Abstract

We discuss ways in which Internet voting differs from e-commerce, analyze the threats to Internet voting, review some key studies and reports, describe several elections and pilots held over the Internet, and reflect on the future of Internet voting.

1 Internet Voting Is Much Harder than e-Commerce

There is a widespread belief that since we use the Internet for commercial activities, we should be able to return voted ballots over the Internet. There are two problems with that belief. First, cybertheft has accompanied the rise of e-commerce, resulting in significant financial losses. The advanced cybertheft tools that we are now encountering could also be used to steal Internet-based elections. Second, even if e-commerce were secure, there are issues with Internet voting, such as the secrecy of the ballot, that do not arise with e-commerce. We begin with a discussion of threats to e-commerce.

A survey of roughly 1000 large organizations stated that in 2008 the average intellectual property loss per company to cybertheft was about $4.6 million [Kanan, Rees, and Spafford (2009)]. According to a December 2009 report from the Computer Security Institute, a different survey of 443 companies and U.S. government agencies found that 64% had reported (malicious software) infections in the preceding year [Lohr (2010)].

A recent major China-based Internet attack on Google and dozens of other companies illustrates that even major corporate sites are vulnerable. The attack targeted Google intellectual property, including systems used by software developers to build code, as well as Gmail accounts of Chinese human rights activists [McMillan (2010)]. As many as 34 companies were attacked, including Yahoo, Adobe, and Juniper Networks. It appears that even defense contractor Northrop-Grumman and Symantec, a major supplier of anti-virus and anti-spyware software, were targeted [Vascellaro and Solomon (2010)]. The attacked companies employ large numbers of computer security experts and have vastly more security expertise and resources than relatively small Internet voting vendors.

The attacks looked as if they came from trusted sources, so that victims would be tricked into clicking on a link or file. Then, using a vulnerability in Microsoft's Internet Explorer, the attacker downloaded and installed malicious software, thereby gaining complete control over the compromised system [Kurtz (2010)].

Government sites are also vulnerable. In a March 2010 talk, FBI Director Robert Mueller stated that the FBI's computer network had been penetrated and that the attackers had "corrupted data" [McMillan (2010)]. General Michael Hayden, former Director of both the CIA and National Security Agency, said in a recent interview: “The modern-day bank robber isn’t speeding up to a suburban bank with weapons drawn and notes passed to the teller. He’s on the Web taking things of value from you and me.” [Hayden (2010)]

Hayden added in response to a question about whether or not policy-makers are up to speed technically: “The technology and the operational art of this thing in cyberspace is way beyond any of the policy lines that we have even begun to think about. Policy has to catch up. And that’s going to take a lot of work and a lot of conversations between tech-savvy people and policy-smart people.”

The implication of comments by Mueller and Hayden, as well as the recent major attack on Google and other companies, is that voting system software could be attacked by outside hackers, including attackers from another country. Perhaps even more disturbing is that the attackers could engage in a “false flag” attack in which the attackers make it appear that the attack is coming from a different source – for example, Iran. Such an attack could be very destabilizing.

1.1 An example: The Zeus Virus

In April 2009 we learned that malicious software, which redirected visitors to an IP address in Amsterdam, had been inserted into Paul McCartney’s website. The Amsterdam site contained software that exploited vulnerabilities on the victims’ machines to download the very dangerous *Zeus virus* onto those machines [McCartney]. The infection, planted shortly before McCartney’s New York reunion concert, was timed to catch as many victims as possible before discovery.

The goal of the Zeus virus is to steal money from on-line financial accounts. When the victim logs onto her bank’s website, Zeus copies her credentials and sends them to a remote location. It then uses those credentials to remove money from her account. Zeus can even simulate the victim’s financial statements [Zeus09]. When the victim checks her on-line statement, everything looks as the victim thinks it should. The victim typically learns of the theft only when checks start to bounce or financial transactions cannot be finalized because of insufficient funds. At that point, it’s too late to retrieve the money, which has been sent off-shore.

The German edition of Wikipedia was also a source of infection [Head (2009)]. A bogus article about another dangerous piece of malware (the Blaster worm) contained a link to software that supposedly would fix the problem. However, anyone who downloaded the “fix” was actually downloading a copy of Zeus.

More recently, the Zeus virus has attacked verification systems used by Visa and MasterCard when enrolling new users [Dunn (2010)]. By mimicking (spoofing) the enrollment process, Zeus obtained sensitive information such as social security number, card number, and PIN or verification code from the unknowing victim, who would think that he was providing that information to the bank. This information, sent to the attacker's computers, was used to defraud the victim. See Figure 1 for an illustration.

According to an August 2010 report from M86 Security Labs, about 3000 bank customers in the U.K. were victimized by a form of the Zeus virus [M86 Security Labs (2010a)]. The announcement accompanying the report's release, which did not provide the bank's name, described the attack [M86 Security Labs (2010b)]:

> Unprotected customers were infected by a Trojan - which managed to avoid detection by traditional Anti-Virus software - while browsing the Internet. The Trojan, a Zeus v3, steals the customer's online banking ID and hijacks their online banking sessions. It then checks the account balance and, if the account balance is bigger than GBP 800 value, it issues a money transfer transaction.
> From July 5, the cyber criminals have successfully stolen GBP 675,000 (c. USD 1,077,000) and the attack is still progressing.

It was estimated in 2009 that Zeus has infected about 3.6 million PCs just in the U.S. [Messmer (2009)].

1.2 Why Voting is Different from e-Commerce

Internet voting is vulnerable to a multitude of attacks from malware such as Zeus. The problem of simply modifying a vote before sending it to its destination is relatively easy compared to stealing money and projecting a rigged bank statement on the victim's monitor. In addition, internet voting has problems that do not arise with e-commerce.

- To prevent my being pressured, threatened, or paid to vote a certain way, my vote must be private. When I vote, I want my vote to be counted, but I don't want my name attached to my vote. (The U.K. is a notable and for me disturbing exception). Ballot secrecy greatly complicates Internet voting, by preventing election officials from verifying to me that they have received a correct copy of my vote. By contrast, when I make an Internet purchase, I want the seller to know who I am and what I'm buying.
- If a commercial website fails, transactions conducted prior to the failure are not impacted, and people can return to the website later. But, failure of a voting website on Election Day could mean that some people get to vote and others do not.

