The Business of Electronic Voting

Ed Gerck[1], C. Andrew Neff[2], Ronald L. Rivest[3],
Aviel D. Rubin[4], and Moti Yung[5]

[1] Safevote.com, egerck@safevote.com
[2] VoteHere.net, aneff@votehere.net
[3] Laboratory for Computer Science
Massachusetts Institute of Technology, Cambridge, MA 02139
rivest@mit.edu
[4] AT&T Laboratories – Research
180 Park Avenue, Florham Park, NJ 07932
rubin@research.att.com
[5] CertCo Inc., New York, NY
moti@cs.columbia.edu, moti@certco.com

Abstract. This work reports on a Financial Cryptography 2001 panel where we concentrated on the emerging business of electronic voting.

1 Preliminaries

The problems associated with traditional voting machines in a national election, their unreliability, inaccuracy and other potential hazards, were put in the public limelight, after the last USA presidential election (especially in the state of Florida). At the same time, but less conspicuously, an industry centered around electronic voting (national, boardroom, company wide, and otherwise) has started to emerge, offering various solutions. Therefore, it seems to be an emerging area where cryptography is crucial to industrial progress, which, in turn, makes it a proper subject within the area of "financial cryptography."

Indeed, for about 20 years, the cryptographic research community has dealt with issues related to security and robustness of e-voting as a fundamental protocol problem. Numerous election protocols with many provable properties have been designed and some have been prototyped as well. This research developed insight, and some of its results will surely influence future systems. However, here we concentrate on issues regarding to "real life" aspects of actual implementations of voting systems.

Obviously, the e-voting problem possesses some of the integrity and secrecy issues that underly many protocol problems in the area of financial transactions. Yet, the problem has certain requirements and characteristics which are unique and perhaps harder to achieve.

The discussion in this report includes the panelists' views of basic requirements and problem specifications, their views of major challenges in the field, their opinions regarding technical and social feasibility and their approaches to

P. Syverson (Ed.): FC 2001, LNCS 2339, pp. 243–268, 2002.

possible solutions. Then the notion of building "businesses" around electronic election is discussed as well. The basic issues are centered around technology, yet legal, social, financial, market and policy issues play important roles in investigating the reality of electronic voting business. The report consists of this preliminary section and a summary section by the moderator (M. Yung) and sections contributed by each of the panelists.

The original issues and problems which the panelists were directed to were: (1) reliability, (2) fairness, (3) scalability (does one solution fit all situations), (4) how much security is actually required? (5) is e-voting for real? (6) how far are we from "real" voting? (namely, is the technology ready?) (7) is the Internet the arena for voting? (8) is there interaction between the technology and its quality and the business success? (9) is it a business at all (namely, is there money to be made and how?) (10) what are e-voting's social and legal implications?

The rest of the paper includes the panelists' sections. Each panelist is fully responsible for his own contribution and each contribution has its own personal characteristics and its own level of optimism. The contributions reflect faithfully the positions expressed in the panel discussion. In section 2 Ron Rivest presents perspectives on electronic voting where he reviews some of the history of voting machines, some of the basic problems with systems and his personal views of the subject. In section 3 Andy Neff presents the difference between general e-voting and Internet-voting, he then presents basic requirements and practical considerations and challenges. In section 4 Avi Rubin considers the feasibility of remote voting, especially when taking into account the current state of the art of platform and Internet security (or, more accurately, insecurity). In section 5 Ed Gerck presents a general background related to the notion of trust and to secure and trustworthy election systems; he then reviews basic requirements for e-voting scheme. The summary then tries to report some of the discussion which followed the presentations by the panelists.

2 Ronald L. Rivest: Perspective on Electronic Voting

2.1 Introduction

Over the years, with varying degrees of success, inventors have repeatedly tried to adapt the latest technology to the cause of improved voting.

For example, on June 1, 1869 Thomas A. Edison received U.S. Patent 90,646 for an "Electric Vote-Recorder" intended for use in Congress. It was never adopted because it was allegedly "too fast" for the members of Congress.

Yet it is clear that we have not reached perfection in voting technology, as evidenced by Florida's "butterfly ballots" and "dimpled chads."

Stimulated by Florida's election problems, the California Institute of Technology and MIT have begun a joint study of voting technologies [CTM00], with the dual objectives of analyzing technologies currently in use and suggesting improvements. This study, funded by the Carnegie Foundation, complements

the Carter/Ford commission [FER01], which is focusing on political rather than technological issues. Electronic voting will be studied.

Among people considering electronic voting systems for the first time, the following two questions seem to be the most common:

Could I get a receipt telling me how I voted?

Could the U.S. Presidential elections be held on the Internet?

The first question is perhaps most easily answered (in the negative), by pointing out that receipts would enable vote-buying and voter coercion: party X would pay $20 to every voter that could show a receipt of having voted for party X's candidate. Designated-verifier receipts, however, where the voter is the only designated verifier—that is, the only one who can authenticate the receipt as valid—would provide an interesting alternative approach to receipts that avoids the vote-buying and coercion problem. See [JSI96] for a discussion of this idea.

The second question—can we vote remotely over the Internet– is more problematic.

We start by noting that "electronic voting" includes a wide range of possible implementations. The California Internet Voting Task Force [Ca00] distinguished between (a) voting at a supervised poll-site using electronic equipment, (b) voting at an unsupervised electronic kiosk (say, in a shopping mall), and (c) "remote voting"— voting from home or business using the voter's equipment.

Before proceeding to comment on the security of electronic voting systems, we should at least pause to consider the desirability of such systems. Is remote electronic voting over the Internet desirable? Why bother?

"Because we can" and "for increased voter convenience" are arguably insufficient justifications for electronic voting. "For increased confidence in the result" might be acceptable, if a convincing case could be made. Political scientists claim that the best justification is "to increase voter turnout."

In the remainder of this note, I discuss the "secure platform problem" as a key impediment to remote voting, and then provide a list of personal opinions regarding the security of electronic voting systems.

2.2 The "Secure Platform Problem"

There is a fundamental problem we must face when trying to design remote electronic voting systems: the "secure platform problem."

Cryptography is not the problem. Indeed, many wonderful cryptographic voting protocols have been proposed; see [Gr00] for a sample bibliography.

The problem is interfacing the voter to the cryptography.

Almost all proposed cryptographic voting protocols *assume* that a voter (e.g. Alice) has a secure computing platform that will faithfully execute her portion of the protocol. The platform (e.g. a PC) will correctly display to Alice her intended vote, and cryptographically submit her vote during the protocol. The platform acts as Alice's "trusted agent" during the voting protocol.

In essence, the platform is *Alice as far as the voting protocol is concerned.*

In reality, the current generation of personal computers running Windows or Unix are not sufficiently secure to act as trusted voting agents. These operating systems and their applications are far too vulnerable to viruses and Trojan horses. A hacker could easily write a virus that would cause Alice's computer to display her voting for one candidate while actually voting for another. If thousands of PC's are similarly infected, an election could be rigged. This is an unacceptable risk.

