UMBC

AN HONORS UNIVERSITY IN MARYLAND

A Study of Vote Verification Technologies

Part I: Technical Study

Prepared for the Maryland State Board of Elections

February 2006

National Center for the Study of Elections
of the
Maryland Institute for Policy Analysis & Research
University of Maryland, Baltimore County

A Study of Vote Verification Technologies
Part I: Technical Study

Prepared for the Maryland State Board of Elections

Principal Investigator: Donald F. Norris
Co-principal Investigator: Andrew Sears
Co-principal Investigator: Charles Nicholas

Editor: Anne V. Roland

National Center for the Study of Elections of the
Maryland Institute for Policy Analysis and Research
University of Maryland, Baltimore County

February 2006

Faculty Researchers

Aryya Gangopadhyay
Stephen H. Holden
George Karabatis
A. Gunes Koru
Chris M. Law
Donald F. Norris
John Pinkston
Andrew Sears
Alan T. Sherman
Dongsong Zhang

Graduate Student Assistants

Markus Dale
Kevin Fisher
Rosa Heckle
Kenny Khoo
Evan Perlman
Mohammad Ammar Sheikh
Tonya Zimmerman

Undergraduate Student Assistant

Chiedozie (Fred) Unachukwu

A Study of Vote Verification Technologies
Part I: Technical Study

Prepared for the Maryland State Board of Elections

Scholars at UMBC, working through the National Center for the Study of Elections of the Maryland Institute for Policy Analysis and Research, conducted a technical review of vote verification systems for the Maryland State Board of Elections (SBE). Initially, the review was supposed to include up to seven systems from the following organizations and individuals: VoteHere (Sentinel); SCYTL (Pnyx.DRE); Prof. Ted Selker, MIT (VVAATT); Diebold's VVPAT; Democracy Systems, Inc. (VoteGuard); IP.Com; and Avante. We determined that IP.Com did not represent a true vote verification technology, and Avante and Democracy Systems, Inc., declined to participate in the study. We also examined the SBE's procedures for "parallel testing" of the Diebold AccuVote-TS (touch screen) voting system in use in Maryland and used this as a baseline against which to evaluate the vote verification systems.

In conducting our analysis, we received demonstrations from the vendors, and we examined the vendors' hardware, software, and documentation to determine if their products did what their vendors claim that they do. That is, do they enable voters who use the touch screen voting system in use in the State of Maryland to verify that their votes were cast as intended, recorded as cast, and reported as recorded, and do they permit post-election auditing? We examined such issues as:

- implementation
- impact on current state voting processes and procedures
- impact on voting
- functional completeness
- security against fraud, attack and failure
- privacy
- reliability
- accessibility

We also compared these systems to one another and to the state's current voting system and procedures, which includes the SBE's use of parallel testing around that system.

We note several specific concerns about these products, including the following:

1. Only one of these products, the Diebold VVPAT, provides for a pure paper solution.
2. All of these products would impose significant one-time implementation and on-going management burdens (cost, effort, security, etc.) on the SBE and the state's 24 Local Boards of Elections.
3. All would increase the complexity of the act of voting.
4. All would increase the amount of time required to vote.
5. All would at least double the amount of effort required to administer elections.
6. All would adversely affect voter privacy.

7. These products would have both potentially positive and potentially negative impacts on security and election integrity.
8. None can be considered as fully accessible to persons with disabilities and none of them fully meets the accessibility standards of Section 508 of the Rehabilitation Act.
9. Integration of these systems will require the cooperation of Diebold to develop and/or ensure the viability of a working interface between the vendors' products and the Diebold system.

Our principal findings are, first, that each of the systems we examined *may* at some point provide a degree of vote verification beyond what is available through the Diebold System as currently implemented. But this is true only *if the system were fully developed, fully integrated with the Diebold DREs and effectively implemented*.

Our second principal finding is that none of these systems is yet a fully developed, commercially ready product. None of these products had been used in an election in the U.S. (SCYTL has been used outside the U.S. and a different version of the Diebold VVPAT has been used in the U.S.).

Were the State of Maryland to decide to acquire any of these products, the vendor would have to invest additional money and effort to produce an actual product and make the product ready for use in actual elections. Indeed, nearly all of these vendors are looking for some level of external support to fully develop and commercialize their products.

In our expert opinion, it is a bad idea for governments to buy products that are not functionally complete and that either do not have positive records in the market place or that cannot be fully and effectively tested in simulated elections to ascertain their performance characteristics.

Therefore, based on the evidence from this study, we cannot recommend that the State of Maryland adopt any of the vote verification products that we examined at this time.

We would note that no election system – regardless of the technology involved – is foolproof nor is any election system completely immune or secure from fraud and attack. Indeed, there is a long and inglorious history of election fraud in the U.S. that involves nearly all methods and technologies of voting, especially paper voting systems. Moreover, it would be prohibitively costly to make any election totally secure.

Finally, regardless of what the State of Maryland does in the near term with regard to vote verification and vote verification systems, in future elections, it should expand the use of parallel testing. The state should also undertake a full-scale assessment of the security procedures and practices around its current voting system. We say this even with the knowledge that current security procedures are reasonable and prudent and that the SBE's system of parallel testing, as currently implemented, reduces considerably the possibility of fraud and attack on the system.

Acknowledgements

To begin with, the study team wishes to thank Ms. Linda Lamone, State Administrator of Elections, and her staff, particularly Ross Goldstein, Deputy Administrator and Nikki Trella, Election Reform Director, for their assistance during the course of this study. They worked very closely with us throughout the study. In particular, they performed an important liaison function between the study team and the vendors, encouraged vendors to participate, provided information to the study team about a variety of issues around voting in Maryland and the roles and functions of the SBE and the LBEs, provided briefings and demonstrations and written documentation about the state's current voting policies and procedures, and, finally, provided useful feedback on our draft report.

We also wish to express our appreciation to Ms. Anne Roland (MIPAR and Public Policy) who edited this report and to Ms. Gay Warshaw (MIPAR) who typed and formatted it. They took the work of a group of disparate academics and made it readable and presentable.

As part of this study, all of the UMBC faulty investigators and research assistants (the study team) attended vendor demonstrations, briefings by SBE staff and regular study team meetings. They received and analyzed documents provided by vendors, SBE staff, and they reviewed scholarly literature relevant to this project. In addition, individual team members had specific responsibilities for various aspects of the study, including lead responsibility for initial drafts of sections of the draft and final project reports. The faculty investigators also directed and guided the research assistants in their assignments throughout the project. Finally, all of the faculty investigators and at least two of the graduate research assistants critically reviewed and provided feedback to the principal investigator on the draft and final reports. In all respects, this report is the product of the joint efforts of the entire study team.

Faculty investigators' responsibilities:

Prof. Norris, the project principal investigator, was in charge of overall project management. With Prof. Holden, he took lead responsibility for the report sections on implementation and impact on elections. He also guided research assistants in examining the literature around electronic voting. Finally, he compiled and edited the draft and final reports from materials provided by all of the team members.

Prof. Nicholas, co-principal investigator, participated in project management activities and in discussion and analysis of various aspects of the project and reviewed and commented on drafts of the report.

Prof. Sears, co-principal investigator, participated in project management activities and, with Mr. Law, conducted the accessibility analyses of each of the vote verification technologies.

Prof. Gangopadhyay, in conjunction with Prof. Karabatis, conducted the data management analysis for each of the vote verification systems.

Prof. Holden assisted Prof. Norris in project management and, with Prof. Norris, took lead responsibility for the report sections on implementation and impact on elections. Prof. Holden also developed description of the SBE's procedures for parallel testing.

Prof. Karabatis, in conjunction with Prof. Gangopadhyay, conducted the data management analysis for each of the vote verification technologies in the study.

Prof. Koru conducted the functionality testing of all of the systems, except parallel testing, and wrote the descriptions of the products from Prof. Ted Selker and SCYTL.

Mr. Law, in conjunction with Prof. Sears, conducted the accessibility analyses of each of the vote verification technologies.

Prof. Pinkston participated in discussion and analysis of various aspects of the project, particularly security procedures, and reviewed and provided comments on drafts of the report, particularly on the security-related matters.

Prof. Sherman was responsible for the analyses of election integrity (security) and privacy of each of the technologies and also wrote the description of the VoteHere product.

Prof. Zhang had responsibility for analysis and description of the VoteGuard product from Democracy Systems, Inc. However, because Democracy Systems, Inc., did not participate in this study, his work was not included herein. Additionally, Prof. Zhang reviewed and provided comments on drafts of the report.

Graduate Research Assistants

Markus Dale (Computer Science and Electrical Engineering) performed a static analysis of the source code of VoteHere reference library.

Kevin Fisher (Computer Science and Electrical Engineering) participated in discussions about the security analyses of the verification systems and created the figure showing the Voter Experience for VoteHere.

Rosa Heckle (Department of Information Systems) assisted in the literature review and analysis of privacy risks for vote verification systems.