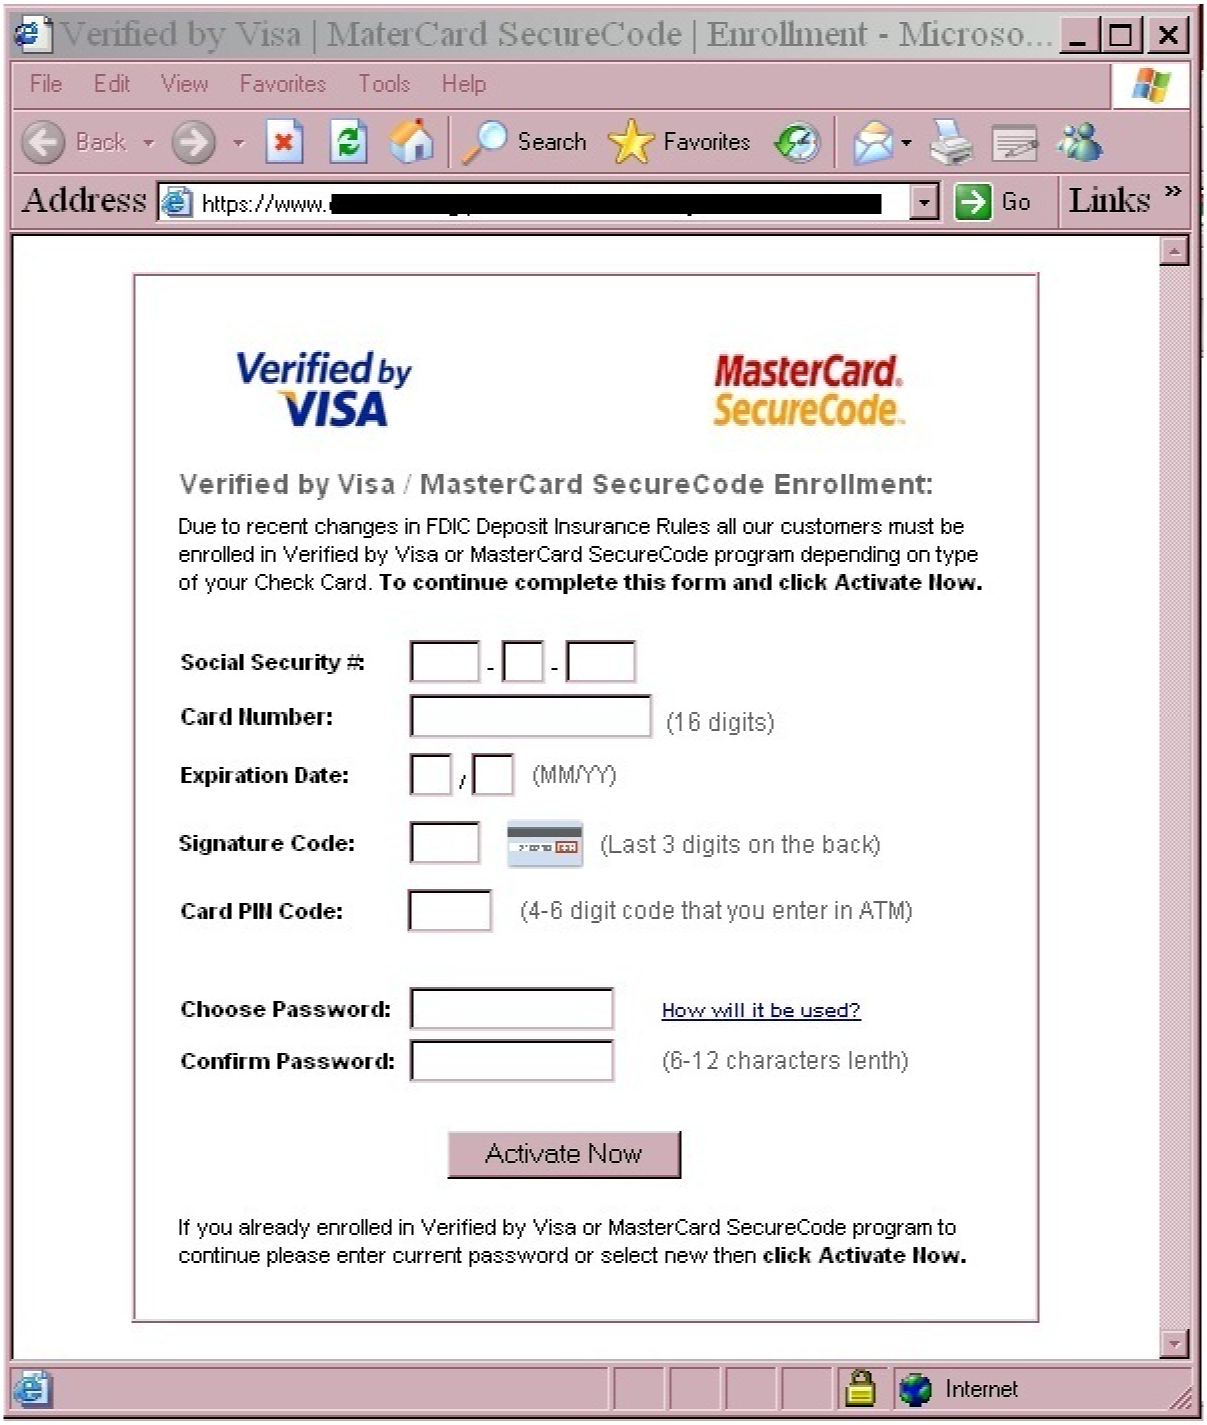

https://www.[redacted]

Verified by VISA

MasterCard. SecureCode.

Verified by Visa / MasterCard SecureCode Enrollment:

Due to recent changes in FDIC Deposit Insurance Rules all our customers must be enrolled in Verified by Visa or MasterCard SecureCode program depending on type of your Check Card. **To continue complete this form and click Activate Now.**

Social Security #: ___ - ___ - ___

Card Number: ___ (16 digits)

Expiration Date: ___ / ___ (MM/YY)

Signature Code: ___ (Last 3 digits on the back)

Card PIN Code: ___ (4-6 digit code that you enter in ATM)

Choose Password: ___ How will it be used?

Confirm Password: ___ (6-12 characters lenth)

Activate Now

If you already enrolled in Verified by Visa or MasterCard SecureCode program to continue please enter current password or select new then **click Activate Now.**

Figure 1: The Bogus Enrollment Screen Used by Zeus for the Credit Card Attack. Photo by Amit Klein, CTO of Trusteer.

- I know if a book ordered from amazon.com does not arrive. With most commercial Internet voting schemes, however, I will never know whether or not my vote was received. I could print out a copy of my ballot, but so long as that paper ballot stays with me and is not in the possession of election officials, it cannot be a check on the electronic version of my ballot, nor can it be used to audit or recount the election. An email confirming my vote is unreliable, since the email could be forged or my infected computer could lie to me, much as the Zeus virus lies to its victims about the contents of their bank accounts.
- In the event of an e-commerce failure, there is a good chance that the situation will be rectified. If I notify amazon.com that my book has not arrived, they probably will send me a replacement book. If I don't receive a replacement, I can stop my credit card payment. But, if my vote is not successfully cast on Election Day, I probably won't know, and I almost certainly will be unable to revote.
- Since Internet voting is done remotely, it is difficult to verify a person's identity. It might be legal for my spouse to make a purchase using my credit card. But in most countries it is illegal for my spouse to vote on my behalf, unless for some reason I am unable to do so myself, e.g. I have a disability that prevents me from voting independently.
- amazon.com might try to sell me additional products, or my browser, when it notices that I'm connected to amazon.com, might display an ad in a pop-up window. While I find such actions annoying, they are not illegal. However, with online voting it would be possible to display political ads while I'm actually voting. That would be like posting political posters in voting booths, something that may not be currently illegal, but probably should be.
- Unlike commercial activities, vote buying and selling is illegal. In the 2000 U.S. Presidential election, an online system designed to broker Nader and Gore votes was created, but it was forced to shut down by the California Attorney General. There is no evidence that any votes were actually traded. Internet voting could eliminate the need for an honor system by allowing voters to sell their voting credentials on a website. The website might even automatically cast the actual ballot.[1]

[1] There are many legitimate situations, such as family members voting on their home computer and people voting from the same work location, in which multiple voters have the same Internet location (IP address). Therefore, while it would be possible to detect multiple votes coming from the same IP address, it would be problematic to prohibit such votes.

2 Other Pro-Internet Voting Arguments

It has become almost a matter of faith in some circles that Internet voting will increase voter participation, especially by young people. For example, in 2000 the Knight-Ridder News Service stated that "Californians could be voting over the Internet in five years with a computerized system that could revolutionize the state's voting process and boost sagging voter turnout" [Chaney (2000)]. However, a May 2009 a local election in Honolulu that allowed people to vote only over the Internet or via phone saw an 83% drop in voter participation when compared to the previous similar election in 2007 [KITV (2009)]. We do not know the causes for the steep drop in voter participation. (The vendor, Everyone Counts, also ran the 2007 Swindon pilot Internet election, discussed below).

Furthermore, despite claims to the contrary, Internet voting will not necessarily save money; in fact, Internet voting can be very expensive. For example, 2009 cost estimates from Everyone Counts were so large that Washington State legislation that would have allowed Internet voting for UOCAVA voters, namely military and overseas U.S. voters, was killed in committee [Bonifaz (2009)]. The estimated costs included [DeGregorio (2009)]:

- 6 weeks consulting to understand state requirements: $40,000 - $60,000;
- 10 weeks consulting to develop processes specific to state requirements: $100,000 - $200,000;
- 12 weeks consulting to conduct initial pilot + 8 weeks engineering/consulting to evaluate the pilot and recommend changes: $300,000 - $700,000;
- 12 weeks consulting to conduct training for new counties, conduct pilots, evaluate the processes, and recommend for next pilot: $500,000 - $900,000;
- 20 weeks consulting to conduct training, build ballots, test ballots, conduct the general election, evaluate the processes, and make recommendations for the next election: $1,500,000 - $2,500,000;
- 4 weeks consulting to provide the final report, etc.: $60,000 - $80,000;
- Primary and general election ongoing annual license fee per county: $20,000 - $120,000 per year + $2 - $7 per registered UOCAVA voter.

