ faces appeared on the screen, Alex Halderman peered back.

Halderman, now a professor at the University of Michigan, had not lost his mentors’ taste for the dramatic. He had just pulled the most flamboyant hack in the short history of the Princeton group. Halderman was called before the D.C. Council, where he got to make the speech he wanted before a captive audience, who were forced to endure this barely 30-year-old’s transported lecture seminar on the dangers of Internet voting.

Halderman shared a private, unreleased video with me that he took from the night of the attack, a project he launched with the help of two graduate students, each barely out of college. In the video, the team huddles around Halderman’s small, beechwood office table, assuming a crouch in a strange coven of furious typing. Hours pass as afternoon tips into evening. Finally, a brown-haired student, Eric, slouched and raccoon-eyed, bolts upright: “Oh my God,” he murmurs. “I have a

shell.” “We’re in!” shouts his blonde-haired compatriot, rubbing his hands. The furious typing resumes.

Halderman explained that the student had used a technique called Shell Injection Vulnerability. He found a single, wayward quotation mark in the code, a crack in the floorboard through which they drove a tractor-trailer of attack commands.

Halderman’s attack is now well-known in the world of elections administration; the Virginia election official I spoke with seemed doubtful that Internet voting could ever take off, citing the conventional view that the risks are too great. “Whether or not Internet voting happens, and whether we will introduce these new risks—I don’t know,” he says. “I’m not holding my breath.”

Internet voting companies have the same incentives as voting tech conglomerates to convince the public they’re worth their mettle; as in the case of HAVA, there would likely be an enormous windfall. In 2004, Michigan deployed Internet voting in its Democratic primary. In 2009, West Virginia greenlighted a pilot to allow overseas military vote online. This year, the entire 2016 Utah primary was conducted online, and an initiative in California to introduce online voting nearly made it onto the state ballot.

Halderman finds it hard to believe he now has to make the same argument about the risk of hacking all over again. “It’s not something only comic book villains can do,” he explains. “These are students right out of college that are doing this.”

The concept of voting in private is an invention in American politics, and a recent one. The first time a secret ballot was widely deployed was the presidential election of 1896—also the first election in which someone was not murdered on Election Day, according to Harvard professor Jill Lepore. The two are not a coincidence: Since the earliest days of the republic, voting was almost entirely a collectivist act. Citizens voted with their feet—standing on one side of a crowd or another, caucus-style—a setup which manipulative party bosses plainly preferred.

The cadre of computer programmers who made their home on the Princeton campus are now in a race, of sorts—against voting machine companies, against Internet voting firms—to invent the future of secure voting. And the most interesting ideas look to this 19th century arrangement not with revulsion, but intrigue. It turns out that, from the perspective of mathematical systems confirmation, Boss Tweed may have had a few things right.

After his testimony in Houston urging the council not to adopt the machines, Wallach, the Rice professor, spent the proceeding years working on research showing vulnerabilities on digital touch

screens, and testifying in state legislatures across the country. But Wallach's focus has shifted from diagnosis to cure, and he's now working with Travis County, where Austin is located, as a leading researcher on the newest innovation in voting technology: Cryptographic voting.

Wallach walks backward through the concept by offering a thought experiment. The most unimpeachable election technique would be to count the votes on an enormous corkboard; every voter would pin his or her vote, and the public would count the results together. Everyone would see the votes, and everyone would agree on the result. Besides the problem of privacy and intimidation (and, ostensibly, killings on Election Day), such a system is ungainly—it's a lot of corkboard. But encrypting the vote would allow a public accounting while keeping the actual votes private: voters would make their selection on a digital processing machine; they'd then receive an encrypted receipt, a random assortment of numbers and letters. Their vote would then be uploaded to a public bulletin board online; any voter could compare their encrypted vote to see if it matched the numbers and letters online. The vote itself would be scrambled and completely secret; a complex function, known as homomorphic cryptography, would count the votes without unencrypting the source.