Kenny Koo (Department of Information Systems) assisted in studying the security threats and countermeasures of DREs using voter verifiable paper trails.

Evan Perlman and Tonya Zimmerman (MIPAR and Public Policy) conducted a search for and review of literature on the subject of electronic voting. They drafted the sections in the report about *Electronic Voting in the U.S.*, *HAVA*, and *Issues around Electronic Voting*. They also reviewed and provided comments on the draft project report.

Mohammad Ammar Sheik (Department of Information Systems) assisted in testing the voting verification systems.

Undergraduate Research Assistant

Chiedozie (Fred) Unachukwu (Department of Information Systems) assisted in the review for electronic voting literature and in particular, electronic voting risks.

A Study of Vote Verification Technologies
Part I: Technical Study

Prepared for the Maryland State Board of Elections

Table of Contents

A Study of Vote Verification Technologies
Part I: Technical Study

Prepared for the Maryland State Board of Elections

Introduction

On August 19, 2005, the University of Maryland, Baltimore County (UMBC) on behalf of the University's Maryland Institute for Policy Analysis and Research (MIPAR) entered into a memorandum of understanding (MOU) with the State Administrator of Elections to provide a technical analysis of commercially developed vote verification technologies. This document constitutes the final report of that analysis.[1] This report constitutes Part 1, Technical Study, of a two part study conducted about vote verification technologies for the State Board of Elections. Part 2, Usability Study, was conducted by researchers at the University of Maryland, College Park (Herrnson, et al., 2006).

We conducted this study at a time when concerns about electronic voting, specifically, voting on Direct Recording Electronic systems or DREs – otherwise known as touch screen voting systems – and independent verification of voting on DREs, have become a focus of national attention. Over the past year or so, a nationwide rush to adopt a solution to the "problem" of touch screen voting appears to have occurred. Twenty-six or more states, for example, have adopted or appear to be in the process of adopting requirements for the inclusion of independent verification systems, nearly all based on the voter verified paper audit trail (VVPAT), for touch screen voting systems. The difficulty is that, at the moment, neither the problem (if there is one) of touch screen voting nor the solution (independent verification systems) is well understood.

The issue almost always raised is whether the touch screen voting systems are reliable. Do they record, store, and count each voter's vote as the voter voted it? Can they be corrupted? Can they be effectively audited?

The apparent favored solution is to require VVPATs to be attached to DREs. However, this option is even less well understood in terms of what it will accomplish in terms of vote verification and auditing, and what its negative consequences might be. Moreover, there are alternatives beyond VVPAT that may offer better solutions, and we discuss some of these alternatives in this report.

The reason that neither the alleged problems nor the apparent solutions are well understood is that no one has conducted the kind of analysis of independent vote verification systems that we have undertaken in this study. Indeed, previous studies have examined touch

[1] The State Administrator of Elections contracted with us for this study. However, we conducted the study completely independently. No one from the SBE influenced our research questions or objectives, methodologies, analytical techniques, findings, conclusions or recommendations although the staff of the SBE were given an opportunity to review and comment on the draft. The results of this study are those of the researchers alone. It is an independent scientific work.

screen systems and/or vote verification systems in isolation. In this study, we examine vote verification systems within the context of the processes and the conduct of elections. This study, therefore, should be helpful in answering some of the questions around touch screen voting and vote verification. The information in this report is intended to help the citizens of Maryland, members of the General Assembly, this Governor's Commission on the Administration of Elections, the State Board of Elections and the Governor in coming to informed decisions about how to proceed with the administration of elections in Maryland.

Our primary emphasis in this study of vote verification systems was to determine how effective each vote verification system is as a means for: (1) independent verification of the vote recorded on the Diebold AccuVote-TS voting system (the voting system used in the state of Maryland); and (2) creating an acceptable audit trail.

Scope of Work

This study includes:

1. A technical review of each of the participating vote verification systems specified by the Administrator of the SBE (see below). This review involved examination and testing of all hardware, software, documentation, and any and all other elements necessary for the verification systems to produce independent verification under actual election conditions;

2. A comparative risk analysis that examined how each vote verification system performed against selected review criteria relative to each other and relative to the Diebold AccuVote-TS voting system with no modification;

3. An analysis of the susceptibility to attack, fraud or failure of each of the vote verification systems;

4. An assessment of the accessibility (e.g., for individuals with disabilities, the elderly, etc.) of each vote verification system (independently, researchers at the Center for Politics and Citizenship at the University of Maryland, College Park conducted a usability analysis);

5. An assessment of the magnitude of effort and cost to implement and integrate each vote verification system with the current voting system and maintain the integrated system (e.g., on-going maintenance and support, cost to operate per election, etc.); and

6. An estimation of the impact, if any, of each vote verification system on the ability of voters to vote in the state's elections, on the State's current election procedures and on the ability of election officials, election judges and volunteers to perform their jobs in actual elections and to adapt, manage, and effectively use these systems.

For purposes of this study, the following terms have the meanings indicated.

> *Auditing* means the ability, through an alternative means and after the election is conducted, to establish that the votes recorded by the Diebold AccuVote-TS voting system correspond to the votes recorded by the independent vote verification system.
>
> *Vote verification* means the ability to independently confirm the accuracy of the Diebold AccuVote-TS voting system.

The systems for possible inclusion in this study were those from the following organizations and individuals: VoteHere (Sentinel); SCYTL (Pnyx.DRE); Prof. Ted Selker's Voter Verified Audio Audit Transcript Trail (VVAATT); Diebold VVPAT; Democracy Systems, Inc. (VoteGuard); IP.Com; and Avante. We also agreed to review the SBE's policies and procedures, including security procedures, around the Diebold touch screen voting system and the SBE's procedure of "parallel testing" of the Diebold AccuVote-TS voting system.[2] We used the Diebold touch screen voting system as currently implemented in Maryland as a baseline against which to evaluate the vote verification systems.

We determined that IP.Com did not meet the criteria of an independent vote verification system and, thus, did not include it in the study. The staff of the SBE contacted Avante, which indicated that it did not want to participate in the study.

Democracy Systems, Inc. (DSI) did not provide its system (VoteGuard) for inclusion in the study. UMBC signed non-disclosure agreements with each of the other vote verification system vendors to have access to their systems. However, DSI required that UMBC also sign a non-compete agreement in order to provide access its system. Non-compete agreements are against UMBC policy. Therefore we were unable to include DSI's VoteGuard system in this study.[3]

In the MOU, we agreed to provide the SBE with a draft report of our findings no later than December 15, 2005, and a final report within 30 days of receiving written comments on the draft from the SBE.[4] However, because there were considerable delays in our receiving access to the vendors' systems, we were unable to meet this deadline. In fact, we gained access to two of the five systems very late in the fall of 2005. We received access to the VoteHere system on November 16, 2005, and we received access to the final vendor's system (Diebold) on December 20, 2005.

[2] In conducting this review, we received briefings and a demonstration from SBE staff as well as written documentation. As extensive as this review was, it was not a full security analysis of the state's election system.

[3] As of a meeting between the president of DSI and UMBC officials on January 11, 2006, a compromise appeared to have been reached under which DSI will provide UMBC with access to the VoteGuard system for the purpose of testing it in the same manner as we tested the technologies reported herein (not including source code). However, this agreement occurred too late for VoteGuard to be included in this report. In the future, should we undertake an analysis of additional vote verification technologies, we intend to include VoteGuard in that study.

[4]We asked for this review because we wanted to provide the staff of the SBE the ability to read the report and provide comments on it to us before we completed the final report. However, we made it clear to the SBE, and the SBE readily agreed, that we retained full substantive and editorial control over the report.

Background

Electronic Voting in the U.S.

Voting methods in American elections have been called into serious question in recent years, specifically as a result of problems that occurred in the 2000 election in Florida (Cranor, 2003; Wang, 2004). This election dramatically brought to the attention of the public the possibility of errors with punch card voting systems. Punch card systems exhibited the now familiar problems of hanging chads and misaligned ballots (Beiler, 1989a; Cranor, 2003). Several older voting technologies (punch cards, optical scans and lever systems) have also been prone to undervoting (not voting in a race), overvoting (voting multiple times for one race) and misvoting (Cranor, 2003). One study noted a particular problem in the 1988 election in Florida, which had large numbers of no-votes (Beiler, 1989b). This was blamed on both placement of the race on the ballot and improper recording of ballots (Beiler, 1989b). Other studies have indicated that punch cards have the worst record of all voting systems in the last four presidential elections, whereas lever machines have the worst record in Senate and gubernatorial elections over the last three election cycles as regards unmarked, uncounted, or spoiled ballots (Cal Tech/MIT, 2001). According to this same study, touch screen voting systems have the second worst record of unmarked, uncounted and spoiled ballots in presidential, Senate and gubernatorial elections (Cal Tech/MIT, 2001).