In other words, upfront cost for only residents of Washington State who were military or civilians living abroad, including running the initial election, would have ranged from $2,500,000 to $4,440,000. After that, each county would have been hit with an annual fee, combined with addition per-voter charges.

3 Internet Voting is Insecure

Suppose you are going to vote in a national election over the Internet. On the morning of Election Day you receive an email alert from your political party, or so it appears, warning of a new virus infecting the Internet and offering virus protection software for the election. Would you download the software? What if the download actually contained an election-rigging virus? What if you were wary, but others were not? Even if you were wise enough not to download questionable software, an election rigging virus might need to infect only a relatively small number of machines to change the outcome of an election.

In addition to the security hazards of current Internet technology, Internet voting combines the risks of absentee voting, such as ballot coercion and vote buying and selling, with the risks of paperless computerized voting machines.

Add to that already insecure mix the political and financial motivation to steal an election, and we are forced to conclude that Internet voting in high stakes elections is a very dangerous idea.

3.1 The Computer Technologists' Statement

Because of the push for Internet voting, computer technologists from major U.S. universities such as Stanford, Berkeley, Princeton, and Yale, as well as research labs and private industry, signed the Computer Technologists' Statement on Internet Voting [Computer Technologists]. The statement observes that "serious, potentially insurmountable technical challenges" would have to be dealt with before secure Internet voting should be deployed. These major challenges include:

- The voting system as a whole must be verifiably accurate in spite of the fact that client systems [the voters' computers] can never be guaranteed to be free of malicious logic [malware]. Malicious software, firmware, or hardware could change, fabricate, or delete votes, deceive the user in myriad ways including modifying the ballot presentation, leak information about votes to enable voter coercion, prevent or discourage voting, or perform online electioneering. . . . [T]here is no guaranteed method for preventing or detecting attacks by insiders such as the designers of the system.
- There must be a satisfactory way to prevent large-scale or selective disruption of vote transmission over the Internet. Threats include "denial of service" attacks from networks of compromised computers (called "botnets"), causing messages to be mis-routed, and many other kinds of attacks, some of which are still being discovered. Such attacks could disrupt an entire election or selectively disenfranchise a segment of the voting population.

- There must be strong mechanisms to prevent undetected changes to votes, not only by outsiders but also by insiders such as equipment manufacturers, technicians, system administrators, and election officials who have legitimate access to election software and/or data.
- There must be reliable, unforgeable, unchangeable voter-verified records of votes that are at least as effective for auditing as paper ballots, without compromising ballot secrecy. Achieving such auditability with a secret ballot transmitted over the Internet but without paper is an unsolved problem.
- The entire system must be reliable and verifiable even though Internet-based attacks can be mounted by anyone, anywhere in the world. Potential attackers could include individual hackers, political parties, international criminal organizations, hostile foreign governments, or even terrorists. The current Internet architecture makes such attacks difficult or impossible to trace back to their sources.

The statement then warns that in light of the known problems, any Internet voting proposal should include a public disclosure of the principles of operation in sufficient detail that anyone with the appropriate background could verify that election results generated by the proposed system can be trusted.

In addition, the statement cautions against "pilot studies," since the appearance of success can be deceiving. Not only could serious problems go undetected, but potential attackers might avoid attacking pilot studies and instead wait until full-scale elections are held over the Internet.

Below we expand on some of the technologists' concerns.

3.2 Viruses and Worms

There are many methods for distributing malware. Just as a human virus spreads from human to human, a computer virus spreads from computer to computer by attaching itself to an email or an input device such as a CD. A virus or worm can infect public or privately owned machines linked to the Internet without the owner's knowledge or permission. If a virus or worm replicates quickly, or if it is widely distributed via popular web pages or candidate websites, attached to spam, etc., it might succeed in infecting a large number of machines prior to detection. Once a computer is infected, the attacker might be able to modify or destroy the voter's ballot, as well as steal information, delete files, and read email. Many viruses or worms allow their creators to seize control of infected machines.

Infected computers that are remotely controlled by the distributor of the malware are called "zombies", and a network of infected computers is called a "botnet" (bot is short for "robot"). There are known botnets that consist of millions of infected computers. According to the FBI, the Mariposa Botnet may have infected 8,000,000 to 12,000,000 computers internationally [FBI (2010)]. The virus used to create the Mariposa Botnet can steal credit card data and online banking passwords, as well as prevent people from accessing a particular website (*denial of service attack*). The creator of the virus also sold customized versions with augmented features.

A large botnet could have a major impact on the outcome of an Internet election. Furthermore, if a virus or worm erases itself after the person votes, it might be impossible to determine that the malware was ever present. The damage, however, will have been done.

Anti-virus software works by checking for known viruses and worms. Therefore, whenever a new virus appears, the anti-virus software must be updated. There can be many days or even weeks between the time when the virus is initially distributed and when it is recognized and analyzed. After that, the virus fix needs to be distributed, and users need to disinfect their machines.

Because anti-virus software has limited capabilities for recognizing unknown malware, a new virus or worm may well escape detection for a while, especially if it remains dormant until, say, Election Day. Even if the virus or worm is detected, removal can be quite difficult, as most PC owners who have had to deal with adware and spyware are well aware.

One study showed that anti-virus software has become less effective over time, with recognition of malware by most anti-virus software falling from 40% - 50% at the beginning of 2007 to 20% - 30% at the end of that year [H Security (2007)]. Another set of experiments conducted at the University of Michigan demonstrated that the number of malware samples detected decreased significantly as the malware became more recent; when the malware was only one week old, the detection rate was very low [Oberheide, Cooke, and Jahanian (2008)]. Given that an election-stealing virus or worm would likely be released far in advance of an election and spread silently with no symptoms so as to infect the maximum number of machines, it could be impossible to determine if votes were modified or even (if the malware erases itself after modifying the vote) which computers had been infected.

Hackers, criminals, or foreign countries can manipulate Internet-based elections by inserting election-rigging software into voters' computers. One of the easiest way to do this is by distributing a virus or worm over the Internet. The Conficker worm, discovered in November 2008, rapidly infected between 9 and 15 million machines [Wikipedia (2010)]. The Conficker worm is especially worrisome, because it has the capability of "calling home" for more instructions. The new instructions could be targeted at specific candidates and elections. In other words, the unknown creator of Conficker can instruct an infected machine to remotely install additional malicious software without the computer owner's knowledge [Symantec Corporation (2009)]. Such software could be fine tuned for each election.

While many viruses and worms are planted without the knowledge of the computer owner, people are sometimes willing to download and install software of highly questionable provenance. In August 2009 a spam message circulate that said, "If You dont like Obama come here, you can help to ddos his site with your installs."[*sic*] (ddos stands for Distributed Denial of Service, an attack that overwhelms a website, making it inaccessible). According to CNET News, people who clicked on the email link were offered money in exchange for downloading denial of service software. They were even told to return to the website for updated versions if their virus detection software deleted the original download [Mills (2009)].

The source of the software is not known. The goal could have been to disrupt websites associated with Obama, to engage in identity theft, or even to infect machines of Obama opponents, something that could be especially useful if Internet voting were to become an option in the United States.

In light of the successful attacks being conducted against governments, major banks, and many of the world's technology leaders, it should be easy to entrap large numbers of voters who are not technologists. Once a voter's computer is infected by a malicious virus, all bets are off. The virus can make the computer screen show the voter a ballot image that correctly represents the voter's intent. But the virus can then send something entirely different over the Internet, and the voter will never see the actual ballot that is sent. In other words, *it is the virus that votes, not the voter*. The attacker could have the ability to control the outcome of the election.