Other studies and reports have reached similar conclusions that current technology is not secure enough to support electronic voting from home. In particular, I note the report of the California Task Force on Electronic Voting [Ca00], Avi Rubin's note [R00], and the Internet Policy Institute Report on Internet Voting[IPI00].

Of course, the secure platform problem is not the only significant security problem that needs to be addressed regarding the possibility of electronic voting from home over the Internet. The Internet itself, while remarkably useful and reasonably robust, is all too vulnerable to flooding and denial of service attacks. The possibility that a foreign power could bring down the Internet on U.S. election day is all too real. For this reason alone, remote electronic voting from home over the Internet would be at best an available alternative, and it would be reasonable to expect existing poll-site voting systems to be prepared to handle everyone should the Internet be taken down.

2.3 Some Personal Opinions

E-Voting Is not Like E-Commerce. Electronic voting is unlike electronic commerce in several important ways, so it is insufficient to argue that secure electronic voting is merely a corollary to secure electronic commerce and that the same security mechanisms should apply.

For example, in electronic commerce there is always time to dispute a transaction if something hasn't worked correctly. With voting, there is a deadline that must be met.

Also, in an electronic commerce transaction, the buyer typically gets a receipt that can be used later to resolve disputes. In contrast, it is important, as noted earlier, that voters do *not* get receipts showing how they voted, since this may enable the voter to sell his vote.

In electronic commerce, transaction records identify the parties involved. In electronic voting, the ballots cast should *not* identify the voters who cast them, as this might violate the voter's privacy and subject them to coercion. (For example, if election officials could see how each voter voted, then the lead election official could see how his employees voted.)

It is more important that no one "has their thumb on the scale" than having a scale that is easy to use or even very accurate. The primary purpose of a voting system is to correctly determine the will of the voters. Given human nature, the likelihood of getting an incorrect result is much higher if

there are significant security vulnerabilities than if the vote-counting is a bit inaccurate. Fraud can be a problem in any election; counting errors affect only close elections. Ease of use is relevant only inasmuch as it affects voter turnout or introduces systematic biases.

Electronic voting from home runs the risk of allowing an adversary to put a "big thumb" on the scale, since the adversary may be able to automate his attack. For example, he could bring down the Internet in Democratic neighborhoods, or create a virus that affects computers with certain characteristics (e.g. those with ".edu" suffix). Such risks threaten the primary purpose of the voting system, and suggest exceptional caution in moving forward with such systems.

The voting system must be simple to understand and operate. Electronic voting systems are often complex. Voting systems must be certified before they are used. Election officials must have confidence that the voting system will prevent fraud and perform reliably.

Complexity is the enemy of security. Complex systems are difficult to understand and debug. Asking an election official to certify that thousands of lines of code provide a secure and trust-worthy election system is an entirely different matter than asking him to certify a set of procedures for managing a collection of paper ballots. Electronic voting systems place a substantial burden on the election officials who must certify the systems, and may weaken the credibility of the entire process in the voters' minds.

Even with poll-site electronic voting, the complexity of electronic voting systems may also challenge the election officials (who are often volunteers) who must install and operate the election equipment. The failure to educate both election officials and voters to use new equipment properly is a major source of election problems.

Physical ballots can provide better audit trails than purely electronic systems. The integrity and trust-worthiness of a voting system is greatly enhanced by having an audit trail recording each ballot cast. Many states require voting systems to have such audit trails.

Audit trails with very high integrity can be obtained when the audit trail is created directly by the voter, as with a paper ballot. Electronic voting systems are *indirect*—they interpose a layer of mechanism between the voter and the audit trail, risking the possibility that the mechanism is not faithfully capturing the voter's preferences.

Nonetheless, paper ballots are not perfect either, and Shamos [Sh93] gives interesting arguments in favor of electronic audit trails. Saltman's classic work [Sa88] discusses in some detail audit-trail requirements for electronic voting systems.

County-level decisions on voting technology has benefits. There are clear and probably compelling advantages to specifying and purchasing voting

systems on a state-wide basis rather than county by county, as is currently the case in the U.S. But we should not lose sight of two arguments to the contrary.

First, just as a woodland's diverse variety of plants can provide better resistance to pathogens than the farmer's single crop, so too can a variety of voting technologies provide resistance to an adversary's attack, as there is no common point of vulnerability for the whole system.

Second, we need ways to gain experience with new voting systems. One good way is to allow individual counties to experiment with techniques that are different than the state-wide norms.

The ability to handle disabled voters will become increasingly important. Existing voting systems tend to be poor at accommodating the needs of disabled voters. For example, blind voters have had to trust election officials to read the ballots and enter their votes. Electronic voting systems are capable of supporting a diversity of interfaces to the voter.

Our largest security problem is likely to be *absentee ballots.* Absentee voting has increased dramatically over the past decade. Indeed, some states, such as Oregon, vote entirely by mail. Remote electronic voting can be viewed as a version of absentee voting.

In my opinion, however, by allowing such an increase in absentee voting we have sacrificed too much security for the sake of voter convenience. While voters should certainly be allowed to vote by absentee ballot in cases of need, allowing voting by absentee ballot merely for convenience seems wrong-headed. I would prefer seeing "Voting Day" instituted as a national holiday to seeing the widespread adoption of unsupervised absentee or remote electronic voting.

2.4 Summary

Some paper-based voting technologies, such as optical scanning, offer reasonable balances of security, ease of use, cost, simplicity, and reliability. (Other paper-based technologies, such as punch cards, should definitely be phased out.)

Electronic voting systems promise benefits in terms of ease of use, especially for disabled voters. Because of the software-based and indirect character of electronic voting systems, these benefits come at the cost of increased complexity and at the risk of decreased security.

While electronic voting from home should perhaps forever remain too risky a fantasy, electronic poll-site voting may provide, even in the near term, worthwhile improvements to paper-based voting technologies. Cryptographic techniques will certainly be essential in any electronic voting technology, as will better methods for addressing the "secure platform problem."

References

Ca00. California Internet Voting Task Force. Final report. Available at http://www.ss.ca.gov/executive/ivote/.

Gr00. Rachel Greenstadt. Electronic voting bibliography, January 2000. Available at http://theory.lcs.mit.edu/~cis/voting/greenstadt-voting-bibliography.html.

IPI00. Internet Policy Institute. Internet voting. Available at http://www.internetpolicy.org.

JSI96. M. Jakobsson, K. Sako, and R. Impagliazzo. Designated verifier proofs and their applications. In Ueli Maurer, editor, *Advances in Cryptology - EuroCrypt'96*, pages 143–154, Berlin, 1996. Springer-Verlag. Lecture Notes in Computer Science Volume 1070.

CTM00. California Institute of Technology and Massachusetts Institute of Technology. Voting technology project. http://www.vote.caltech.edu/.

FER01. National Commission on Federal Election Reform. Available at http://www.reformelections.org.

R00. Avi Rubin. Security considerations for remote electronic voting over the internet, 2000. Available at http://avirubin.com/e-voting.security.pdf.

Sa88. Roy G. Saltman. Accuracy, integrity, and security in computerized vote-tallying. Technical report, Computer Science and Technology, National Bureau of Standards, Gaithersburg, MD 20899, August 1988. NBS Special Publication 500-158. Available at http://www.itl.nist.gov/lab/specpubs/500-158.htm.