Due to these reported problems with other systems and as a result of the issues surrounding the 2000 presidential election in Florida, there has been movement toward electronic or touch screen voting (Cranor, 2003). According to one study, the proportion of voters using electronic systems is expected to have increased from 13 to 29 percent between 2000 and 2004 (Wang, 2004, and Election Data Services, 2004). Touch screen voting systems are also popular because it is felt that the systems are easy to use, more accessible to persons with disabilities, better able to accommodate multiple languages, prevent overvoting, provide quick results (with less human error), and eliminate costs associated with printing ballots (Burmester and Magkos, 2003; Wang, 2004). A principal concern about touch screen voting systems is whether the underlying software of these systems can be trusted, especially whether the software can be trusted to record and count votes as cast by voters.

Help America Vote Act (HAVA)

The election controversies of November 2000 also prompted a response at the federal level. In 2002, Congress passed the Help America Vote Act (HAVA). This legislation attempted to bring voting procedures, which until that point had been the responsibility of individual state governments, under the purview of the federal government, (Kurlantzick, 2004).

The bill was designed to combat a host of issues plaguing the voting process. Through a mix of new guidelines, requirements, and federal programs and funding, this legislation provides assistance for states as they update and improve their voting processes. Among other provisions, it requires states to upgrade away from the older voting systems, in this case mainly away from lever and punch card systems, and toward new touch screen and optical scan systems (Holt, 2005).

HAVA also provided for the formation of the Election Assistance Commission (EAC), a federal body designed to promote the goals of the 2002 bill. Among other duties, the EAC was charged with helping the states successfully make the upgrade to new voting technology. The EAC would offer administrative and technical support, as well as provide grants to develop and test new election systems. It would also develop a program to test, certify and decertify election systems as they were introduced (EAC, 2005).

What raised the concerns of critics of touch screen voting in this instance was that HAVA does not provide guidelines for states regarding performance tests on the newly approved voting technologies (especially touch screen voting systems), nor does it contain a requirement for any sort of independent verification systems (Kurlantzick, 2004; Holt, 2005; Pynchon, 2005).

There are also problems with the implementation timetable as far as providing access for disabled voters. Although the law did not go into effect until January 1, 2006, some voting system vendors are selling systems now, to be used for the foreseeable future, which do not meet the access standards of HAVA (Pynchon, 2005).

To address these criticisms, Representative Rush Holt (D-NJ) introduced the Voter Confidence and Increased Accessibility Act (HR 2239 in the 108th Congress, HR 550 in the 109th Congress). According to Representative Holt's web site, the bill is chiefly concerned with a requirement that all voting machines produce a verifiable paper trail and a more general requirement for an "accessible voter-verification method." Finally, it addresses concerns raised by some security experts who warn of hacks and attacks if a voting system is ever connected to the Internet. Holt's bill prohibits such systems from being connected to the Internet or being attached to any non-secure communication device, (Holt, 2005). HR 550 has been relegated to a subcommittee, and it is unknown whether or how soon it will emerge, but it is important to note the bill because it encapsulates many of the concerns that lawmakers have about touch screen voting and limitations in HAVA's scope.

Issues Around Electronic Voting

Several potential or actual problems have been identified around touch screen voting. First, problems exist on an individual level that might affect elections. These problems include voter trust in the system, readability of the touch screen systems, problems with smart cards (which prevent persons from voting more than once), issues around instructions and assistance to voters, ability to write in candidates, issues concerning the ability to administer the system, and privacy (Herrnson, et al., 2005; Rubin, et al., 2005). An additional, and to many a greater concern involves the security of touch screen systems. Security issues including malicious programming, unintentional but nevertheless bad programming, equipment errors, system malfunctions including crashes, the inability to independently recount votes, and issues about the correct capture of votes (Burmester and Magkos, 2003; Cranor, 2003; Hall and Alvarez, 2004; Machlis, 2004; Rubin, et al., 2005; Selker, 2004; Wang, 2004; The Economist, 2002). To correct these security issues, calls have been made for open source coding (which would allow for independent examination of programming of electronic voting systems), voter verifiable paper trails (which could also be used for auditing), and active testing programs of the equipment and software (Beiler, 1989a; Brunvard, et al, 2004; Selker, 2004; Wang, 2004; The Economist, 2002).

Finally, the concerns about touch screen voting have propelled a number of states to enact legislation regarding independent verification. Twenty-six (26) states have adopted legislation requiring some form of paper ballot or paper trail to be used in their voting systems (see Appendix A: "Paper Solution States").

Maryland's Voting System[5]

In 2001, the Special Committee on Voting Systems and Election Procedures recommended to the Governor and General Assembly that a statewide voting system be implemented in Maryland. Subsequently, House Bill 1457 (2001) which required a statewide, uniform voting system for polling place voting and a uniform system for absentee voting was introduced into the General Assembly and adopted. The law did not specify the type of voting system, just that there be a uniform system. After the law became effective, and as the result of an open, competitive bid process, the State Board of Elections selected Diebold Election Systems, Inc., to provide a DRE voting system for polling place voting and an optical scan voting system for absentee voting.

This voting system was implemented in phases. Phase I counties (Allegany, Dorchester, Montgomery and Prince George's Counties) implemented the voting system for the 2002 elections. These counties were selected for Phase I because they used the oldest voting systems in the State: punch card for Montgomery County and lever machines for the others. The contract for Phase I was signed in 2002.

The Phase II contract was signed in 2003, and Phase II counties (which include the remaining counties but not Baltimore City) implemented for 2004 elections. Phase III, which includes Baltimore City, will be implemented for the 2006 elections. (Two other counties, Cecil and Caroline Counties, were eligible to implement in Phase III, but chose to implement in Phase II.) With the completion of Phase III, Maryland will have almost 20,000 DRE voting units.

Table 1 presents the costs of the implementation of the touch screen voting system in Maryland. It shows that by FY 2009, a total of about $95.5 million will have been spent by state government on this system. Of that amount, about $45.6 million will have been spent on hardware and maintenance and almost $50 million on a variety of necessary support services including security measures, warehousing, transportation, voter outreach, support services, technical support, testing of various kinds, and project management. This amounts to a state government cost of almost $2.82 for every Maryland resident (5.6 million) and just over $5.10 for every Maryland registered voter (3.1 million) per year for each of the six years. This cost includes providing every jurisdiction in the state with all of the equipment needed to conduct an election as well as some level of technical assistance and voter outreach.

[5] Information in this section is based on documentation provided by the staff of the SBE, 2006.

Table 1: Maryland Touch Screen Voting System Costs[6]

STATEWIDE COSTS*

	FY03	FY04	FY05	FY06	FY07	FY08	FY09	Total
Hardware Payments	2,131,933	5,034,100	9,654,582	8,142,292	6,412,403	6,411,015	6,409,343	44,195,668
Maintenance		0	0	449,880	1,509,318	1,628,890	828,455	4,416,543
Warehouse		237,797	237,215	348,166	321,996	318,874	316,674	1,780,722
Transportation		280,776	280,776	177,198	775,126	517,695	355,177	2,386,748
All Training (SBE/LBE/Judges)		341,271	56,271	47,877	183,624	111,385	88,633	829,061
Absentee Ballot Printing		72,023	72,023	178,222	214,462	116,476	110,428	763,634
Voter Outreach		1,000,000	354,394	500,000	444,120	50,000	50,000	2,398,514
Support Services		1,064,140	1,119,422	1,757,665	2,018,814	1,696,040	2,110,978	9,767,059
Total Services	914,704	2,996,007	2,120,101	3,009,128	3,958,142	2,810,470	3,031,890	18,840,442
Technical Support		1,497,263	1,529,513	3,520,456	5,207,019	3,607,748	1,529,145	16,891,144
Acceptance Testing		491,400	0	163,830	75,000	75,000	75,000	880,230
IV&V		380,000	191,673	1,000,000	1,200,000	600,000	500,000	3,871,673
Set-up/Breakdown		182,000	182,000	0	545,200	326,601	239,701	1,475,502
DRE Ballot Preparation		37,950	37,950	0	105,800	47,050	43,550	272,300
Project Mgmt		733,688	816,515	800,000	750,000	750,000	750,000	4,600,203
Total Optional Services		3,322,301	2,757,651	5,484,286	7,883,019	5,406,399	3,137,396	27,991,052
Total	**3,046,637**	**11,352,408**	**14,532,334**	**17,085,586**	**19,762,882**	**16,256,774**	**13,407,084**	**95,443,705**

*This chart assumes that the State continues to use the voting system without any additional verification methodology.

Notes:

1. Hardware total for Phase III which begins in FY2006 totals $7,570,750, financed over five years, estimated at $1.9 million/year.

2. FY2003, 2004, and 2005 were not altered since they occurred in the past.

3. Voter Outreach: This chart assumes that SBE will provide the brochure the LBEs pay for printing/distribution costs as needed. $50,000 is included for misc unexpected costs.

4. Diebold currently provides all services except for Project Mgmt and IV&V and Acceptance Testing (which are all provided through separate vendors). The Diebold contract is through half of FY09; the services will be rebid at that point.