3.3 Counterfeit Websites/Spoofing

Another risk to Internet voting involves counterfeit websites, or *spoofing*. Because fake websites can be made to look very much like legitimate ones, spoofing is used to entrap victims into revealing sensitive information such as credit card and social security numbers. In the case of Internet voting, spoofing could trick a voter into downloading malware on the voter's computer.

For example, suppose I want to defeat the Whig Party. I create a website called SupportWhigCandidates.com that looks like an official Whig Party website. I also post a notice on the website saying, "For more information, click here." Next, I send fake newsletters to the people who respond. These newsletters include an attached video about the Whig Party, accompanied by software for viewing the video. Unfortunately for the unwary voter, the software also contains a virus that emulates the official voting website. When voters who downloaded the video attempt to vote, they think they are connected to the official website, but in fact they are communicating only with malicious software that had been installed on their machines. If the voter selects a Whig candidate, the software can change the vote and send the modified vote to the official website. The voter will never know what happened.

A variation on spoofing is called *phishing*. Phishing involves email appearing to be from a legitimate organization, for example a credit card company. The phony email contains an authentic looking link that appears to go to the legitimate organization's website, but actually goes to the criminal website. Because the emails and websites tend to be well designed, the victim who goes to the fake website frequently ends up providing sensitive information, such as a credit card number. Regrettably, phishing has defrauded many people [Kerstein (2005)].

While phishing usually steals personal information, the same techniques could be used to undermine an Internet election. Someone wishing to defeat, say, Edmund Burke could send phishing emails that look like they are from the Burke campaign saying, "Click here to vote for Edmund Burke." The link would be to a fake website that looks exactly like the real one used for voting.

Yet another use of the fake voting website is to create a "man in the middle" attack, which is a form of active eavesdropping. In this case, private information gathered on the fake website, such as a password, is used to impersonate the voter. By posing as the real voter, the attacker has the capability of casting the legitimate voter's ballot for the Whig opposition.

A good way to avoid counterfeit websites is to rely on a trusted entity to authenticate that a website is legitimate. This job typically is done by the user's browser, which checks that the website has a valid "certificate" issued by a trusted entity known to the browser. The certificate digitally identifies the organization or individual with whom it is associated. If the browser does not trust the issuer of the certificate, it will ask the computer user if she still wants to access the questionable website. If a user does not understand the significance of what the browser is asking, she may naively instruct the browser to access the possibly counterfeit website.

But even if the voter is very careful and links only to what she believes are legitimate websites, she could still be victimized. First, it is possible to trick many browsers into going to the attacker's, rather than the legitimate, website [Zetter (2009)]. Second, the voter could be vulnerable to an attack on the routing infrastructure of the internet [Marsan (2009)]. Such an attack could take her to a fake voting website without her knowledge. This attack is a bit like stealing mail (in this case the voted ballot) out of the voter's mailbox.

3.4 Denial of Service Attacks

Another risk of Internet voting is a Denial of Service (DoS) attack that prevents people from accessing a critical website. A DoS attack occurs when so many corrupted computers attempt to access a website that legitimate users are unable to do so. A Distributed Denial of Service (DDoS) attack involves large numbers of computers that typically are controlled by some widely distributed malware. There are many well documented instances of DDoS attacks, such as the massive 2007 DDoS attack on Estonia, discussed below. There was speculation that both the Estonian attack and an attack on Georgia that occurred around the time of the Georgian invasion of South Ossetia originated in Russia. DDoS attacks have not been limited to countries; other victims have included Twitter, Facebook, Google, Yahoo, eBay, and Amazon.

A DDoS attack might be deployed during an Internet election to prevent certain groups from voting. A DDoS attack also could disrupt an entire election, as may have occurred in a leadership vote by the New Democratic Party (NDP) in Canada. Internet voting for the NDP election lasted from January 2 through convention day of January 25, 2003. On January 25, the same day that the Slammer worm was attacking large numbers of (unpatched) Windows 2000 servers on the Internet, the NDP voting site reportedly was down for several hours [CBC News (2003)]. The vendor claimed that the site was down for only half an hour. If, however, the site was so slow that for several hours people attempting to use it thought it was down, it was in effect down for those hours.

Because of the secrecy surrounding the technical aspects of the NDP election, we do not know if the NDP voting site was brought down by a denial of service attack or if it was a victim of the Slammer worm. The vendor, election.com, claimed to have patched the servers against the Slammer worm and maintained that the attack was a denial of service attack. However, election.com provided neither logs nor other proof that the servers were patched; nor did they permit expert examination of the records. There was no transparency, and no way for an independent outsider to determine what had happened.

3.5 The Server

Internet voting, whether web-based or via email, involves having a computer (client), typically that of the voter, communicate via the Internet with a receiving computer (server). Consequently, yet another risk of Internet voting is that the server could be manipulated by insiders or attacked over the Internet by individual hackers, political operatives, foreign governments, or even terrorists. Since national races, especially in the U.S., cost vast amounts of money, a very small fraction of which would be an exceedingly large bribe, there could be considerable financial, as well as ideological, incentives for individuals to perpetrate attacks. As attacks on large databases containing personal information demonstrate, insider and outsider attacks are serious threats [Vijayan (2006)].

3.6 No Post-Election Audit

Accountability is a critical and difficult aspect of any election. In order to address security concerns, we need to be able to verify after the election that the technology has operated correctly and that the right person has been declared the winner. Just as businesses routinely conduct audits, so too should elections be routinely audited, or, if necessary, recounted – as happened in the 2008 Minnesota Senatorial race. In order to conduct a post-election audit or recount, election officials must possess accurate copies of all legally cast ballots, something that is impossible with current Internet voting technology. Allowing the voter to print a copy of her ballot for her own use is meaningless, because that copy may not match the electronic version that reached the election official.

There are some proposals to use encryption to allow the voter to verify that an accurate version of her ballot has been received and counted. Unfortunately, it is very difficult for encryption to protect a voter whose computer has already been compromised by malware, since the malware can control the software without the voter's knowledge [Desmedt and Estehghari (2009)]. Encryption also does not protect against insider and Denial of Service attacks, spoofing, and many kinds of ordinary software bugs.

It is difficult to prove that election fraud has occurred in elections for which it is impossible to conduct a valid audit or recount. However, in July 2007 Computerworld reported that old-fashioned ballot box stuffing was used in an Internet election to fill five seats on the board of directors of the U.S. chapter of the IT Service Management Forum (ITSMF USA) [Thibodeau (2007)]. After receiving an anonymous tip that included the names of 15 people who were supposed to have voted, but didn't, ITSMF USA hired Kroll Inc. to investigate the election. Kroll initially was able to confirm 13 fraudulent votes out of about 500.

3.7 Societal Implications

Not all elections are the same. Private elections can be quite different from national, state, and local elections. Shareholder elections allow proxy voting, do not require a secret ballot, and do not adhere to the one person one vote principle. Usually there is little incentive to subvert a private election, although in recent years there have been some highly contested shareholder elections.

Just as a small bank in a rural town may not have as large a safe and as many armed guards as a major bank in a financial center, so too a relatively insecure Internet voting system might be adequate for some small scale relatively unimportant local elections. For example, Internet voting might be suitable for frequently held referenda in the canton of Geneva, Switzerland, as we discuss later, but grossly insecure for use in a U.S. presidential election.

The key point is this: *The larger the target, the more an attacker can benefit from a successful attack.* Because attackers are willing to put far more effort into attacks on large targets, small targets may be able to get away with relatively weak defenses – independent of the public or private nature of the election. Whether we are electing the board of directors or the mayor, the village of Setauket, N.Y. can run elections far more informally than New York City.

Because a large target is more attractive than a small target, so-called “pilot” studies of Internet voting are essentially meaningless. If I want to steal an election, I am highly unlikely to attack a small pilot voting project. Instead, I shall wait to launch an attack until policy makers and election officials have become convinced of the safety of the Internet voting system and deployed it widely.