Sh93. Michael Shamos. Electronic voting—evaluating the threat. Presented at CFP'93. Available at http://www.cpsr.org/conferences/cfp93/shamos.html.

3 C. Andrew Neff: E-Voting: Proceed with Caution

3.1 Introduction

Election 2000 has shown the need for a well-defined audit process that can be independently verified – at very least by multiple parties with disparate interests in the outcome of the election; better yet, by anyone who cares about the results. There is a growing groundswell of opinion that computers could be used to make elections more accurate and efficient, but they bring with them their own pitfalls. Since electronic data is so easily altered, simple-minded electronic voting systems cannot produce an audit trail that is as strong as a paper audit trail. Further, limitations on the reliability of the supporting electronic hardware and software are also an important concern. However, it would be both naive and unscientific to conclude that it is *impossible* to implement an electronic voting system which meets or exceeds the integrity standards of our best conventional systems. If designed properly, an electronic voting system can actually produce an audit trail that is even stronger than conventional ones – including paper based systems. This is exactly what was needed in the 2000 U.S. Presidential

Election to avoid the costly and time consuming dispute process that ensued; a system that is automated and indisputable, while preserving ballot secrecy.

3.2 Fundamentals – E-Voting vs. I-Voting

In any discussion of "high-tech" voting solutions, it is important to make a distinction, from the start, between systems which use digital data to *capture the original voter selections* and/or act as *official record* thereof – *"e-voting"* systems – and systems which use the remote connectivity of the Internet, or other public network to cast, collect and tabulate ballots – *"i-voting"* systems. Clearly the class of e-voting systems includes all i-voting systems, but there are e-voting systems *already in use* today – so called DRE equipment – which are obviously not i-voting systems. Unless one is careful, it is easy to lose track of this distinction in the arguments for and against either one of them.

Assuming that all systems must meet certain standards for integrity and dependability, i-voting clearly presents a much greater set of problems than e-voting alone. However, the problems associated simply with e-voting are already challenging ones. To be sure, electronic computers can manage election data far more easily and efficiently than physical methods, the trouble is that computers are inherently only as trustworthy as the people who administer them. They are also completely opaque, unlike the physical paper ballot box that can be watched at all times; and electronic data is far more easily altered or destroyed than physical ballots – especially in large quantities.

The key to building systems that do not suffer these problems is to leverage the strengths of digital data, rather than trying to employ the same procedures that are used to protect conventional systems. Attention must be focused on ways to guarantee integrity of the election data itself, rather than on the custody of machines that are handling it. If a complete record, or *transcript*, of all election data – from who has voted to the specific computation steps used to arrive at the final tally – can be collected and represented in such a way that not one bit of it can be altered without creating intrinsic inconsistency, then the goal of an indisputable electronic election can be achieved. This is really the only way an e-voting recount makes any sense at all. Running a count over and over on the same machine, or set of machines, proves nothing about the true election results. It only proves that the software that is running can display the same numbers repeatedly. Of course, the transcript must also not compromise voter privacy.

An election transcript provides a tool for better election audit than ever. Without the power of modern electronic systems, a central audit of this scope was out of reach. With conventional systems, each election participant – voters and candidates – have to trust parts of the audit that are only enforced by local procedural requirements. There is no way to verify after the fact that these procedures have been sufficiently implemented, or if they have been subverted.

3.3 Requirements Must Be Scientific and Unbiased

As new election systems are proposed, it is important to keep the debate from sinking to unscientific levels. New systems should not be shoehorned to fit old ones simply because we are used to the old ones. A good example of this is the debate over media. Paper ballots have many desirable qualities, but the requirements for new systems should be phrased in terms of the *fundamental qualities* that should be maintained, rather than artificially insisting that they continue to use paper for ballot recording. For example, a reasonable storage requirement might be that the storage media have "99.99% chance of surviving a 8.0 Richter scale earthquake". An example of an unreasonable, biased requirement is "the system should be capable of printing *paper ballots* for hand recount."

Once a suitable set of fundamental system requirements is agreed upon, the unsettling vulnerabilities of our conventional election systems become apparent. In this light, the benefits of moving to new election system technology may begin to outweigh the risks. Typically, the definition of specific requirements are the responsibility of legislative, or administrative bodies. However, VoteHere's experience with technology and its capabilities can be a source of useful information to legislators. To that end, we suggest an outline for the general requirements decisions that must be made at the highest level.

1. **Fairness**. Only votes from *distinct*, *eligible* voters should be counted in the final tally.
2. **Accessibility**. No eligible voter should be prevented, or "deterred" from casting his/her vote, either by malicious or accidental forces.
   - The difference between "deterred" and "inconvenienced" will be difficult to pinpoint. Many voters already feel that the voting process is an inconvenience.
3. **Accuracy**. The final published election results should be, mathematically, an exact "count" (i.e. sum, or in general, aggregation when using more complex tabulation rules) of the collection of *intended choices* made by all the participating voters. This requirement breaks into two pieces:
   3.1. The results should be an exact "count" of the ballots recorded in the "ballot box", or on the "ballot media".
   3.2. Each vote in the "ballot box" should be an accurate representation of the voter's intended choice.
      - As with accessibility, a requirement of this type can only be specified with "reasonable precision". No system is capable of preventing *all* types of voter errors.
4. **Privacy**. The contents of each ballot should be known only to the voter who cast it. This is simple, but there are subtleties that can only be addressed with a complete threat model. For example
   4.1. No system can protect a voters ballot secrecy from a collusion of all the other voters.

4.2. Current "mail in" absentee voting does not protect privacy against all threats. In fact, privacy can be broken by a collusion of two at the time that envelopes are opened. Given this, is cryptographically protected privacy sufficient, especially if only the voter has the corresponding key? There are many ways to spy on voter choices that are easier and less expensive than breaking RSA, for example.

5. **Receipt Freeness**. To discourage both vote buying and coercion, it is important that a voter not be able to *prove* how he/she voted.
   - Cooperative vote selling will never be prevented by legislation or by any election system. The best that can be done is to significantly limit the attractiveness of such activity, a goal that is accomplished when a prospective buyer, or coercer, can not be sure of the success of his/her efforts. (It should be noted, however, that current "mail in" absentee voting does nothing to address this problem.)

Specific requirements imposed on election systems should be chosen so as to most effectively *implement* the general considerations above. They should not be imposed as a means of *a priori design or engineering*.

3.4 Practical Considerations

In addition to the general considerations of the previous section, there are practical issues that any new voting technology must take into consideration.

Integration with other systems. It would be unreasonable to expect that conventional voting systems will disappear overnight, even in isolated precincts. It will take time for voters to gain comfort with new systems, and for counties to migrate to them. As a result, new voting systems must be able to fit as part of an aggregate system – the way that, in many jurisdictions, "mail in" voting fits together with poll site voting. This, then, requires some care to be sure that overall election integrity remains in tact – voters shouldn't be able to vote once with each system, for example.

Lack of PKI. For i-voting especially, it is crucial to be able to associate eligible voters – *people* – with digital credentials. For now, an infrastructure to support this which is both widely used, and robust is missing. This may change in the near future, but for now, election systems will need to build custom solutions. This means some integration with voter registration systems.