5. If the 2006 primary is moved forward, the bulk of the training costs will occur in FY06 rather than FY07.

[6] Provided by the staff of the SBE, January 2006.

Election Administration in Maryland

The State of Maryland has a dual election system in which the SBE and the local boards of election (LBEs) share authority and responsibility for administering elections in the state. The following description was provided by the staff of the SBE (January 2006).

The State Board of Elections is a five-member body charged with managing and supervising elections in the State and ensuring compliance with election laws. With the advice and consent of the State Senate, the Governor appoints the members to staggered four-year terms. No more than three members may be from the same political party. The law for board membership requires no specific knowledge or training. The Governor may remove a member for incompetence, misconduct, or other good cause. A member of the SBE may not serve more than three consecutive terms.

The State Board appoints the State Administrator who may be removed for incompetence, misconduct, or other good cause by a vote of not fewer than four members. As the State's chief election official, the State Administrator is charged with oversight of the State Board functions as well as with supervising the operations of the local boards.

The primary responsibilities of the State Board and the Administrator are:

- Manage and supervise elections in the State
- Maintain a central voter registration database
- Oversee the local election management system
- Maintain an electronic campaign finance reporting system
- Maintain a records management program for State and local election records
- Direct, support, monitor, and evaluate the activities of each local board office
- Develop a program of instruction for election judges and oversee its implementation locally
- Select and certify a statewide voting system, in consultation with local boards
- Adopt regulations/uniform guidelines governing voter petitions, canvassing of the vote, recount procedures, and absentee voting
- Conduct biennial meetings with local boards and staff
- Maintain a web site, including prompt and accurate posting of election results

Each jurisdiction in Maryland (23 counties and Baltimore City) has a local board of elections (LBE). Under Maryland's Election Law Article, the local board and its staff are subject to the direction and authority of the State Board and are accountable to the State Board for its actions in implementing the requirements of Maryland's election law.

The LBEs generally consists of three regular members and two substitute members. Two regular members and one substitute member are registered voters of the same political party as the Governor of Maryland, and one regular member and one substitute member are registered voters of the other principal political party. These members are appointed by the Governor and confirmed by Senate of Maryland (except for

Caroline, Dorchester, and Kent Counties, for which confirmation is made with the Maryland House of Delegates).

The local boards are responsible for overseeing the conduct of all elections held in its county and ensuring that the election process is conducted in an open, convenient, and impartial manner. The local boards also have the authority to appoint an election director and counsel, serve as the local board of canvassers and certify the results of each election conducted by the local board, establish precinct boundaries, determine polling places, assign voters to precincts, and provide information about voter registration and elections.

The election director in each jurisdiction manages the operations and supervises the staff of the local board, trains election judges, provides notice to voters about voter registration and elections, follows procedures issued by the State Board relating all aspects of election administration, verifies petitions, and conducts a canvass for each election conducted by the local board.

Study Methodology and Evaluation Criteria

In conducting this study, we examined both non-technical and technical aspects of the vote verification technologies. The non-technical aspects included implementation of these systems within the framework of the current Maryland election policies and procedures and the impact of the vote verification technologies on current Maryland election procedures and on voters and voting in Maryland elections. The technical aspects included: data management, reliability, accessibility, election integrity (security) and privacy. We describe our study methodologies and evaluation criteria below.

Implementation

Assessing implementation impacts from adding vote verification systems to the existing base of the Diebold Accuvote TS units is difficult for several reasons. First, the lack of maturity of many of the vote verification products makes it difficult to estimate impacts, especially one-time and on-going costs. Because none of these products is available in the U.S. market at this time as functionally complete solutions, pricing discussions between study team members and vendor representatives often devolved into oral commitments by the vendors "to be competitive." Second, several of the verification systems vendors have requested access to the proprietary software that runs the Diebold Accuvote TS to estimate integration costs. Understandably, since Diebold has its own solution for providing vote verification, it has little business incentive to cooperate with its competitors to integrate their verification solutions with Diebold's DREs. As a result, cost estimates for integration are speculative due to lack of data.

This analysis relies on data from a variety of sources. In response to requests from the SBE prior to the commencement of this study and subsequent requests for data from the study team, vote verification system vendors provided data on estimated costs of implementation. Vendors also met with the study team on at least one occasion to demonstrate their vote verification solution. There are also some cost estimates from other organizations, most notably the Asian Pacific American Legal Center, which provided an estimate of the impact of voter verified paper trail submitted to the California Assembly in 2004. In some cases, the study team was left to make assumptions about how the state would implement a verification solution since the vendor did not have a "standard" configuration. Any state, including Maryland, might make different configuration decisions based on on-going risk analyses, changes to policy and process, and availability of resources to mitigate identified risks.

We gave the products rankings between 1 and 5, with 5 being the easiest and least costly to implement within current Maryland election technologies and procedures, and 1 being the most difficult and most costly.

Impact on Current Election Procedures and on Voters

Here, we examined the ways in which adding vote verification technologies to the current Maryland system of elections would affect that system and could affect voters, exclusive of usability and accessibility (which are addressed elsewhere). As noted below, the impacts are numerous and some are potentially significant. Moreover, as we discuss with respect to each verification technology, some have more and more serious impacts than others.

The addition of *any* vote verification system or technology to the Diebold AccuVote-TS system that Maryland currently uses in elections *will at least double the complexity* of the act of voting and the administration of elections. Some vote verification systems will add even greater complexity.

This complexity will affect voters who, if they so choose (or may be required), will have to review (and, in some cases, verify) the results of the votes that they cast on the touch screen as displayed on the vote verification equipment and also verify their votes on the touch screen.[7] This adds somewhat more time to the act of voting and adds from one to several more steps to the voting process for the voters.[8] For some, perhaps many voters, this will also add a degree of confusion that is independent of any potential equipment malfunctions or failures that may occur on the verification systems. Finally, it will undoubtedly result in more voters (perhaps significantly more) requiring assistance from election judges.[9]

[7] In the precincts that were observed in Las Vegas in the Nevada election in 2004 that used VVPATs, something less than 40 percent of voters actually did the verification. That is, less than 40 percent actually looked at the printer screen, compared what it displayed to the DRE, and touched the DRE to indicate that they had verified. The rest did not look at the printer at all and simply touched the screen to say that they had verified. (DVD provided by the Los Angeles County, CA, Register/Recorder 2004)

[8] Because of the absence of standardization among tested systems, researchers at the University of Maryland, College Park (UMCP) conducting the usability study were unable to compute the additional amounts of time required for voters to cast ballots using the vote verification systems. This was because of time, cost and vendor cooperation issues, and it meant that none of the vote verification systems could be integrated with the Diebold AccuVote TS system for the study. Thus, all we can say is that each one adds some increment of time to the act of voting. (Herrnson, et al., 2006).

[9] Researchers at UMCP noted that a substantial fraction of voters requested help when voting when independent verification systems were employed in field tests (Herrnson, et al., 2006).

The addition of any vote verification system will also affect election administration on both a one-time and on a continuing basis. O*ne time impacts* include at least the following:

1. The SBE will have to acquire up to 20,000 vote verification devices (the actual number of devices may be greater) to install onto and interface with the Diebold DREs;

2. The SBE will have to contract with some entity to make the necessary physical adjustments to the DREs to accept installation of the devices;

3. The SBE will have to contract with some entity to write the interface between the DRE and the verification device (except for the VVAATT) so that the device will receive the necessary and correct information or instructions from the DRE for taking the actions for which the device is selected, e.g., printing a vote for the voter to verify. This will involve both cost and time;

4. The SBE will have to securely transport the devices to the LBEs;

5. The LBEs will have to securely store the devices;

6. The SBE will have to provide training for the LBE officials (and volunteers) in the operation of the devices, including training in how to connect them physically to the DREs, how to address issues of device malfunctions and failures, and how to assist voters under various problem scenarios;

7. The SBE will have to develop security procedures regarding the devices, promulgate those procedures and train the LBE officials in those procedures;

8. The LBEs will have to implement these security procedures;

9. Periodically, the SBE will have to monitor the implementation of security procedures by the LBEs and take corrective action when deficiencies are found; and

10. The SBE will have to provide at least initial (if not also on-going) technical assistance or support to the LBEs regarding all aspects of the implementation, management, security, and use of the devices involved.