A serious failure of Internet voting in a major election would have significant societal repercussions. Since our legal system is not designed to rerun compromised elections, an Internet voting attack detected during the election would create chaos. If there were no back-up voting scheme, the country could be thrown into disarray. If voters depended on the compromised system to vote, then even if the election were to be rerun, we would be confronted with holding another election on a system already shown to be insecure. It is not clear what would happen if the failure were detected after the winners had been declared.

4 Early Reports

Many of the risks we have discussed were already recognized in some early reports.

In 2000 and 2001 three reports on Internet voting, as well as a more general report issued by a commission co-chaired by former Presidents Carter and Ford, were released. All of the reports expressed significant reservations about voting in federal elections over the Internet. The security risks described in these reports remain as relevant today as they were when the reports were issued, as has been demonstrated by the serious problems caused by post-2001 viruses and worms such as Blaster, Slammer, MyDoom, Conficker, and Zeus.

The California Internet Voting Task Force Report was the first major report to address security issues of Internet voting [California Task Force (2000)]. The report, commissioned by the California Secretary of State and released in January 2000, warned of "the possibility of 'Virus' and 'Trojan Horse' software attacks on home and office computers used for voting." While the report optimistically stated that such attacks are preventable, it realistically observed that attacks "could result in a number of problems ranging from a denial of service to the submission of electronically altered ballots." In light of the many security problems, the report recommended against deploying a system that allows voting from home, office, or any Internet connected computer "until a satisfactory solution to the malicious code and remote control software problems is offered."

The Report of the National Workshop on Internet Voting (NWIV): Issues and Research Agenda was published in 2001 in response to a 1999 request to examine the possibility of Internet voting made by President Clinton to the National Science Foundation. Below is a quote from the Executive Summary [National Workshop (2001)]:

> *Remote Internet voting systems pose significant risk to the integrity of the voting process, and should not be fielded for use in public elections until substantial technical and social science issues are addressed.* [italics in original] The security risks associated with these systems are both numerous and pervasive, and, in many cases, cannot be resolved using even today's most sophisticated technology.

The Caltech/MIT Voting Technology Report: What Is, What Could Be, published in July 2001, stated: "Remote Internet voting poses serious security risks. It is much too easy for one individual to disrupt an entire election and commit large-scale fraud." [Caltech-MIT (2001)]

The National Commission on Federal Election Reform, co-chaired by former Presidents Carter and Ford, strongly recommended against Internet voting in its August 2001 report: [National Commission (2001)]

> Our concerns about early and remote voting plans are even stronger as we contemplate the possibility of Internet voting. In addition to the more general objections, the Commission has heard persuasive testimony that Internet voting brings a fresh set of technical and security dangers all its own. This is an idea whose time most certainly has not yet come.

5 Some Elections Held on the Internet

Internet voting began in earnest around 2000, with one of the earliest and largest election being international.

5.1 The ICANN Election

The most ambitious international Internet-based election was conducted by the Internet Corporation for Assigned Names and Numbers (ICANN) in 2000. Since the organization deals with worldwide issues of Internet governance, ICANN felt that some Board members should be elected via an Internet-based election.[2]

Anyone age 16 or older with an email address was eligible to vote. The world was divided into five regions, and a nominating committee selected candidates from each region. There was also a member nomination process. The voter registration stage was managed by ICANN; administration of the election was contracted out to election.com.

Problems with accessing the official website occurred in every phase of the election, starting with voter registration. According to analysis by the Carter Center: [Carter Center (2001)]

> The serious problems encountered occurred in the final days of the registration period. A strong surge of registrants, particularly from the Asia Pacific region, overwhelmed the ICANN server and caused a denial of service to people attempting to register.. . . [T]he legitimacy of any election can be questioned if the voter registration process grossly fails to enfranchise the great majority of persons wishing to vote.

In spite of problems, about 158,000 people completed the first stage of voter registration, though registration was unevenly distributed around the world. As the Carter Center observed, there may have been an attempt to augment the voter list:

> [I]n the European region the number of registering voters from Germany was much higher relatively than from other EU member countries . . . In the Asia Pacific region media coverage, apparently state-sponsored in large part or corporate-sponsored in one instance, generated a higher level of voter registrant interest in Japan, China, Korea

[2] The author has an intimate knowledge of this election, having been the runner-up for the North American seat won by Karl Auerbach.

> and Taiwan.... Further information ... indicated that a Japanese corporation and certain Japanese government agencies apparently did actually solicit registrations and votes on behalf of a Japanese candidate.

Because ICANN significantly underestimated the number of people wishing to participate in the election, there were insufficient computing resources. The problem was exacerbated because people tended to wait until the last minute to attempt to access the ICANN website. Each person who successfully registered was sent a unique Personal Identification Numbers (PINs) via the postal service. After receiving their PINs, voters were required to activate their membership, a process intended to reduce the risk of fraud. However, the activation requirement was confusing to many voters. Furthermore, some voters lost their PINs or never even received them – a problem in countries with unreliable mail systems.

Of the roughly 158,000 people who registered, only about 76,000 activated their membership. The number who voted in the election was only about 34,000. These numbers suggest large-scale disenfranchisement of voters at every stage of the process, with fewer than 1/4 of the initially registered voters actually voting.

Some people claimed to have voted multiple times. According to the Carter Center:

> Individuals intent on registering more than once using more than one email address would not find it too difficult to defeat the controls and beat the system. The possibility of doing the same on a large scale organized basis therefore also exists, introducing the risk of fraud capable of changing the electoral outcomes. Batch registration in the Asia region apparently occurred and raises questions about people registering for other people and voting on their behalf as well.

Many aspects of the ICANN election were far from transparent. Attempts to get ICANN to disclose the means by which they decided whether or not someone attempting to register was a legitimate voter were unsuccessful [Froomkin (2007)].

In addition, requests for information about the hardware and software used for the election were ignored. Ted Byfield, who has written extensively about ICANN, describes his attempts to obtain details about the software and hardware used for the signup process [Byfield (2000)]:

> [I] pressed ICANN to release information about the systems supporting the ... signup process ... whose failures were widely noted .... In light of those failures, ... I requested: (1) “the hardware configuration

> of the server(s) . . . "; (2) info on "who or what company wrote the software . . . "; (3) "documents associated with the specification of the hardware/software configuration . . . "; and (4) a statement as to whether "the implementation . . . [was] subject to an open and/or competitive bid." In its inimitable style, ICANN hasn't refused these requests: instead, after much hand-waving, *it has refused to refuse them* [italics in original].

The Carter Center report stated that "the technical weakness in the registration system made it virtually impossible to assess the integrity of the voters' list, the security of the PINs, and secrecy of vote." It also observed that technical problems during the first day of voting, which ran from October 1 – 10, 2000, created a "credibility problem."

5.2 The Secure Electronic Registration and Voting Experiment (SERVE)

Despite the warning in a report it had commissioned earlier, FVAP (the Federal Voting Assistance Program) allocated $22 million to build the SERVE system so that U.S. military and civilians living abroad could vote over the Internet in the 2004 primaries and general election. SERVE had ambitious goals. It was intended to allow voters covered by FVAP to register and vote over the Internet in the 2004 primaries and in the general election. Participation by states and counties within those states was voluntary.

SERVE required the voter to have Windows 95 or above; Microsoft Explorer 5.5 or above, or Netscape Navigator 6.x through 7; an Internet connection; and the ability to download an ActiveX component.[3] Voters could use their own computers, so long as they ran Windows, or they could use public computers, such as those found in libraries and cybercafes, in any country. Voters were responsible for the security of whatever computers they used. The vendor for SERVE was Accenture.

In 2003, FVAP assembled a group of experts called the Security Peer Review Group (SPRG) to evaluate SERVE. Following two three-day meetings with FVAP officials and lead technical staff of SERVE, the four computer scientists who attended both meetings of SPRG released *A Security Analysis of the Secure Electronic Registration and Voting Experiment (SERVE)*, on January 20, 2004 [Jefferson, Rubin, Simons, and Wagner (2004)]. The conclusion stated:

[3] ActiveX is Microsoft technology used for developing software components that can be automatically downloaded and executed by a web browser. Since Netscape did not support ActiveX, Netscape users would run a Java applet that performed the same tasks as the ActiveX component.