Mechanism of verification. Election systems in the theoretical literature generally fall into two categories – those that are *universally verifiable* and those that are *voter verifiable*. The former produce election transcripts that allow complete verification of election integrity by any independent entity, while the latter require that voters check the validity of their vote in the ballot box to be sure that data has not been changed. The integrity of a voter verifiable system is based on the premise that "enough" voters *will* verify their ballots. In practice, these systems are undesirable for two reasons

1. Some voters will be mischievous, or forgetful. There is then no good mechanism for deciding whether the system has been compromised, or if just some group of voters wishes to make the system appear compromised.
2. It is hard to get a large fraction of voters to vote in the first place. It will be much harder still to get them to verify their ballots.

Redundant infrastructures not viable. A brute force attempt at assuring election integrity is through redundancy. The idea is that each voter submits multiple copies of her ballot to multiple, *independent* tabulation authorities. Results are then determined by way of "majority rule", or some variant. In practice, such solutions are seriously flawed.

1. From a business perspective, how will the authorities be maintained in a truly independent manner? It is not cost effective, competitive, or interesting to each of the replicated authorities. We could propose that the government subsidize them all, but isn't this even worse?
2. Dispute resolution is complicated and potentially costly.
3. Voters now have a new mode of malicious behavior – submit conflicting votes to the various authorities.
4. Network reliability problems go up exponentially with the number of authorities.

In this respect, the right model for e-voting comes from the conventional poll site model itself: *One poll site, lot's of observers*, leads to *One tabulation center, lot's of crypto keys.*

Client trust. Theoretical voting protocols presume that the voter does "her own" computation. In fact, the computation is done by a computing device, which may itself not be trusted. Supervised e-voting systems may be able to prevent this threat through careful procedure, but it is much more difficult to address in the situation of i-voting.

Network weaknesses. Even if all election data can be protected from compromise, there is still the practical problem of getting it from one place to another. The Internet is vulnerable to *Denial of Service* attacks, and these must be taken seriously. However, the criteria put on e-voting systems for reliability should be reasonable. All systems are subject to some risk of DoS. Conventional poll sites can be forced closed for several reasons, such as earthquake, fire, or bomb threat. Voters can be prevented from getting to them by something as common as a traffic jam. As long as voters using an e-voting system as "first choice" can use other systems as a fall back, the standards for reliability should not be absurdly high.

3.5 VoteHere Philosophy

The philosophy at VoteHere, Inc., is to build systems completely transparent to public review and scrutiny. Our *ambition* is to achieve this by way of published and accepted protocol, thereby eliminating the need for trusted, audited or otherwise independently inspected software and hardware. Where components

cannot clearly meet this goal, they must be "sun lighted" by careful, independent certification. This combination of open cryptographic protocol and component review creates the requisite level of trust and dependability vital to any public election system.

4 Avi Rubin: The Feasibility Of Remote E-Voting

The feasibility of remote electronic voting in public elections is currently being studied by the National Science Foundation by request of the former President of the United States (see `http://www.netvoting.org/`). Remote electronic voting refers to an election process whereby people can cast their votes over the Internet, most likely through a web browser, from the comfort of their home, or possibly any other location where they can get Internet access. There are many aspects of elections besides security that bring this type of voting into question. The primary ones are:

coercibility. the danger that outside of a public polling place, a voter could be coerced into voting for a particular candidate.
vote selling. the opportunity for voters to sell their vote.
vote solicitation. the danger that outside of a public polling place, it is much more difficult to control vote solicitation by political parties at the time of voting.
registration. the issue of whether or not to allow online registration, and if so, how to control the level of fraud.

The possibility of widely distributed locations where votes can be cast changes many aspects of our carefully controlled elections as we know them. The relevant issues are of great importance, and could very well influence whether or not such election processes are desirable. However, in this paper, we focus solely on the security considerations as they relate to conducting online public elections. In particular, we look at remote online voting, as opposed to online voter registration, which is a separate, but important and difficult problem. We also focus solely on public elections, as opposed to private elections, where the threats are not as great, and the environment can be more controlled.

4.1 The Platform

On the platforms currently in the most widespread use, once a malicious payload reaches a host, there is virtually no limit to the damage it can cause. With today's hardware and software architectures, a malicious payload on a voting client can actually change the voter's vote, without the voter or anyone else noticing, regardless of the kind of encryption or voter authentication in place. This is because the malicious code can do its damage before the encryption and authentication is applied to the data. The malicious module can then erase itself

after doing its damage so that there is no evidence to correct, or even detect the fraud. To illustrate, let's look at a program that exemplifies the level of vulnerability faced by hosts.

The program we describe, Backorifice 2000 (BO2K) is packaged and distributed as a legitimate network administration toolkit. In fact, it is very useful as a tool for enhancing security. It is freely available, fully open source, extensible, and stealth (defined below). The package is available at `http://www.bo2k.com/`. BO2K contains a remote control server that when installed on a machine, enables a remote administrator (or attacker) to view and control every aspect of that machine, as though the person were actually sitting at the console. This is similar in functionality to a commercial product called PCAnywhere. The main differences are that BO2K is available in full source code form and it runs in stealth mode.

The open source nature of BO2K means that an attacker can modify the code and recompile such that the program can evade detection by security defense software (virus and intrusion detection) that look for known signatures of programs. A signature is a pattern that identifies a particular known malicious program. The current state of the art in widely deployed systems for detecting malicious code does not go much beyond comparing a program against a list of attack signatures. In fact, most personal computers in peoples' houses have no detection software on them. BO2K is said to run in stealth mode because it was carefully designed to be very difficult to detect. The program does not appear in the Task Menu of running processes, and it was designed so that even an experienced administrator would have a difficult time discovering that it was on a computer. The program is difficult to detect even while it is running.

There can be no expectation that an average Internet user participating in an online election from home could have any hope of detecting the existence of BO2K on his computer. At the same time, this program enables an attacker to watch every aspect of the voting procedure, intercept any action of the user with the potential of modifying it without the user's knowledge, and to further install any other program of the attackers desire, even ones written by the attacker, on the voting user's machine. The package also monitors every keystroke typed on the machine and has an option to remotely lock the keyboard and mouse. It is difficult, and most likely impossible, to conceive of a web application (or any other) that could prevent an attacker who installs BO2K on a user's machine from being able to view and/or change a user's vote.

4.2 The Communication Infrastructure

A network connection consists of two endpoints and the communication between them. The endpoints here are a user's host and an elections server. While it is in no way trivial, the technology exists to provide reasonable protection on the servers. This section deals with the communication between the two endpoints.

Cryptography can be used to protect the communication between the user's browser and the elections server. This technology is mature and can be relied

upon to ensure the integrity and confidentiality of the network traffic. This section does not deal with the classic security properties of the communications infrastructure; rather, we look at the availability of the Internet service, as required by remote electronic voting over the Internet.