11. The SBE and the LBEs will have to undertake extensive initial voter outreach and education about the new technology.

Continuing impacts include at least the following:

1. The devices will have to be serviced and maintained both during an election and between elections;

2. The devices will need to be securely stored between elections;

3. The devices need to be securely transported from storage to the election precincts prior to an election and returned immediately after an election;

4. The SBE will need to provide continuing training to the LBEs in all aspects of managing and using the devices and in the security around them;

5. The SBE will need to provide continuing technical assistance or support to the LBEs regarding all aspects of the device technology;

6. In setting up for an election, the LBEs will have to physically install the devices onto the DRES and will have to securely manage all of the physical elements associated with the devices (e.g., printer rolls, flash cards, audio tapes);

7. During an election, the LBEs will be responsible for the operational functionality of the devices on all DREs and will have to address (replace or repair) all malfunctions;

8. During an election, the LBEs will be responsible for assisting voters with using the devices on all DREs and will have to have knowledgeable staff and volunteers available in all precincts who can instruct and assist voters in how to use the devices, who can address technical problems that will inevitably occur with the devices, and who can address any voter's belief that the device and the DRE recorded different voter selections and the possibility that the device and the DRE actually recorded different selections; and

9. When closing an election, the LBEs will have to remove the devices from the DREs, remove any physical elements associated with the devices and the physical elements associated with them (e.g., printer rolls, flash cards, audio tapes), secure the devices and the physical elements, transport the devices and their physical elements to the LBE office; inventory and store the physical elements securely.

10. The LBEs will have to understand how and be able effectively to use the technologies in the event of a recount.

11. The SBE and the LBEs will have to undertake a continuing program of voter outreach and education about the verification technology.

At a minimum, these actions will at least double the amount of effort required for election administration. For some vote verification technologies, the added effort will be greater than one order of magnitude. Since most of these costs are "soft" – that is, they involve personnel and do not, for example, involve purchasing hardware or services – they are difficult to estimate. But they certainly would increase the cost of election administration considerably. Consider the following examples. First, the SBE and the LBEs will have to train all election judges[10] to: understand and be able to use the selected vote verification technology; be able instruct voters to use the technology; understand and be able to correct problems that may occur with the technology during an election; in case of malfunctions or failures, take the technology off line and replace it with functioning technology; also in case of malfunctions or failures, address what happens to the voter's vote on both the touch screen and the vote verification technology (i.e., did it get recorded and counted as intended by the voter); at the end of an Election Day, remove and securely handle both the vote verification technology and all physical elements associated with the technology (e.g., paper rolls, memory cards, etc.); and much more. Second, the LBEs will have to securely transport the vote verification devices and all physical elements associated with them to and from their voting precincts for Election Day, and between elections they will have to provide secure storage for the devices and all physical elements associated with them.

In our examination of the impacts on elections of the individual vote verification technologies later in this study, we use this generic discussion of impacts as a baseline. That is, at the minimum, all of the vote verification technologies will have at least these impacts on elections and voters. We examine and discuss only additional impacts on elections that each technology may produce.

We gave the products rankings between 1 and 5, with 5 signifying that a technology had the least adverse impact on voters and election administration, and 1 signifying that it had the greatest adverse impact.

Data Management

Regardless of how well a file system is designed and implemented, it is always vulnerable to hardware failures, malicious attacks, and or natural disasters that may render it inaccessible or useless. It is therefore, important to devise plans which ensure that the data will remain accessible or restorable if the system is compromised.

Table 2: Data Management Issues

Event	Problem
Voter casts vote	
Voter verifies vote as cast in DRE screen	
Voter verifies vote as cast in verification system	
Vote recorded in verification system	Inconsistent records in DRE and verification system
Verification system flags an error	Can the vote be aborted in DRE?

[10] Election judge is the formal title for election poll workers in Maryland.

When a voter casts his vote, it must be recorded in a buffer (temporary location) in the computer memory (possibly at the time the voter is verifying the vote) before it is actually committed to a persistent data file under the control of the operating system. In the case of a failure in the operating system or a power loss after the time verification is done by the voter but before the data is written in a persistent data file, the vote in the main memory will be lost. Given that the DRE and the verification systems are independent of each other there are several conceivable scenarios when this problem of "lost vote" may take place.

We explain the above with an example of events shown in Table 2. Events such as those noted in the table may occur due to systems failures resulting from natural causes. In this example, the vote is recorded in the DRE but not in the verification system. Hence, a possible way to recover from this situation is to discard the vote from the DRE and start over again, asking the last voter to re-vote. Such actions would require the DRE to be able to delete the last vote but not the other votes cast prior to that. The other associated problem is that the verification system must convey the message to the DRE that it is unable to record the correct vote. Any automatic means for doing this would require a signal back to the DRE, entailing collaboration between DRE and verification system.

Events such as the above may occur due to systems failures which do occur as a result from natural causes. However, it is possible to envision a "denial-of-service" attack that systematically causes one of the two independent systems (DRE or the verification system) to flag a failure after a pre-determined or randomized number of votes. This will make the entire system fail caused by an adversary who gained control over just one of the two systems. Thus, the addition of a redundant system may jeopardize the very advantage it was designed for.

Atomicity

Atomicity refers to a data management issue related to recording the votes cast. In the voting booth, the DRE and verification system comprise a two-unit distributed system. Each of the two units contains a repository of "votes cast" (in a form of persistent memory). In such system, a voter casts a vote as a transaction spanning the DRE and the verification system. The two repositories should remain "consistent," that is, each one should record the same vote cast by the voter. Each vote in the DRE should be present in the verification system, to ensure that both DRE and verification system have recorded the same vote. This property, called atomicity, as a feature of the voting transaction, refers to recording the vote in both DRE and the verification system or in neither of them. It is also known as the "all-or-nothing" property. If atomicity is not enforced, the two repositories may not record the same votes, leading to a discrepancy in the vote counts between the DRE and the verification system, i.e., exhibiting the undesirable phenomenon of "lost" or "not recorded votes." Therefore, if votes are not correctly recorded in both systems, the DRE vote count (total number of votes recorded in the DRE) would not equal the verification system vote count (total number of votes recorded in the verification system).

A sample scenario exhibiting a case of not recorded vote: A voter in the booth votes in the DRE, and the vote is transmitted to the verification system. The voter then proceeds to check the vote on the verification system, ensures that the verification system has recorded the correct vote, and then walks away. However, the voter did not go back to the DRE to complete the transaction by pressing the "Final OK" or “commit” button on the DRE screen, so the DRE does not record that vote.

This way, a vote was recorded on the verification system, but not on the DRE, leading to discrepancies between the total counts on the DRE and the verification system. And vice versa, there can be similar scenarios where a vote would be recorded in the DRE but not in the verification system. The result is the same: When atomicity is not enforced between DRE and verification system, there may be not recorded votes in either the DRE or the verification system. Such scenarios could occur several times with different voters, on the same or several machines during an election!

A solution that guarantees atomicity calls for the collaboration between the DRE and the verification system during the integration phase of the two units (DRE and verification system) of a distributed system.

The above problems have been addressed in various areas in the Database and Operating System community. There are solutions to the above problems, and they require specific steps to be taken in both DRE and verification systems. As of December 2005, these steps are completely absent in the versions of the systems that we evaluated (with the exception of the Pnyx.DRE system from Scytl, although still not complete), therefore the problems created by the absence of atomicity still exist.

To implement such solutions, modifications have to be made to both DRE and verification system systems. In addition, there is a need to integrate both systems and some of the above problems can be solved during the integration process. By “integration process” we mean the existence of Application Programming Interface (APIs) that will contain functions which can be called by both DRE and verification system systems. This entails modifications to software at the DRE and Verifier systems.[11]

Therefore, in this study we examined each system for its data management capabilities and ranked them between 1 and 5 depending on the extent to which they achieved a desirable state of data management. We adopt the following criteria from the guidelines for security, accuracy, recovery, and integrity provided by the FEC voting system standards, based on their relevance to data storage and management.

1. Single-point failure: The FEC voting standards recommends that the systems should “protect, by a means compatible with these Standards, against a single point

[11] It is important to note that atomicity is independent from security. Possible solutions on atomicity may have an effect on the security of the system and vice versa. That is, solutions implemented for the security of the system may have an effect on data management.

of failure that would prevent further voting at the polling place." We refer to this as single-point failure. This can occur unless the system is capable of rolling back individual votes in case of a problem.

2. Atomicity (lost votes): "Protect against the failure of any data input or storage device." We refer to this as atomicity (lost votes). This can occur unless a two-phase commit protocol ensures atomicity of votes across the DRE and the verification system. The two-phase commit protocol requires a two-way communication between the DRE and the verification system in order to check consistency of voting records at the level of individual votes. If there is no communication between the two systems this cannot be implemented in any way. Simply displaying the votes to the individual voters, either by electronic or paper media, does not guarantee consistency in the way votes are recorded internally in the systems.

3. Tamper-proof data entry: "Protect against any attempt at improper data entry or retrieval." We refer to this as tamper proof data entry. This could be caused by an adversary deliberately attempting to tamper with the voting system. The inclusion of a separate verification system adds to the vulnerability of the overall system against this type of attack. Without a two-phase commit protocol between the DRE and the verification system it is possible to launch a programmatic attack by getting hold of any one of the two systems (either the DRE or the verification system).

4. Event recording: "Record and report the date and time of normal and abnormal events." This is an important functionality for auditability of the system. However, this may also raise privacy concerns. We refer to this as event recording. This is only possible if the events are recorded with a timestamp in the verification system.