> Because the danger of successful, large-scale attacks is so great, we reluctantly recommend shutting down the development of SERVE immediately and not attempting anything like it in the future until both the Internet and the world's home computer infrastructure have been fundamentally redesigned, or some other unforeseen security breakthroughs appear.

At the time the security analysis report was issued, 50 counties in seven states were planning to participate in SERVE. FVAP had estimated that the maximum overall vote total would be approximately 100,000, including both primaries and the general election.

On January 30, then Deputy Defense Secretary Paul Wolfowitz issued a memo stating that the Pentagon "will not be using the SERVE Internet voting project in view of the inability to assure legitimacy of votes that would be cast using the system, which thereby brings into doubt the integrity of election results" [Weiss (2004)]. The Wolfowitz memo also said that the department "may continue efforts to demonstrate the technical ability to cast ballots through the use of electronic voting systems" [Department of Defense (2007)]. On February 5, news that SERVE would not be used in the 2004 election became public [FVAP (2004)]. SERVE was subsequently terminated.

FVAP has referred to the authors of the security report as "a minority of members of a peer review group for SERVE"; however, there was neither another peer review group nor a "majority" report. [FVAP (2006)]

While the goals of SERVE were laudable, and the team working on SERVE dedicated, it is fortunate for our democracy that SERVE was halted. If SERVE had been deemed successful in 2004, it is likely that a similar program would have been made available first to the military and then to the entire country – despite the impossibility of proving the absence of serious software bugs or successful election rigging attacks. And even if somehow we could have known with certainty that SERVE had not been attacked in 2004, there is no guarantee that a similar system would not be attacked in a future election, especially since an attacker could have far more impact if many more people were voting over the Internet. Nonetheless, Internet voting for the military is an idea that is being actively pursued in the U.S.

5.3 The UK - Swindon

The UK conducted its first Internet voting pilot in 2002, when the town of Swindon provided voters with the option of voting over the Internet in a local election [BBC (2007)]. Swindon subsequently offered Internet voting for local elections in 2003 and again in May, 2007. Four other councils provided an Internet voting option in the 2007 election, but given its history, we shall focus on the Swindon election.

The 2007 Swindon pilot election allowed people to vote via telephone or over the Internet. (Anyone wishing a paper ballot had to vote at her assigned polling station). The Internet voting could be remote, via electronic voting from five locations prior to polling day, or from one of 64 electronic polling stations that provided a "vote anywhere" environment only on polling day. The Internet voting software was provided by Everyone Counts.

The entire Swindon pilot election was analyzed by the Electoral Commission, an independent body established by the UK Parliament [The Electoral Commission (2007)]. They concluded that:

- On the whole, the pilot scheme facilitated and encouraged voting.
- The pilot scheme did not facilitate the counting of votes.
- The pilot scheme had a negligible effect on the turnout of voters.
- On the whole, voters found procedures easy to follow.
- The pilot scheme does not appear to have led to any increase in . . . offences or malpractice.
- The pilot scheme led to an overall increase in expenditure by the Council. . . . The average cost of the 2007 pilot scheme per elector was £8.33, compared with £2.30 for a conventional election, while the cost per electronic vote cast was £102.50.

In other words, each Internet vote was over 44 times more costly than a conventional vote.

The Electoral Commission also found serious usability problems, including non-intuitive screen layouts that could cause the elector to "very easily miss additional unseen candidates at the bottom of the ballot paper," a user interface that "did not include high-contract, clearly labelled buttons" (thereby increasing the chance for error), and a voter interface that "seemed to contain several more steps than the equivalent manual voting process." As a result of the additional steps, electronic voting "took an average of three minutes compared with 30 seconds for manual voting."

Someone choosing to vote remotely, either via the Internet or phone, was sent a registration form on which she provided her date of birth and a unique self-generated personal identification number (PIN). The voter then signed, dated, and returned the registration form. Once the registration process had been completed, the voter was sent a 10-digit ballot code in a secure mailer that also served as the official poll card. When voting by Internet or telephone, voters had to provide their birth dates, PINs, and ballot codes. The votes were encrypted prior to transmission.

After voting over the Internet remotely or at a supervised location, voters were given the option of creating an encrypted receipt that could be used to verify that their votes had been accurately received and recorded. Over half of the ballots were voted over the Internet.

5.3.1 ORG's observations about Swindon

The 2007 elections were the first in the UK to permit accredited observers. The non-profit Open Rights Group (ORG) organized teams of volunteers to oversee several of the pilot projects, including Swindon. Their *May 2007 Election Report* expressed serious concerns about the various pilots [ORG (2007)]:

- The accredited observers encountered difficulties with the observation process and at times had to deal with arbitrary decisions and limited access.
- There was no independent or governmental analysis of any of the software.
- No manual audit of electronically counted ballots was conducted, though obviously such an audit could not be conducted where Internet voting occurred.
- Audit trails were inadequate.

ORG also noted that "the supplier concerned confirmed that only 'about 100 people [in Swindon] had used the service' to check their receipts" [ORG (2007)].

Security of the laptops used for voting in Swindon seemed lax at best. Presiding officers were allowed to take home laptops used for voting prior to the election. There also were a number of unsecured laptops at the voting locations. Observers were told that those laptops would not be used to tabulate votes, but they were not told how the laptops would be used. In addition, according to ORG, a majority of the polling stations experienced problems with the laptops that were used for e-voting.

An ORG volunteer visited a polling station, after having received conflicting reports about what appeared to be malfunctioning (Internet-linked) voting equipment. There she met Lori Steele, CEO of Everyone Counts. When asked if the laptops were experiencing problems, such as had been observed elsewhere in Swindon, Steele responded that the machines were "up and running." However, after Steele left, the observer asked the same question of the Presiding Officer, who responded that there were "technical glitches." When pressed about Steele's comment, the Presiding Officer said that she was being "diplomatic" [ORG07].

5.4 Geneva

Because a large number of its citizens live outside the country, in 2001 Geneva initiated its eVoting project as a third option in addition to polling station and postal voting. The goal was to develop an Internet voting system that would be as secure as postal voting, which is to say only somewhat secure.

Geneva mandated that the software, which it owns, not contain any secret code. Vendors were informed that "it must be possible for specialists external to the State of Geneva ... to examine thoroughly all the software dealing with the ballots" [Republique de Geneve (2007)]. In 2003 it was estimated that the development costs would ultimately be about two million Swiss francs ($1.35 million at the time) [Hensler (2003)].

Before each election the voter is sent a *voting card* via postal mail. The card contains an identification number, which is not hidden, and a key or PIN, concealed under metallic paint that must be scratched off in order to vote on the Internet.

A voter choosing to vote over the Internet first logs on to the voting website and enters the identification number from the card. The website uses that information to verify that the voter has not yet voted. The voter then enters her choice on an electronic ballot, and the system returns the voter's selection for the voter to modify or verify. (Currently, Internet voting is being used only for initiatives and referenda; typically there is only a single item on the ballot.) The voter authenticates the choice by providing some personal information and the PIN from the voting card. Because the voter has a unique ID, there is a risk that the voter's choice might become known. The prohibition against secret software provides some protection against this risk.

To protect the voter from mistakenly accessing counterfeit websites, the system returns a "verification code" that should match the code that has been included on the voting card. Voters who do not receive the correct verification code are instructed to notify an election official. Presumably, notification would trigger a new election, though it is not clear what would happen if only a very small number of people claimed not to have received a correct verification code. There is also the risk that some voters may not bother checking the verification code or might be fooled by a spurious error message.

Paper is intrinsic to the system; only the link from the voter to the election website is electronic. Paper is used to distribute the voter's ID number, PIN, and verification code.

The Geneva system does not provide a good defense against an attack by malware that has been installed surreptitiously on the user's machine. The malware could intercept all communications between the voter and the voting website and modify the voter's selection. The voter, who still would receive the correct verification code, would not know that her vote was changed.

While the Geneva system attempts to protect against a denial of service attack, there is the recognition that this might not be possible. In the event that a denial of service attack were to make Internet voting difficult or impossible, Internet voting would be halted, and voters would be instructed to vote at polling stations. Citizens residing in Geneva would have little trouble going to polling stations to vote, but those living outside the country who had not yet cast their votes could be disenfranchised.