“Crypto,” as it's known in the field, would secure our elections something close to permanently. But it would change fundamentally the way we vote. It would make the act of gawking at random source code a civic requirement. And it would abolish the concept of a countable “ballot,” forcing us to trust that incomprehensible code is the equivalent of a ballot. Cryptographic voting is still years away from ready. But it also begs the question of whether the concept has simply transferred a technocratic leap of faith from one part of the electronic system to another one. It seemed difficult to believe, after a bruising decade of invisible votes and disappearing ballots, that voters would put their faith in something so abstract. After four explanations from Wallach, I was still dumbfounded.

Wallach and other researchers point to another safeguard that is closer to application-ready, a new method of auditing. The technique is called Risk Limited Auditing, statistical innovation worked out by Philip Stark, a statistics professor at the University of California, Berkeley. The auditing techniques of most states aren't sophisticated enough to detect a subtle attack—every 100th vote switched from Trump to Hillary Clinton, for instance. “The whole point of a Risk Limiting Audit is not to find the tally down to the last digit,” explains Wallach. “The problem you're trying to figure out is if the error rate is big enough that I could change who won.” RLA would enhance the auditing prospects of most states, 25 of which have inadequate auditing procedures, according to Verified Voting. Colorado is expected to implement RLA next year.

But there may be a simpler hack at hand. Appel, the Princeton cybersecurity expert—master of numbers, merry prankster of machines—proposes a radical idea to this 15-year nightmare: What if

we took a page from the town criers of two centuries ago, and simply read the precinct results out loud?

“There’s a very simple and old-fashioned recipe that we use in our American democracy,” Appel says. “The vote totals in each polling place are announced at the time the polls closed, in the polling place, to all observers—the poll workers, the party challengers, any citizen that’s observing the closing of the polls.” He goes on to describe how the totals in that precinct would be written on a piece of paper—pencils do just fine—then signed by the poll workers who have been operating that polling site.

“Any citizen can independently add up the precinct-by-precinct totals,” he continues. “And that’s a very important check. It’s a way that with our precinct-based polling systems, we can have some assurance that hacked computers could not undetectably change the results of our election.”

There could be a greater lesson in Appel’s point. Technology didn’t create the problem. Perhaps technology is intrinsic to the problem—our lack of trust that has metastasized in a surveillance culture was bound to aggrandize the problems of voting, the most trusting civic act we know. It seems unlikely to expect a singular cure to the American presidential election, not because of the incomprehensibility of cryptography or the untrustworthiness of tech companies, but because there is no such thing as the singular election: 8,000 jurisdictions in a leaky mess of federalism and poorly spent dollars. The neat results and cable announcements on election night represent an optical illusion, like a series of ones and zeroes, whizzing beyond our apprehension.

Wallach’s encomium on cryptography reminded me of another tech item: The concept of shared fate, sometimes referenced in drone research. Researchers have long suggested our planes and trains could be made safer were they run by highly precise robots, or drone pilots—cool customers who don’t have to save a burning plane while worrying about turbulence and screaming passengers. It may be one of the most enduring examples of psychology trumping technocracy: Even though systems would run better—even save lives—everyone knows this arrangement is unworkable. Humans require knowing that there’s someone, like us, in the cockpit. We need to know we’ll endure a shared fate.

If this century has shifted our trust from away from our neighbors toward machines, it might be time to switch back again. Eight countries in Europe that once flirted with digital voting have seen six go back to paper; Britain counted its Brexit votes by hand. Even if the vote were never hacked—and it is an exceedingly implausible event—the remotest possibility is an albatross on democracy and a boon for mischief-makers, and not just the cyber attackers. Trump’s most recent jujitsu—pointing out that by virtue of the fact that the election is hackable, it could be rigged against him—illustrates this risk. Technology has amplified not only the threat of hacking, but the threat of a

hack.

The Princeton alums can warn us—but they can't protect us. "We are in a collision-course between the technology we use in election administration and the growing reality of politically motivated, statelevel cyberattacks," Halderman tells me, arm propped on his red office chair, sunlight pouring through his westward window. "We sit around all day and write research papers. But these people are full-time exploiters. They're the professionals. We're the amateurs."