5. Exception recorder: "Detect and record every event, including the occurrence of an error condition that the system cannot overcome, and time-dependent or programmed events that occur without the intervention of the voter or a polling place operator." We refer to this as the exception-recorder, which is useful in auditing and maintenance of the system. This is very similar to the previous case with the exception that this refers to error conditions only.

6. Self-management: "Include built-in measurement, self-test, and diagnostic software and hardware for detecting and reporting the system's status and degree of operability." We refer to this as self-management.

Reliability

During an election, failures of vote verification systems can lead to considerable loss of time and effort and can also lead to lost or improperly counted votes. In addition, such failures would prevent the vote verification systems from fulfilling their original purpose, which is to increase our confidence in the voting process.

Therefore, it is desirable to understand the reliability of the alternative vote verification systems. At the time of this study, only three of the systems had reached an operational level where it was possible to make a judgment about their reliability. These three systems were Diebold's AccuVote TS, Prof. Ted Selker's Voter-Verified Audio Audit Transcript Trail (VVAATT) system, and SCYTL's Pnyx.DRE system. For reasons discussed below, we were unable to evaluate the VoteHere system for reliability. Our reliability assessments for these three systems were subjective assessments on a scale of 1 to 5 as shown in Table 3.

Table 3: Reliability Assessment

Ranking	Reliability Assessment
0	No assessment could be made
1	Very frequent failures
2	Frequent failures
3	Occasional failures
4	Infrequent failures
5	No failure

To assess reliability, we tested the high-level usage scenarios for the systems. Our test cases covered the main functionalities provided by the vote verification systems. Obviously, there may be additional and more detailed failure scenarios for these systems. Testing cannot prove the absence of errors. However, it increases our confidence level.

For the remaining system, VoteHere, the vendor did not provide us with an operational system that fulfilled its promises. The system lacked functional completeness. The equipment provided by VoteHere was a bare prototype that lacked operational (or functional) features and was only developed to give the "look and feel of the intended system for demonstration purposes." For example, the system provided did not have the capability to store votes. More importantly, there was no web site set up for the purpose. These are important omissions, and because of them, there is no way to verify the claims made by VoteHere or to test for the reliability or usability of the VoteHere system.

Functional Completeness

In our evaluation, we used an additional measure, *functional completeness,* which refers to the degree that a system actually provides the functions given in its specifications and allows users to run those functions. Functional completeness provided by the vote verification systems is an important criterion for their acceptance. If a system does not actually provide the adequate functions it originally promised, reliability becomes less observable.

Functional completeness is important because all of the vote verification systems that we examined, except VVAATT, involve software. Complexity is an inherent characteristic for software systems (F. P. Brooks, 1987). Many researchers previously observed that software projects suffer from complexity and failure risks are high for software projects compared with those in other manufacturing and engineering disciplines (J. F. P. Brooks, 1995). The software products often require a substantial field (operational) testing and maintenance period to eventually become robust systems (Musa, 1998).

To assess functional completeness, we used a five-scale ranking scheme shown in Table 4. It is a subjective assessment of the ratio of system functions actually implemented and made available to the users to those promised in the system specifications.

Table 4: Assessment of Functional completeness

Ranking	Functional completeness (Usable functions)
0	No assessment could be made
1	Nothing
2	Less than half
3	Half
4	More than half
5	Complete

Accessibility

To examine the accessibility of the four vote verification devices we used the standards of Section 508 (www.section508.gov) of the Rehabilitation Act of 1973, as amended in 1998, for procurement of electronic and information technology (Authority: 29 U.S.C. §794d.). Specifically, we used the following subsections of the 508 Standards (see Appendix B):

Subpart B – Technical Standards
Self contained, closed products
Subpart C – Functional Performance Criteria
Subpart D – Information, Documentation, and Support

We evaluated the four devices in a laboratory environment at UMBC. This was an expert review based on accepted standards. The user-based assessment regarding use by individuals with disabilities is part of the study undertaken by researchers at the University of Maryland, College Park.

An accessibility expert conducted the evaluation, judging each of the applicable Section 508 criteria for each device. We then assigned a summary evaluation to each device based on a scale from 1 (hardly accessible) to 5 (very accessible). In the Section 508 criteria, Subpart B §1194.25 (a) requires that "Self contained products shall be usable by people with disabilities without requiring an end-user to attach assistive technology to the product. Personal headsets for private listening are not assistive technology." We evaluated the voting systems as self-contained products and, therefore, the assistive technology clauses in Subpart C §1194.31 "Functional performance criteria" were not appropriate.

We ranked the systems according to the following scale:

Table 5: Accessibility Evaluation

Ranking	Accessibility
1	Hardly accessible at all
2	A little bit accessible
3	Somewhat accessible
4	Accessible
5	Very accessible

Security and Privacy

To evaluate each vote verification and audit system for security and privacy, we consider how well each system mitigates threats to election integrity, voter privacy, and reliable operations. We report our findings in two ways: a narrative of the strengths and weaknesses of each system, and a summary comparison table.

In our evaluations, we focus primarily on election integrity and voter privacy. We concentrate less on resistance to malicious disruption, because most voting systems (and most information systems) are highly vulnerable to such disruption. For example, an adversary could disrupt an election with a bomb, bomb threat, fire, dispersing toxic substances, physically tampering with machines, or disrupting electrical power. Most of these threats must be mitigated with procedures beyond the verifier. Corrupt DREs or verifiers could try to disrupt elections through maliciously generating incorrect displays, results, or printed outputs.

How well a system enhances election integrity can be completely separate from how well the system protects voter privacy. Therefore, we make no attempt at combining these two characteristics into a single security and privacy rating. Also, because it is impossible to compare security and privacy in any simple strictly quantitative objective fashion, our summary comparisons necessarily have a strong subjective qualitative element. The narratives explain the basis of our comparisons.

In the rest of this section we summarize specific evaluation criteria and metrics for election integrity and vote privacy. In Appendix C, we discuss further issues regarding the methodology for examining security and privacy.

Goals

There are three main security and privacy goals of any voting system: (1) to give voters high confidence in the integrity of the election results, (2) to preserve voter privacy to the maximum extent possible consistent with law, and (3) to reduce chances of disruptions to the election as much as reasonably possible.

> *Election integrity* means that each voter casts his ballot as intended; the system records the ballot as cast; the system tallies the votes as recorded; and the Election Boards certify the results as tallied. Thus, an adversary, must not be able to modify the election results. Election integrity is central to giving voters confidence in election results and, for many people, is the most crucial characteristic of any voting system.
>
> *Voter privacy* means that no one (besides the voter) can learn how the voter voted. Many voters view privacy as a fundamental right that is vital to prevent coercion, vote selling, and bribery. If a system provides the voter with a receipt to take home, the receipt must not reveal how the voter voted.
>
> *Resistance to disruption* means that it is difficult for an unintentional or malicious adversary to cause a delay, rescheduling, or stoppage of the election process. Such denial-of-service attacks are virtually impossible to prevent entirely for any type of voting system, even for an unsophisticated yet determined attacker.[12]

A major difficulty in conducting elections is to achieve both election integrity and voter privacy. It would be much easier to achieve election integrity without voter privacy (e.g., by posting all marked ballots on a public bulletin board).

Summary Evaluation Criteria

In our summary evaluations for security and privacy, we subjectively rate each system with a score on a five-point scale (1=poor, 5=excellent) by various criteria. The evaluations are for the entire composite systems, when used in combination with Maryland's existing Diebold voting system and the security policies and procedures, including parallel testing, that have been adopted around it. Thus, the score for parallel testing refers to Diebold with parallel testing, and the score for Sctyl refers to Diebold with Sctyl and parallel testing. The scores provide a useful relative comparison of the security and privacy characteristics of the study systems.

We organize our summary evaluation criteria in two groups: goals and means.

[12] We do not discuss resistance to disruption in the product analyses included in the section below entitled "Evaluation of Vote Verification Technologies." Instead, see Appendix I.

Evaluation Criteria (Goals)

- election integrity
- voter privacy
- resistance to disruption

We do not condense these criteria into a single overall criterion, because scores for these three high-level goals can be independent. Some of the important factors influencing these scores are the following.

Evaluation Criteria (Means)

- votes cast as intended
- votes recorded as cast
- non-reliance on complex DRE hardware and software for integrity
- physical security
- protection of removable data devices
- proper use of cryptography
- sound key management
- software implementation best practices
- transparency
- prevents correlation of votes

Transparency measures the extent to which the system architecture, details of operation, and software are clear and publicly available for inspection. It is a well-accepted principle of security engineering that security ought to be based on the fundamental strength of the design, not on the obscurity of the implementation details. This point, however, is separate from the issue of whether or not such design and implementation details are publicly released.