Most people are aware of the massive distributed denial of service (DDOS) attack that brought down many of the main portals on the Internet in February, 2000. While these attacks brought the vulnerability of the Internet to denial of service attacks to the mainstream public consciousness, the security community has long been aware of this, and in fact, this attack was nothing compared to what a dedicated and determined adversary could do. The February attack consisted of the installation and execution of publicly available attack scripts. Very little skill was required to launch the attack, and minimal skill was required to install the attack.

The way DDOS works is that a program called a daemon is installed on many machines. Any of the delivery mechanisms described above can be used. One other program is installed somewhere called the master. These programs are placed anywhere on the Internet, so that there are many, unwitting accomplices to the attack, and the real attacker cannot be traced. The system lies dormant until the attacker decides that it is time to strike. At that point, the attacker sends a signal to the master, using a publicly available tool, indicating a target to attack. The master conveys this information to all of the daemons, who simultaneously flood the target with more Internet traffic than it can handle. The effect is that the target machine is completely disabled.

We experimented in the lab with one of the well known DDOS programs called Tribe Flood Network (TFN), and discovered that the attack is so potent, that even one daemon attacking a Unix workstation disabled it to the point where it had to be rebooted. The target computer was so overwhelmed that we could not even move the cursor with the mouse.

There are tools that can be easily found by anyone with access to the web that automate the process of installing daemons, masters, and the attack signal. People who attack systems with such tools are known as script kiddies, and represent a growing number of people. In an election, the adversary is more likely to be someone at least as knowledgeable as the writers of the script kiddy tools, and possibly with the resources of a foreign government.

There are many other ways to target a machine and make it unusable, and it is not too difficult to target a particular set of users, given domain name information that can easily be obtained from the online registries such as Register.com and Network Solutions, or directly from the WHOIS database. The list of examples of attacks goes on and on. A simple one is the ping of death, in which a packet can be constructed and split into two fragments. When the target computer assembles the fragments, the result is a message that is too big for the operating system to handle, and the machine crashes. This has been demonstrated in the lab and in the wild, and script kiddy tools exist to launch it.

The danger to Internet voting is that it is possible that during an election, communication on the Internet will stop because attackers cause routers to crash,

election servers to get flooded by DDOS, or a large set of hosts, possibly targeted demographicly, to cease to function. In some close campaigns, even an untargeted attack that changes the vote by one percentage point could sway the election.

4.3 Conclusions

A certain amount of fraud exists in the current offline election system. It is tolerated because there is no alternative. The system is localized so that it is very unlikely that a successful fraud could propagate beyond a particular district. Public perception is that the system works, although there may be a few kinks in it here and there. There is no doubt that the introduction of something like remote electronic voting will, and should, come under careful scrutiny, and in fact, the system may be held up to a higher standard. Given the current state of widely deployed computers in peoples' homes and the vulnerability of the Internet to denial of service attacks, we believe that the technology does not yet exist to enable remote electronic voting in public elections.
A full paper on this topic is available at
`http://avirubin.com/e-voting.security.html`.

5 Ed Gerck: Voting System Requirements

This section presents a set of voting system requirements that are consistent, technologically neutral, can be applied to paper, electronic and network (Internet) voting, and exceed the current requirements for paper-based ballots and electronic voting DRE (Direct Recording Electronic) machines. The requirements are based on the principles of "Information Theory" and of "trust as qualified reliance on information." The principles favoring multiple, independent channels of information over one purportedly "strong" channel. However, adding multiple channels can also decrease reliance if the design principles laid out in these requirements are not followed.

5.1 Background

As defined by Alan Turing some fifty years ago, a mathematical method is effective if, loosely speaking, it can be set out as a list of instructions which a human clerk who works obediently with paper and pencil can follow, for as long as is necessary, but without insight or ingenuity. Together with Alonzo Church, Turing, in fact, argued that every effective mathematical method can be carried out by a sufficiently powerful computer (represented by the universal Turing machine).

The above mentioned Voting System Requirements were born out of the desire to create products that would allow modern computer-based technology to automate and truly emulate the secure desirable properties motivated by what

has been collected throughout centuries of public voting. Put differently, we ask: *can we use a perfect clerk in elections—one who works obediently with paper and pencil, for as long as is necessary, but without insight or ingenuity?*

Indeed, if perfect clerks were to conduct an election using paper-ballots, this would provide the best model we have for a public election. Such an election would be, for example: anonymous (avoiding collusion, coercion), secret (all cast votes are unknown until the election ends) and yet correct (all votes are counted) and honest (no one can vote twice or change the vote of another participant), oftentimes also complete (all voters must either vote or justify absence). In such an election system, if we know the voter (e.g., in voter registration) we cannot know the vote and if we know the vote (e.g., in tallying) we cannot know the voter. After an election, all votes and all voters are publicly known—but their connection is both unprovable and unknown.

But: real-life clerks are not perfect! Neither are computer systems! Thus, we need to introduce the concept of qualified reliance on information in terms of providing proofs (e.g., proof of voting, proof of correctness) that can be objectively evaluated and not just subjectively accepted or taken at face value.

To discover and rate such proofs, the requirements employ the idea that one should favor multiple, independent communication channels over one "strong" channel. Such an idea was successfully used by the Moguls in India some 500 years ago in the context of combating corruption [1], and was mathematically described by Claude Shannon some 50 years ago in the context of combating noise when he introduced his Information Theory [2], a well-known general theory of communication processes.

Thus, for example, how can a voting system prove that the vote received at the ballot box is the same vote seen and cast by a voter? This question is not easier to answer if the voter is close to the ballot box then if he is far away. Distance plays no role, contrary to what one might think at first. The *fundamental problem of voting* is that the voter cannot see his tallied vote, hence the voter has no way of knowing if information sent through the communication channel (which may be very short) equal that which was received and tallied. This problem is oftentimes called the "vote-gap problem" by the author.

To solve this question in electronic voting, some advocate printing a paper copy of the ballot, which the voter can see and verify that it is identical to the ballot she intended to cast, and then sending the paper copy to ballot box A while an electronic copy of that same ballot is sent to ballot box B. The idea is that ballot box B can be tallied quickly while ballot box A will be used as a physical proof for a manual recount. Such a suggestion is oftentimes advanced as the sine qua non solution to voting reliability in electronic voting.

But what makes the introduction of a paper ballot special is not the fact that it is paper instead of bits. It is the fact that the voter is actually casting his vote twice. We now have two independent channels of information for the ballot, one from the terminal as source B, the other one from the printer as source A. We denote the multiplicity of such channels N (In our case $N = 2$).

In other words, this design provides for two outputs: ballot A and ballot B. However, in the event of a discrepancy between the two, no resolution is possible inside the system. The situation can thus be summarized:

- N = 1: If the system is always similar to a perfect clerk then N = 1 (one channel) suffices, whether paper or electronic. But if we use a system with N = 1, we cannot define any level of reliance on the final result except that which was assigned a priori.
- N = 2: If we add one independent channel (e.g., the paper ballot) to a system that already provides one channel (e.g., electronic ballot), this creates a system with N = 2. However, this additional channel makes the system indeterminate and still incapable of, by itself, defining any level of reliance on the final result except that which was assigned a priori (e.g., paper is more trustworthy).