In February 2009, Geneva voters ratified a constitutional amendment guaranteeing Internet voting in Geneva [Republique de Geneve (2009)].

5.5 The Baltic States

5.5.1 Estonia

The stated goal of the Estonian system is to make Internet voting "at least as secure as regular voting" [National Election Committee (2005)]. In 2007 Estonia became the first country to hold national parliamentary elections that provided the option of voting over the Internet to all citizens. Internet voting was possible from the morning of February 26 to the evening of February 28, with polling-place voting occurring on March 4 [Reuters (2007a)]. Repeat voting was allowed, including polling place voting on Election Day, with the understanding that the vote with the latest time stamp would be the official vote [Valimiskomisjon (2009)].

Most Estonians had been issued ID cards that were also smart cards. In addition to the national ID cards and computers with smart card readers, the Estonian system requires the existence of a kind of technological infrastructure (public key), which Estonia has. During the election the ID card was used both for authentication and to allow the voter to sign her ballot digitally. Over 30,000 people, or roughly 3.5% of registered voters, voted online [Borland (2007)].

After connecting to the Network Server, the voter authenticated herself by using her ID card. The system then provided the voter with the candidate list on the website. After making her selections and finalizing her ballot, the voter's computer generated a random number r, which was concatenated to the ballot. The purpose of the random number was to prevent an adversary from using exhaustive search to determine the contents of encrypted ballots. The system next encrypted and digitally signed the concatenated ballot, which was sent to the Network Server. The Network server verified the voter and stored the encrypted ballot. Before the ballots were counted, the voter's signature was stripped from the ballot. The system was developed by Cybernetica, an Estonian company.

Unfortunately, the system was far from providing the same level of security as regular voting. The only known protection against an insider attack is the ability to conduct a manual count of the paper ballots or records that represent the voters' choices. That capability is clearly lacking in the Estonian system. Hence, as with other paperless Internet voting systems, threats from insider attacks and malicious computer viruses remain.

The system is also vulnerable to denial of service attacks. Indeed, a massive denial of service attack against web sites in Estonia began on April 27, 2007 and continued through part of May [Kirk (2007)]. Because Estonia had become so dependent on Internet communications, the attack created serious problems for governmental and private institutions. There was speculation that the attack, which originated in Russia, was triggered by Russian anger over Estonia's decision to move a Soviet war memorial [Finn (2007)]. While voters would have had the option of polling place voting in the event that a successful denial of service attack had occurred during the election, voters who had planned to vote over the Internet could have been disenfranchised.

The 2007 attack notwithstanding, Estonia continues to allow Internet voting, most recently for the June 7, 2009 European Parliament elections and local government elections a few months later [Valimiskomisjon (2009)]. On-line voting was available from 9 a.m. May 28 until 8 p.m. June 3. The computing requirements are an Internet connection; Windows, Mac OS X, or Linux operating system; ID-card software; and a card reader.

The Estonian system depends on the voter's computer to be free of election-rigging malware. In spite of the limitations of anti-virus software in identifying new unknown viruses, voters were told to [Valimiskomisjon (2009)] ;

> Make sure that your computer is virus-free! If no anti-virus software has been installed into your computer please click here to install it free of charge.

5.5.2 Lithuania

Estonia is not the only Baltic state that continues to believe in Internet voting, in spite of the cyberattack it experienced. Nearby Lithuania wanted to prove that it too could conduct an Internet election. In fact, the Lithuanian Prime Minister, Gediminas Kirkilas, told a news conference, that "I hope that in this area we will catch up with Estonia, and by doing this we will surpass most European Union states" [Reuters (2007b)]. In order to catch up with Estonia, Lithuania spent about 580,000 euros to implement Internet voting for their 2008 elections. However, in January

2008 the legislation that would have enabled Internet voting was defeated in the Lithuanian Parliament [Heise online (2008)].

Lithuania did follow in Estonia's path, but not in a way that they had wished. In July 2008 Lithuania was subjected to a massive cyberattack in which Soviet symbols was posted on corporate and government websites [Rhodin (2008)].

6 The Future of Internet Voting

Proposals to conduct voting pilots or experiments using real elections continue to reappear both in the U.S. and elsewhere, seemingly independent of warnings from computer security experts. While the appeal of Internet voting is obvious, the risks, unfortunately, are not, at least to many decision makers.

It is dangerous to draw conclusion from what appears to be a successful Internet voting experiment. If the election is insignificant, there is little to no motivation to subvert the election. Even if the election is significant, a malicious player might not attack the first election in the hope that the election will be declared a success, thereby resulting in future Internet elections. Having claimed success, independent of any verifiable proof of the accuracy of the election, Internet voting enthusiasts could push to extend Internet voting to a broader group of voters, thereby seriously undermining the security of our electoral system.

An attacker who does launch an attack is not going to publicize it – just as the credit card thieves who break into U.S. and European banks don't advertise that they are stealing credit card numbers. When election officials or policy makers ask for proof that voting systems have been attacked, it is important to keep in mind that well devised attacks by their very nature can be difficult or impossible to detect. Furthermore, most Internet voting systems do not include a proactive search for compromise before or after an election, thereby making the attacker's job all the easier.

In spite of the well documented risks of Internet voting, many governments cling to the misguided notion that an Internet voting pilot can prove that Internet voting is secure. Allowing voters to return voted ballot via electronic transmission (fax, email, and web-based) is being or has been considered by several countries, including the U.S., as well as by a disconcertingly large number of states.

Perhaps some day a paperless encryption-based Internet voting will be secure and accurate. However, as stated by the National Commission on Federal Election Reform, co-chaired by Presidents Carter and Ford, Internet voting "is an idea whose time most certainly has not yet come" [National Commission (2001)].

References

BBC (2007): "Swindon's enthusiasm for e-voting," *BBC News*, http://news.bbc.co.uk/2/hi/uk_news/england/wiltshire/6608809.stm.

Bonifaz, J. (2009): "Washington State Internet Voting Bill Defeated," *Voter Action*, http://www.voteraction.org/print/807.

Borland, J. (2007): "Online Voting Clicks in Estonia," *Wired*, http://www.wired.com/politics/security/news/2007/03/72846?currentPage=all.

Byfield, T. (2000): "ICANN: transparency through obscurity," The Roving Reporter Blog, http://www.tbtf.com/roving_reporter/icann2.html#10.

California Task Force (2000): "A Report on the Feasibility of Internet Voting," http://www.sos.ca.gov/elections/ivote/final_report.pdf.

Caltech-MIT (2001): "The Caltech/MIT Voting Technology Report: What Is, What Could Be," Technical report, Caltech-MIT Voting Technology Project, http://web.archive.org/web/20070219091734/http://www.vote.caltech.edu/reports /2001report.htm.

Carter Center (2001): "Report on the Global, On-Line, Direct Elections for Five Seats Representing At-Large Members on the Board of Directors of the Internet Corporation for Assigned Names and Numbers (ICANN)," http://www.markle.org/downloadable_assets/icann2_report.pdf.

CBC News (2003): "Computer Vandal Delays Leadership Vote," CBC News, http://www.cbc.ca/news/story/2003/01/25/ndp_delay030125.html.

Chaney, T. (2000): "Using the Internet to Improve Voter Turnout (Netocracy)," http://www.spark-online.com/may00/esociety/tyson_chaney.html.

Computer Technologists (2008): "Computer Technologists' Statement on Internet Voting," http://www.verifiedvoting.org/article.php?id=5867.

DeGregorio, P. (2009): "UOCAVA Voting Scoping Strategy," obtained under a public record request by John Gideon, http://www.votersunite.org/info/WA-PRR-ScopingStrategy.pdf.

Department of Defense (2007): "Department of Defense: Expanding the Use of Electronic Voting Technology for UOCAVA Citizens - As Required by Section 596 of the National Defense Authorization Act for Fiscal Year 2007," http://servesecurityreport.org/DoDMay2007.pdf.