Publicly releasing all security details and software can have two-sided implications. On the one hand, such release enables the public to scrutinize the system and detect any possible weaknesses. Also, programmers tend to write better software when they know that other people might closely examine their software. On the other hand, releasing such details can help potential attackers, if the system has any imperfections: such details are needed for most attacks, and releasing the details eliminates the attacker's challenge of obtaining them. Even if releasing security details and software might have a net negative effect on security, there might be a political advantage in doing so to increase voter confidence in the system. Regardless, technical experts will still be needed to inspect and validate the security of any voting system.

Evaluation Metrics

In carrying out our evaluations, we informally considered a number of metrics that measure the levels of protections offered by a voting system. Most security evaluations cannot simply be reduced to exact quantitative measures. Nevertheless, the following metrics are useful in understanding and comparing the security and privacy properties of the systems, even if informally applied.

Attack Metrics

- number of conspirators required
- number of votes affected
- number of machines, precincts, local election boards affected
- cost (in dollars)
- time to carry out attack
- computer resources needed (computer time, memory space)
- probability of detection
- probability of success
- required knowledge, skills, and equipment
- required level of sophistication (Level 1, 2, 3 – see discussion of threats)

The greater the number of conspirators required, the more difficult it is to mount the attack, and the greater the chance that the attack will be detected. Attacks that affect large numbers of votes can have a larger influence on some elections, and risk greater chances of detection; in close elections, however, affecting a relatively small number of votes might be sufficient to change the outcome. Overall cost in dollars is a useful summary metric of how much effort is required for an attack to succeed.

Limitations of the Study

Like any study, this one has certain limitations. The first is that our examination was limited to parallel testing and four vote verification the systems. Two vendors who were invited to submit their systems to this study, Democracy Systems, Inc. (VoteGuard) and Avanti, chose not to participate. Policy makers in Maryland should also be aware that there are additional, promising vote verification technologies that may deserve consideration in the future.

We were also limited to the specific hardware and software provided us by the participating vendors. Some of the systems, as should be clear from our analyses, were incomplete. Second, we were able to examine the source code only for VoteHere and Pnyx.DRE. VVAATT does not employ software, so this was not an issue. But we were not able to examine the Diebold source code. Third, two of the vendors provided their systems to us very late, making it very difficult for us to undertake as rigorous an examination of the systems as would have been optimal. Fourth, none of the systems that employed software (all but VVAATT) provided complete applications for examination (in part because complete software applications are not available from all of the vendors). Finally, none of the systems was integrated into a Diebold AccuVoteTS system. Thus, we could only estimate how they would function if so integrated.

Nevertheless, and even considering these limitations, we believe that this study should help decision-makers in Maryland to better understand these vote verification systems and to make more informed decisions about them and about the state's current election system.

Vote Verification Technologies

In the following paragraphs, we describe the SBE's method of parallel testing and the four vote verification technologies that we examined in this study. Members of the study team wrote descriptions of the technologies based upon viewing demonstrations of them, reading documentary materials provided by the vendors, communications with vendors and expert observation of the operation of these systems. We sent our initial descriptions to the respective vendors and asked for their feedback in order to ensure that the descriptions were accurate. The vendors all cooperated and provided comments on and corrections. We incorporated those comments and corrections, as we felt appropriate. The purpose of these descriptions is to provide the reader with as clear a portrait as possible of the technologies and how they would work in an election.

Some descriptions are longer than others. This is because some systems are more complex than others and required greater prose to describe them. The relative length or brevity of a description does not imply a subjective judgment about or endorsement or lack of endorsement of a system by the study team.

Parallel Testing

While not identified as a vote verification system, *per se*, the State of Maryland has instituted a process, with the support of the Maryland League of Women Voters (LWV), to provide assurances to the public and SBE that the touch screen voting process accurately counts votes as cast.

In discussions with SBE officials, they asserted that parallel testing demonstrates the accuracy and integrity of the vote recording and counting processes of the Diebold TS Accuvote units. Background documentation provided by staff of the SBE defines "parallel testing" as "method of testing an electronic voting unit by producing an independent set of results that can be compared against the results produced by the voting unit."

The SBE sponsors parallel testing at two stages in the election process. In both cases, LWV volunteers vote paper ballots, which are photocopies of real ballots that could be used for an absentee or provisional voting produced by the Diebold GEMS servers. Because these are copies of real ballots, though, the optical scan units used for recording and tallying absentee and early ballots will neither accept nor read these facsimile ballots. What is parallel about the process is that the LWV volunteers are split into three groups of two individuals. One of the volunteers of Group A calls out the votes recorded on the paper ballot. A volunteer from Group B records each vote and tallies them manually. At the same time, a member of group C uses a Diebold Accuvote TS machine chosen randomly by the volunteers to record and tally the votes. Both members of Group C cross-checking the tally between the manual and machine vote each ten votes to ensure the accuracy of the manual count. At the end of the parallel test, the volunteers compare the tallies from the hand and machine count. Both parallel tests use the same configuration of hardware, software and ballot definition used in the election. In neither case are the results of the parallel test uploaded to a GEMS server.

The Code of Maryland Regulations (COMAR 33.10.02.16) requires that all counties' LBEs conduct public demonstrations of election technologies prior to an election. In 2004, six counties incorporated the parallel test plan described herein as part of their public demonstrations. In 2005, the SBE amended the regulations to require all 24 counties to conduct the parallel tests at their public demonstrations. The public demonstrations use Diebold Accuvote TS machines that have completed logic and accuracy (L&A) checks[13] and are ready to be used in the upcoming election. A LWV volunteer selects one machine and then individuals attending the public demonstration vote 50 paper ballots, which are then recorded and tallied through both a manual and machine process. This typically takes place within ten days of the election. At one of the public demonstrations, the LWV also chooses four Diebold Accuvote TS machines that will be used on Election Day for the parallel test described below.

Second, on Election Day, the SBE hosts a test that the LWV administers. This parallel test uses four Diebold units that the League selected (and continue to be sealed from the L&A test) from a LBE, which was Montgomery County for the 2004 general election. The volunteer from the LWV records the serial number and seal number for each of the four units, which an employee of the SBE then transports to the SBE for secure storage until Election Day. Before Election Day, League of Women Voters vote 300 paper ballots per voting unit. These paper ballots are copies of absentee ballots made immediately after the LWV selects the units. On Election Day, volunteers check to confirm that the four units being opened for the count were the same units selected on the public demonstration day based on the log of serial numbers and seal numbers from the four machines. The parallel test then proceeds as described above at the SBE office. For the Election Day parallel test, the SBE videotapes the voting unit screen in the event there is a need to discover any mistakes.

[13] Logic and Accuracy (L&A) is the entire process of configuring, testing, and certifying the accuracy of voting equipment for an election. The process involves setting up an election using AccuVote-TS machine and all the related peripherals so that the technology is ready to be used on Election Day in the precincts. L&A is performed on the GEMS Application Server, AccuVote-TS and AccuVote-OS units to ensure: 1) each unit is fully functional and free from mechanical problems, 2) each unit contains the appropriate ballot styles available for voting in each vote center, 3) results are tabulated accurately on each election component, and 4) results can be transmitted accurately to the election database on the GEMS Application Server. The process to achieve these goals includes: 1) downloading the election to an AccuVote-TS PC Memory Card and AccuVote-OS Memory Card 2) testing of the AccuVote-TS unit and AccuVote-OS units, 3) uploading L&A test results, 4) generating L&A test results reports, 5) resetting election equipment for an election, 6) conducting a public display, 7) preparing the AccuVote-TS units for delivery, and 8) securing the AccuVote-TS and AccuVote–OS units until Election Day.

VVAATT

Professor Ted Selker's Voter Verified Audio Audit Transcript Trail (VVAATT) is a simple and easy to install system that requires little pre-election work. Basically, it includes a voice-operated tape recorder (VOR) that is connected to DRE (see Figure 1). VVAAT is delivered to the polling place packaged in a sealed box. The box contains documentation, which is available when two election judges unpack it. The documentation consists of three pieces of paper; a packing list, a voters' guide, and a election judges' guide. The election judges mark a checklist to assure that the batteries and tape are factory sealed, insert them into the VOR, connect the provided labeled audio cable to the marked input of the VOR, and the earphones to the marked output of the VOR. However, before this, the system should first be tested by running the output of a test microphone into the input of the VOR and by recording the names of the two people that set it up and the number of the seal that was used to seal the VVAAT. Then the VOR is put back into the VVAATT and its record tab is pressed, and it is secured with a screwdriver.   The VVAATT is then plugged into the output of the DRE. The election judges then sign the checklist to include the seal number on it and place it in voting materials transfer container. They must also write down the seal number in the polling