Clearly, before considering other well-meant suggestions (which might be similarly ill-fated), what is necessary is to seek a logically provable solution to reliability problems caused by imperfect communication systems.

Such a solution needs to consider not only machine-machine communication channels but also human-machine communication channels because the voter can act as a source and as verifier in more than one part of the system. Further, human-human communication channels must be considered because we do not want machines to have the potential to "run amok", unchecked.

Information Theory [2] can be used to describe such communication channels and, as previously noted, the concept of qualified reliance on information can be introduced as a formal definition of trust [3] in order to rate such channels in terms of providers of proofs.

As a result, the only provable solution to increased reliability in communications (e.g., the communication between the voter as a sender and the ballot box as a receiver) turns out to be increasing the number/capacity of independent channels until the probability of error is as close to zero as desired (direct application of Shannon's Tenth Theorem in Information Theory [2]). To be complete, the solution should consider not only machine-machine communication channels but also human-machine and human-human ones. Thus, if an electronic system is able to provide N proofs (human and machine based), these N proofs for some value of N larger than two will become more reliable than one so-called "physical proof"—even if this one proof is engraved in gold or printed on paper.

An undefined system also presents opportunities for fraud (e.g., someone can change and/or delete some paper ballots after the election in order to cast doubt on the integrity of the entire election). It is also open to attacks (e.g., a group of voters might agree beforehand to call out a "discrepancy" after they vote and thereby disrupt an election, which is similar to a "denial of service" attack).

Thus, we need a real-world voting system—not one that is based on perfect parts (N = 1) or one that produces an undefined result in the case of a single error (N = 2). In order to provide for qualified

reliance on information, such a voting system needs to have multiple independent channels.

In plain English, the greater the number of independent channels for the verification of a result, the greater trustworthy the result is.

However, suppose the terminal where the voter enters his choices changes them to something else and then sends this information over N different channels, what difference does it make if it is N = 1, 2 or 500?

None! In such a case N would still be 1 for the ballot channel. The 2 or 500 channels are not independent for the ballot because they all originate as copies of that single stored potentially corrupted ballot. So, it does not make a difference in terms of ballot reliance. This would, however, make a difference in terms of communication reliance, in which there are now different transmission channels, 2 or 500 channels for which each channel could behave as a correction channel for the others. Namely, in this case the ballot box would more probably receive the right ballot (even though it may have been corrupted before transmission) and more so for N = 500 than for N = 1.

What is needed is therefore a requirement to include several truly independent ballot, transmission and audit channels—whether or not electronic transactions are used. These channels should be employed in rating the reliance on each node of an end-to-end balloting system, even during the election and in real time. There should be several ways to implement this requirement and channels could be added also in time and context, not just in space. Channels can also transport information by reference, not just information by value.

What is also needed is a way to allow the voter to verify results, for example the presence/absence of her ballot at the ballot box and whether her ballot at the ballot box is a valid one. This is useful because sufficient indirect verification does produce trust. "Trust but verify" is a mode typically preferred by our collective wisdom and it is definitely applicable here. It is important to note that even if just a fraction of the voters (e.g., 5%) do verify the results, the capability of verification is already a deterrence to fraud because a cheater has no way of knowing who will verify, or not.

Another characteristic of a good voting system is that the only person whom you prove the vote to is the voter. If the proof can be shown to someone else, then the vote can be coerced or sold. Therefore, when using multiple channels of information, they either have to be deniable by the voter or else temporary so that the voter cannot be threatened or hurt as a result of the vote.

Regarding the use of paper, it is important to note that the reason to distrust a paper/electronic voting system with N = 2 is not based on distrust in paper. Paper is just another communication channel. The reason is that adding paper does not solve the problem and, in fact, makes the problem indeterminate. This is so, since we need N larger than 2. Certainly, paper can be one of the channels, if desired, because the channel make-up is irrelevant. But a cost-benefit analysis might result in the use of non-paper channels.

Next, another question that must be addressed is the possibility of all-electronic voting systems. Should we trust them and why?

Nowadays, all-electronic systems and computers are used in flying commercial and military jets. And yet, no one in the public is afraid that a terrorist will introduce a virus in the system and will down all commercial jets worldwide, or all U.S. military jets. Why? Because there is a designed redundancy at many levels in the system. For example, there are three independent laser inertial navigation sensors and any decision on the plane's position depends on the agreement of at least two of them, which decision is further verified by a GPS system, as well as flight time and speed calculations.

Thus, voting systems—like any other type of systems—derive their trustworthiness from the fact that they work consistently, both conceptually and perceptually. However, in the absence of an easy conceptual understanding of the system (e.g., a laser inertial navigation sensor) that the average user is able to grasp, a sufficiently coherent perceptual understanding (e.g., observing that the system works) is enough to eventually build trust in the system.

Trust may also be denied by the design itself, because disasters may occur at any time if the principles of communication reliance (i.e., trust itself) are not taken into account. To visualize this, imagine a plane that would be flown with just two navigation sensors, one compass-based and the other electronic—we would then have an idea of the disastrous consequences of using a paper/electronic voting system with N = 2, even though a physical channel is used (compass, paper).

Thus, we can conclude that the deciding factor in trusting a system is not whether it includes one or even two sources of information that can be touched or seen in physical form (e.g., a paper in your hands, a paper behind a screen, a compass needle behind a screen).

A factor that mitigates against an all-electronic voting system is the fact that although paper and electronic records are both vulnerable to subversion, it is a lot easier to change what is in an electronic record than it is to change what is on paper.

Thus, electronic records need to be bound to other references in a manner that is demonstrably inaccessible to an attacker, both through physical access controls and through cryptographic protocols.

Moreover, there really needs to be a step-by-step description of the voting process, so that when someone asks, "What if the intruder succeeds in breaking into the system to change X?" this can be clearly answered, for example, by:

> (i) to change X would cause a subsequent binding failure, thus it would be detectable except with parallel access to Y and Z, which are independently inaccessible, or
> (ii) knowledge of an alternate (and attacker-desirable) value for X is insurmountably difficult to achieve, and the effort could not be leveraged to any other X.

Put most plainly, people know that ordinary voting systems can be subverted by someone who can bribe enough individuals to collude, but the physical fact

of several tons of paper ballots still represents somewhat of an obstacle to an "easy subversion" in the eyes of many.

In contrast, people are well aware that electronically one can modify a million records with as little as a few keystrokes. This is the "fear" that needs to be addressed in an all-electronic system. Further, such a subversion can be massive and rapid, executed from the safety of a remote laptop, etc. so that it would be unavoidable.

Of course, one alternative to reduce fear would be education. To educate voters regarding the very nature of distributed cryptographic assurances and at a level where the concepts are not hidden behind excessive abstractions.

But cryptography is not, by itself, the critical issue, nor is it the silver bullet. Further, no amount of education will stop attackers, on the contrary, it may aid them.

Instead, *voting systems can use the concept of multiple independent communication channels to make it as impossible as desired to tamper with the electronic ballot both before and after it is cast.*

Here, the question is not how many copies of paper or bits one has, but how many independent channels the attacker needs to subvert versus how many independent correction channels one has available during such an attack. Of course, if the attacker is able to subvert the correction channels while attacking the other channels, then they will not be independent.