Desmedt, Y. and S. Estehghari (2009): "Hacking Helios and its Impact," Crypto 2009, rump2009.cr.yp.to/1b884ce772d84af05f0f4b07bf019053.pdf.

Dunn, J. E. (2010): "Trojan attacks credit cards of 15 US banks ," *Techworld*, http://news.techworld.com/security/3232010/trojan-attacks-credit-cards-of-15-us-banks/.

FBI (2010): "Press Release: FBI, Slovenian and Spanish Police Arrest Mariposa Botnet Creator, Operators," http://www.fbi.gov/pressrel/pressrel10/mariposa072810.htm.

Finn, P. (2007): "Cyber Assaults on Estonia Typify a New Battle Tactic," *The Washington Post*, http://www.washingtonpost.com/wp-dyn/content/article/2007/05/18/AR2007051802122.html.

Froomkin, M. (2007): Email to the Author.

FVAP (2004): "e-Voting Initiatives," http://web.archive.org/web/20040606205521/ http://fvap.gov/services/evoting.html.

FVAP (2006): "Report on IVAS 2006 As Required by Section 596 of the National Defense Authorization Act for Fiscal Year 2007," http://www.votetrustusa.org/pdfs/2006 IVAS Assessment Report Dec - Final.pdf.

H Security (2007): "Antivirus protection worse than a year ago," http://www.h-online.com/security/news/item/Antivirus-protection-worse-than-a-year-ago-735697.html.

Hayden, M. (2010): "Hayden: Hackers Force Internet Users to Learn Self-Defense," *PBS Newshour*, http://www.pbs.org/newshour/bb/science/july-dec10/cyber_08-11.html.

Head, W. (2009): "Hackers use Wikipedia to spread malware," SC Magazine, http://www.securecomputing.net.au/News/67796,hackers-use-wikipedia-to-spread-malware.aspx.

Heise online (2008): "Controversy around the introduction of internet voting in Lithuania," http://www.heise.de/english/newsticker/news/102124.

Hensler, R. (2003): "The Geneva Internet Voting System," Republique et Canton de Geneve, www.mailclad.com/articles/pre_projet_eVoting_eng.pdf.

Jefferson, D., A. D. Rubin, B. Simons, and D. Wagner (2004): "A Security Analysis of the Secure Electronic Registration and Voting Experiment (SERVE)," http://news.findlaw.com/ cnn/docs/voting/nsfe-voterprt.pdf.

Kanan, K., J. Rees, and E. Spafford (2009): "Unsecured Economies: Protecting Vital Information," *McAfee, Inc.*, http://www.cerias.purdue.edu/assets/pdf/mfe_unsec_econ_pr_rpt_fnl_online_012109.pdf.

Kerstein, P. L. (2005): "How Can We Stop Phishing and Pharming Scams," *CSO Magazine*, http://web.archive.org/web/20080117092752/http://www.csoonline.com/talkback /071905.html.

Kirk, J. (2007): "Estonia recovers from massive denial-of-service attack," *InfoWorld*, http://www.infoworld.com/d/security-central/estonia-recovers-massive-denial-service-attack-188.

KITV (2009): "Voting Drops 83 Percent In All-Digital Election," http://www.kitv.com/politics/19573770/detail.html.

Kurtz, G. (2010): “Operation “Aurora” Hit Google, Others,” McAfee Security Insights Blog, http://siblog.mcafee.com/cto/operation-%E2%80%9Caurora%E2%80%9D-hit-google-others/.

Lohr, S. (2010): “Companies Fight Endless War Against Computer Attacks,” *The New York Times*, http://www.nytimes.com/2010/01/18/technology/internet/18defend.html?pagewanted=all.

M86 Security Labs (2010a): “Cybercriminals Target Online Banking Customers,” http://www.m86security.com/documents/pdfs/security_labs/cybercriminals_target_online_banking.pdf.

M86 Security Labs (2010b): “M86 Security Labs Discovers Customers of Global Financial Institution Hit by Cybercrime,” http://www.m86security.com/i/M86-Security-Labs-Discovers-Customers-of-Global-Financial-Institution-Hit-by-Cybercrime,news.1430~.asp.

Marsan, C. D. (2009): “Feds to Shore Up Net Security,” *About.com*, http://pcworld.about.com/od/securit1/Feds-to-Shore-Up-Net-Security.htm.

McCartney (2009): “McCartney site serves up Zeus malware,” Internet and Network Security, http://www.infosecurity-us.com/view/1178/mccartney-site-serves-up-zeus-malware/.

McMillan, R. (2010): “FBI Director: Hackers have corrupted valuable data,” *IDG News Service*, http://www.macworld.com/article/146904/2010/03/hackers.html?lsrc=rss_main.

Messmer, E. (2009): “America’s 10 Most Wanted Botnets,” *Network World*, http://www.networkworld.com/news/2009/072209-botnets.html.

Mills, E. (2009): “Spam offers to let people use their PC to attack Obama site,” CNET News, http://news.cnet.com/8301-27080_3-20013246-245.html.

National Commission (2001): “To Assure Pride and Confidence in the Electoral Process,” http://www.reformelections.org/ncfer.asp#finalreport.

National Election Committee (2005): “E-Voting System Overview,” http://www.vvk.ee/elektr/docs/Yldkirjeldus-eng.pdf.

National Workshop (2001): “The Report of the National Workshop on Internet Voting: Issues and Research Agenda,” Technical report, Internet Policy Institute, http://news.findlaw.com/cnn/docs/voting/nsfe-voterprt.pdf.

Oberheide, J., E. Cooke, and F. Jahanian (2008): “CloudAV: N-Version Antivirus in the Network Cloud,” Technical report, 17th USENIX Security Symposium, http://www.usenix.org/events/sec08/tech/full_papers/oberheide/oberheide.pdf.

ORG (2007): “May 2007 Election Report: Findings of the Open Rights Group Election Observation Mission in Scotland and England,” Technical report, The Open Rights Group, http://www.openrightsgroup.org/wp-content/uploads/org_election_report.pdf.

Republique de Geneve (2007): "The Geneva Internet Voting System," http://www.geneve.ch/evoting/english/presentation_projet.asp.

Republique de Geneve (2009): "E-Voting - Internet voting in Geneva, Frequently asked Questions (FAQ)," http://www.geneve.ch/evoting/english/faq-internet-voting.asp.

Reuters (2007a): "Estonia to hold first national Internet election," http://news.cnet.com/2100-1028_3-6161005.html.

Reuters (2007b): "Lithuania plans voting via Internet," http://www.reuters.com/article/technologyNews/idUSL1136307020070711.

Rhodin, S. (2008): "Hackers Tag Lithuanian Web Sites with Soviet Symbols," *The New York Times*.

Symantec Corporation (2009): "The Conficker Worm," http://www.symantec.com/norton/theme.jsp?themeid=conficker_worm#do.

The Electoral Commission (2007): "Electoral pilot scheme evaluation Swindon Borough Council," Technical report, The Electoral Commission.

Thibodeau, P. (2007): "ITIL group gets clear evidence of voting fraud in online election," *Computerworld*.

Valimiskomisjon, V. (2009): "Elections and E-Voting," http://www.valimised.ee/teema_eng.html.

Vascellaro, J. E. and J. Solomon (2010): "Yahoo Was Also Targeted in Hacker Attack," *The Wall Street Journal*.

Vijayan, J. (2006): "Breach at UCLA exposes data on 800,000," *Computerworld*.

Weiss, T. R. (2004): "Pentagon drops online votes for armed forces," *ComputerWeekly.com*, http://www.computerweekly.com/Articles/2004/02/06/200087/pentagon-drops-online-votes-for-armed-forces.htm.

Wikipedia (2010): "Conficker," http://en.wikipedia.org/wiki/Conficker.

Zetter, K. (2009): "Vulnerabilities Allow Attacker to Impersonate Any Website," *Wired.com*, http://www.wired.com/threatlevel/2009/07/kaminsky/.

Zeus09 (2009): "Measuring the in-the-wild effectiveness of Antivirus against Zeus," Trusteer, http://www.trusteer.com/files/Zeus_and_Antivirus.pdf.