Therefore, the same mechanism that protects the casting of a ballot must also be used to protect presenting the ballot. And this needs to be given as a set of requirements which work together in an end-to-end design.

These requirements are therefore general principles, valid for any physical implementation of a "ballot"—whether as print marks on paper, pits on a CD-ROM surface, electrons hitting a video screen (electronic ballot), modulated electromagnetic waves, bits in a network protocol or any other form of information transfer to and from the voter. They also apply to any form of voting, including majority voting and single transferable votes. The requirements may be applied in their entirety or merely a subset may be used.

To achieve these goals, the requirements should be able to handle voting rules of any type and should apply to voting systems anywhere in the world. However, the main objective here is for the requirements to be as complete and as independent from one another as possible, without sacrificing consistency. It is understood that "completeness" is an elusive goal that might never be reached when we consider the diversity of election needs [4], while "consistency" is a necessary feature for the requirements to work together in a particular election. In short, this was the reason to stop after 16 requirements. Increasing the number of requirements can risk decreasing their consistency, in general [4]. Of course, other requirements may be added, or deleted, as needed.

Some of the words used in the requirements may have different (and equally valid) meanings in other contexts (e.g., "voter privacy"). Therefore, the requirements also include the operational definitions of the main words used. Three words are, however, used without a definition even though they could also be

misunderstood. These words are "trust" [3], "manifold" and "meshwork" [5], as defined in the references.

5.2 Summary of Requirements

A voting system of any type and media needs to satisfy various requirements which are summarized in the following 16 points.

1. Fail-safe voter privacy. Definition: "Voter privacy is the inability to link a voter to a vote." Voter privacy MUST be fail-safe—i.e., it MUST be assured even if everything fails, everyone colludes and there is a court order to reveal all election data. Voter privacy MUST be preserved even after the election ends, for a time long enough to preserve backward and forward election integrity (e.g., to prevent future coercion due to a past vote, which possibility might be used to influence a vote before it is cast).

2. Collusion-free vote secrecy. Definition: "Vote secrecy is the inability to know what the vote is." Vote secrecy MUST be assured even if all ballots and decryption keys are made known by collusion, attacks or faults (i.e., vote secrecy MUST NOT depend only on communication protocol and cryptographic assumptions, or on a threshold of collusion for the keyholders).

3. Verifiable election integrity. Definition: "Election integrity is the inability of any number of parties to influence the outcome of an election except by properly voting." The system MUST provide for verifiability of election integrity for all votes cast. For any voter the system MUST also provide for direct verifiability that there is one and only one valid ballot cast by the voter at the ballot box.

4. Fail-safe privacy in verification. If all encrypted ballots are verified, even with court order and/or with very large computational resources, the voter's name for each ballot MUST NOT be revealed.

5. Physical recounting and auditing. We MUST provide for reliability in auditing and vote recounting, with an error rate as low as desired or, less strictly, with an error rate comparable or better than conventional voting systems [8]. The auditing and vote proofs MUST be capable of being physically stored, recalled and compared off-line and in real-time during the election, without compromising election integrity or voter privacy, and allowing effective human verification as defined by election rules.

6. 100% accuracy. Every vote or absence of vote (blank vote) MUST be correctly counted, with zero error [8].

7. Represent blank votes. We MUST allow voters to change choices from 'vote' to 'blank vote' and vice-versa, at will, for any race and number of times, before casting the ballot.

8. Prevent overvotes. As defined by election rules. We MUST provide automatic "radio button" action for single-vote races. If overvoting is detected in multiple-vote races, we MUST warn the voter that a vote has to be cleared if changing choices is desired. This warning MUST be made known only to the voter, without public disclosure.

9. Provide for null ballots. As defined by election rules, we MAY allow voters to null races or even the entire ballot as an option (e.g., to counter coercion; to protest against lack of voting options). Overvoting, otherwise prevented by requirement #8, MAY be used as a mechanism to provide for null ballots.

10. Allow undervotes. As defined by election rules, the voter MAY receive a warning of undervoting. However, such a warning MUST NOT be public and MUST NOT prevent undervoting.

11. Authenticated ballot styles. The ballot style and ballot rotation (changes between individual ballot representations) to be used by each voter MUST be authenticated and MUST be provided without any other control structure but that which is given by the voter authentication process itself.

12. Manifold of links. We MUST use a manifold [5] of redundant links and keys to securely define, authenticate and control ballots. We MUST avoid single points of failure—even if improbable. If networks are used, we MUST forestall Denial-of-Service (DoS) and other attacks with an error rate comparable or better than conventional voting systems [8].

13. Off-line secure control structure. We MUST provide for an off-line secure end-to-end control structure for ballots. We MAY use digital certificates under a single authority. Ballot control MUST be data-independent, representation-independent and language-independent.

14. Technology independent. We MUST allow ballots and their control to be used off-line and/or in dial-up and/or in networks such as the Internet, with standard PCs or hand-held devices used to implement their components in hardware or in software, alone or in combination for each part.

15. Authenticated user-defined presentation. We MUST enable the ballots to dynamically support multiple languages, font sizes and layouts, so that voters can choose the language and display format they are most comfortable with when voting as allowed by law and required by voters with disabilities, without any compromise or change to the overall system, from an authenticated list of choices defined by the election rules.

16. Open review, open code. We should allow all source code to be publicly known and verified (open source code, open peer review). The availability and security of the system must not rely on keeping its code or rules secret (which cannot be guaranteed), or in limiting access to only a few people (who may collude or commit a confidence breach voluntarily or involuntarily), or in preventing an attacker from observing any number of ballots and protocol messages (which cannot be guaranteed). The system SHOULD have zero-knowledge properties (i.e., observation of system messages do not reveal any information about the system). In fact, only keys MUST be considered secret.

5.3 Comments

Implementations and examples [9] are discussed in the full paper, available in [10].

These requirements include comments and references from Tony Bartoletti, Thomas Blood, Netiva Caftori, Gordon Cook, Hal Dasinger, Hugh Denton, Rosario Gennaro, Jason Kitcat, Brook Lakew, Elaine Maurer, Don Mitchel, Erik Nilsson, Michael Norden, Marcelo Pettengill, Roy Saltman, Bernard Soriano, Gene Spafford, Einar Stefferud, Arnold Urken, Eva Waskell, Thom Wysong, the IVTA tech WG (`http://www.mail-archive.com/tech@ivta.org/`), the CPSR-activists list, several cryptography lists, contributions from comments collected at Safevote's website, and from articles published in The Bell (`http://www.thebell.net`).

5.4 References

[1] *"... one of the earliest references to the security design I mentioned can be found some five hundred years ago in the Hindu governments of the Mogul period, who are known to have used at least three parallel reporting channels to survey their provinces with some degree of reliability, notwithstanding the additional efforts."* Ed Gerck, in an interview by Eva Waskell, "California Internet Voting." The Bell, Vol. 1, No. 6, ISSN 1530-048X, October 2000. Available online at `http://www.thebell.net`.

[2] Shannon, C., "A Mathematical Theory of Communication." Bell Syst. Tech. J., vol. 27, pp. 379-423, July 1948. Available online at `http://cm.bell-labs.com/cm/ms/what/shannonday/paper.html`.
Shannon begins this pioneering paper on information theory by observing that

"the fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point." He then proceeds to thoroughly establish the foundations of information theory, so that his framework and terminology have remained standard practice. In 1949, Shannon published an innovative approach to cryptography, based on his previous Information Theory paper, entitled Communication Theory of Secrecy Systems. This work is now generally credited with transforming cryptography from an art to a science. Shannon's Tenth Theorem states (cf. Krippendorf and other current wording): *"With the addition of a correction channel equal to or exceeding in capacity the amount of noise in the original channel, it is possible to so encode the correction data sent over this channel that all but an arbitrarily small fraction of the errors contributing to the noise are corrected. This is not possible if the capacity of the correction channel is less than the noise."*

[3] *"When we want to understand what trust is, in terms of a communication process, we understand that trust has nothing to do with feelings or emotions. Trust is that which is essential to communication, but cannot be transferred in the same channel. We always need a parallel channel. So the question is having redundancy. When we look at the trust issue in voting, it is thus simply not possible to rely on one thing, or two things even if that thing is paper. We need to rely on more than two so we can decide which one is correct. In this sense, the whole question of whether the Internet is trusted or not is simply not defined. The Internet is a communication medium and whatever we do in terms of trust, it is something that must run on parallel channels."* Ed Gerck, testimony before the California Assembly Elections & Reapportionment Committee on January 17, 2001, in Sacramento. Assemblyman John Longville (D), Chair. For an application of this model of trust to digital certificates, see "Trust Points" from `http://www.mcg.org.br/trustdef.txt` excerpted in "Digital Certificates: Applied Internet Security' by J. Feghhi, J. Feghhi, and P. Williams, Addison-Wesley, ISBN 0-20-130980-7, p. 194-195, 1998.

[4] This is similar to the situation found in Goedel's incompleteness theorem. The requirements form a logical system of some complexity and thus we do not expect such a system to be both complete and consistent.

[5] "Manifold" means a whole that unites or consists of many diverse elements and connections, without requiring these elements and connections to depend upon one another in any way. "Meshwork" is used to denote a manifold in the context of the Multi-Party protocol designed by Safevote to implement the requirements. A meshwork builds a meta-space in relationship to a space—a meshwork describes relationships about a space, not the space itself.

[6] *"We say that information-theoretic privacy is achieved when the ballots are indistinguishable independent of any cryptographic assumption; otherwise we will say that computational privacy is achieved."* In Ronald Cramer, Rosario Gennaro, Berry Schoenmakers, "A Secure and Optimally Efficient Multi-Authority

Election Scheme," Proc. of Eurocrypt'97. (available online at
http://www.research.ibm.com/security/election.ps).

[7] E. Gerck, "Fail-Safe Voter Privacy", The Bell, Vol.1, No.8, p. 6, 2000. ISSN 1530-048X. Available online at
http://www.thebell.net/archives/thebell1.8.pdf.

[8] *Accuracy* and *Reliability* are used here in the sense of standard engineering terminology, even though these different concepts are usually confused in non-technical circles. Lack of accuracy and/or reliability introduces different types of errors:

> (i) Reliability affects a number of events in time and/or space, for example, errors in transfers between memory registers. We know from Shannon's Tenth Theorem [2] that reliability can be increased so that the probability of such an error is reduced to a value as close to zero as desired. This is a capability assertion. It does not tell us how to do it, just that it is possible. This is the realm of requirements #12 and also #5, where one can specify an error rate as low as desired or, less strictly, an error rate "comparable or better than conventional voting systems".
> (ii) Accuracy affects the spread of one event, for example whether a vote exists. Here, requirement #6 calls for 100% accuracy. The requirement is that no "voter-intent" or "chad" or "scanning" issue should exist—which is feasible if, for example, each voting action is immediately converted to a standard digital form that the voter verifies for that event. Accuracy error can be set to zero because 100% accuracy is attainable in properly designed digital systems that (e.g., by including the voter) have no digitization error.

For an illustration of the above definitions of accuracy and reliability, see the four diagrams in http://www.safevote.com/caltech2001.ppt.

[9] "Contra Costa Final Report" by Safevote, Inc. Available upon request. Summary available at http://www.safevote.com.

[10] "Voting System Requirements", The Bell newsletter, ISSN 1530-048X, February 2001, archived at
http://www.thebell.net/archives/thebell2.2.pdf.

6 Summary

Given the presentations above and the relative popularity of the subject, a large numbers of questions and remarks were raised by the conference participants. We thank the participants for their active role and contributions to the usefulness of our panel. Since many opinions were expressed it is hard to report on all of them (and surely important comments are omitted herein).

Many shared the caution expressed by some of the panelists. The paraphrased comment "I am not sure what the next election technology is going to be, but

the next to next such technology will definitely be a paper technology" (by Matt Blaze) perhaps best represents this healthy skepticism. Specific concerns about the Internet reliability and immunity to attacks were raised by many participants. A doubt was raised (by Stefan Brands), claiming that the current theoretical work does not provide a sufficient level of privacy. The specific concern was that votes can be revealed if administrating machines collaborate, even in the most advanced protocols.

Social concerns regarding the suitability of modern technology to running democracy were discussed. The fear that aggressive modernization may generate a "voting divide" between those who use computers and those who do not, was expressed. The idea that politicians may not like the technology and will oppose its introduction, was raised as another potential hurdle to technological progress in political processes like voting.

Some more optimistic views were also expressed. They pointed out the problems of current systems on the one hand, and on the other hand they reminded us that electronic voting was somewhat successful in trials in the USA and in actual votings in countries like Brazil. Also noted is the fact that the current technology cannot remain forever the technology of choice for election, since this technology, even though it has been evolving very slowly, has been nevertheless evolving.

Overall, the panel represented the current state of the art of the business of electronic voting. We covered the commercial possibilities of supplying modern technology with sufficient levels of privacy, security, reliability and flexibility. We covered the basic requirements and the challenging issues that we need to cope with before the adoption of the new technology. We heard various opinions and interesting remarks regarding the diversified aspect of electronic voting. Naturally, e-voting industry researchers were more optimistic than their colleagues. The mix of opinions and variety of perspectives were instrumental in understanding the basic issues and problems regarding the reality of nation-wide e-voting, especially through the Internet. In fact, due to its popularity, the discussion centered around national (political) voting, and ignored other potential applications of the technology: e.g., inter-organizational small scale voting.

It is obvious that since there is quite uneasiness with the current election technology, the opportunity for using more modern technology exists. The actual adoption of electronic voting within electronic government and electronic democracy is going to stay a "hot issue" in the coming years. What technology is going to be adopted, and what level of cryptographic support will be used for the election itself and in securing election platforms, are open issues. Regardless of the differences of opinions and the various points of view, during the panel we learned about new angles to look at voting problems. We realized what are the burning issues in the area; issues of all kinds (social, business, systems, technology, policy and politics, etc.). Our hope is that we have stimulated further thinking, research and technological development which will motivate further studies of new subjects in all the relevant research areas.