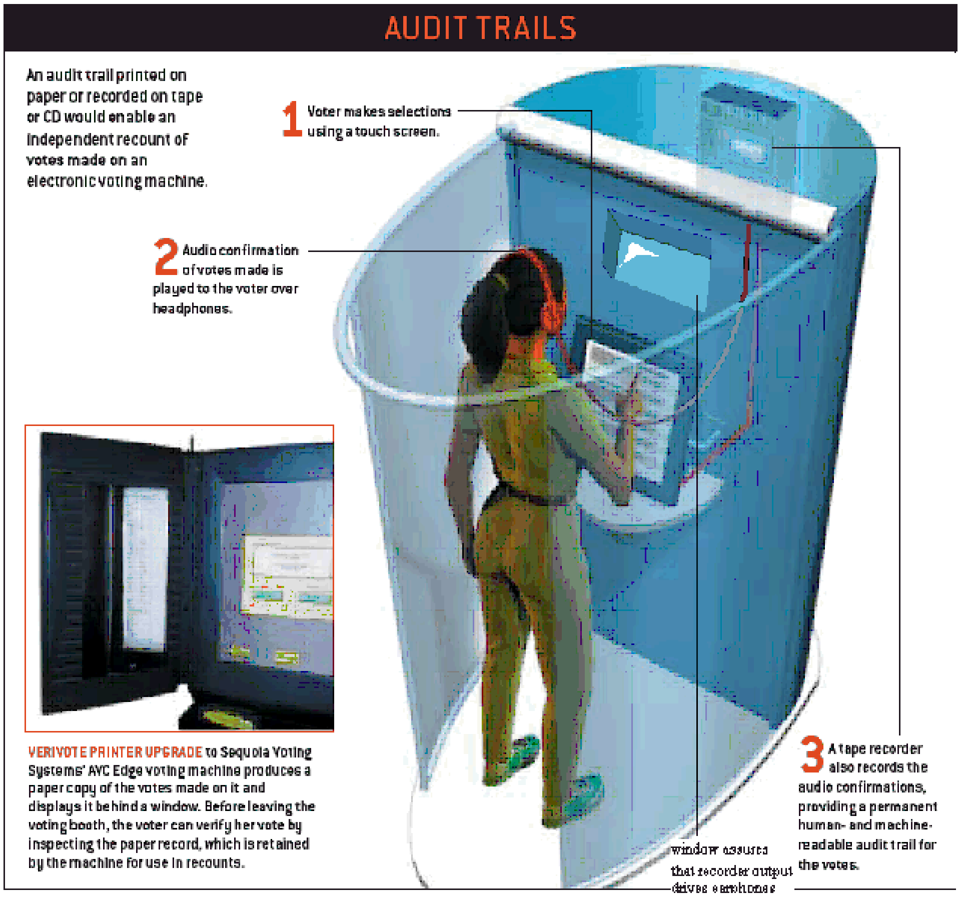

An audit trail printed on paper or recorded on tape or CD would enable an independent recount of votes made on an electronic voting machine.

1 Voter makes selections using a touch screen.

2 Audio confirmation of votes made is played to the voter over headphones.

VERIVOTE PRINTER UPGRADE to Sequoia Voting Systems' AVC Edge voting machine produces a paper copy of the votes made on it and displays it behind a window. Before leaving the voting booth, the voter can verify her vote by inspecting the paper record, which is retained by the machine for use in recounts.

window assures that recorder output drives earphones

3 A tape recorder also records the audio confirmations, providing a permanent human- and machine-readable audit trail for the votes.

Description of VVAATT in use (provided by Prof. Ted Selker).

place ledger and sign that together to assure a permanent record that the VVAAT was ratified as set up correctly. Spare VVAAT recording devices must be on hand to replace any that fail in each voting place.

The voter steps into the voting booth and puts on the provided headphones. Then, the voter makes selections using the touch screen of the DRE. Every activity of the voter is spoken to the voter by means of an audio feedback mechanism in the DRE. Each time the voter makes a selection (or deselects); an audio confirmation message is given to the headphones through the tape recorder's successful recording of it. For example, when the voter selects candidate A, the DRE says "selected candidate A," or when the voter unselects candidate B, it says "unselected candidate B." As the voter listens to the audio feedback, the audio feedback is sent to the portable tape recorder and recorded on the tape. During this procedure, the voter verifies his vote through headphones. He votes on the touch screen, and can verify that whatever he chose on the screen is the same as what he hears. If the two forms match, then the vote is verified through the touch screen and cast.

Everything that a voter hears is also recorded on the VOR. Whenever there is an output from the voting machine, it activates the VOR and starts recording. When there is no output, the VOR stops. It restarts whenever there is again an audio signal.

At completion of a voting day or if a VVAAT fails, the following procedure is used to secure the tape. Two election judges must be present to corroborate the event. The screws on the VVAATT are removed. The VOR is taken out and the record tab on the tape is pressed out. The seal is affixed to the audiocassette with a label that bears the precinct number and the names of the people that removed it. A record of the tape is put in the polling place ledger and signed by both election judges. The audiocassette is then sealed in its cassette holder and placed in the polling place election materials transfer container. Batteries are removed and discarded.

If the election results are challenged or if an audit team wants to audit the results, the audit team listens and counts the votes for each race and candidate until all the tapes used in voting are played back. Then, the tallies are calculated and compared to those obtained from the DRE. (Although the vendor did not provide this option for us to examine, it is also possible to audit the results from the audiotapes using computer technology and software.)

Pnyx.DRE

SCYTL's Pnyx.DRE consists of a verification module that is a simple hardware device composed mainly of a screen, an audio output (for voters with hearing impairments) and two buttons, a "CONFIRM" button and a "CANCEL" button (see Figure 2). The vote verification module enables voters to verify and confirm their votes before casting them. The device is also equipped with an internal memory for safely storing votes confirmed by voters, together with an integrity register of all these votes. Finally, the device has sufficient computing capability for the execution of the cryptographic protocol that protects the confirmed votes and fulfils the basic security requirements for electoral processes; that is, integrity of the results and voter privacy.

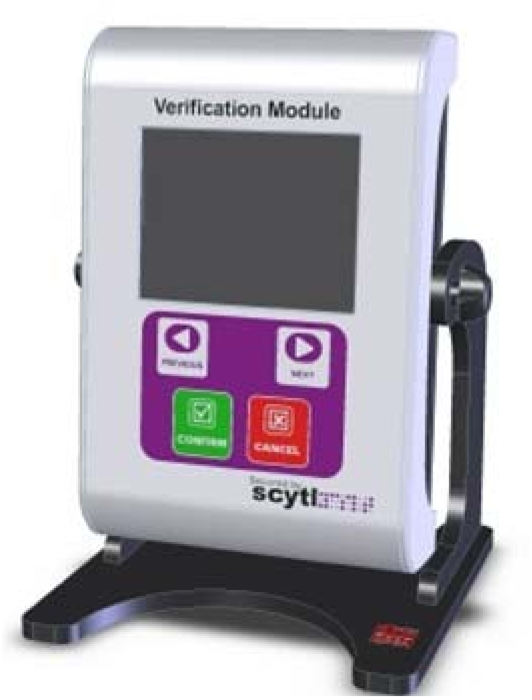

SCYTL’s Pynx.DRE verification module (provided by SCYTL).

The vote verification module is connected to the DRE to allow voters to verify that their desired votes will be accurately cast and recorded. The cryptographic protocol implemented in the vote verification module protects the votes against any potential manipulation afterwards (e.g., in the DRE voting terminal). Essentially, the voting process is as follows. First, the voter interacts with the DRE voting terminal, and, after viewing the different options on the terminal screen, he or she selects the desired voting options. Then, the voter verifies visually whether the options selected in the DRE match those transferred to the vote verification module. If this is the case, the voter confirms (with the “CONFIRM” button). Otherwise, the selection process is restarted (with the “CANCEL” button) in the DRE. Blind and visually impaired voters can complete this process by means of a headset connected to the audio output of the vote verification module and the touch pad attached to the vote verification module.

The vote verification module digitally encrypts the verified vote, signs, and stores it. It also keeps track of an integrity record of the verified votes. After each verified vote, the integrity record is updated in the vote verification module.

After the election, during the vote counting process, an integrity record is calculated from the votes stored in the DRE. Then, the integrity record obtained from the DRE is compared with the integrity record of the vote verification module. If the two integrity records match, this means that the votes in the DRE were not changed from what the voters voted on the DRE. If the two records do not match, then decrypting the votes stored in the vote verification module can perform a parallel recount and can compare them to the votes stored in the DRE. To decrypt the digitally encrypted and signed votes, a certain majority of election authorities (i.e., LBEs) should provide their portion of the election key. This key is generated before the election during the election initialization process, divided into pieces, and given to the election authorities. Comparison of the votes in DRE and those in the vote verification module is automated via software, which makes auditing faster and less error-prone.

VVPAT

The Diebold voter verifiable paper audit trail option (VVPAT) consists of a small printer encased in a sealed take-up unit housing that attaches to the side of the DRE (see Figure 3). The take-up unit housing has a glass front or screen approximately 2 inches by 4 inches. At the beginning of the election, a new spool of paper is loaded into the printer, and the printer unit is physically locked. The printer receives a direct feed from the DRE through an integrated serial databus (RS232) port.

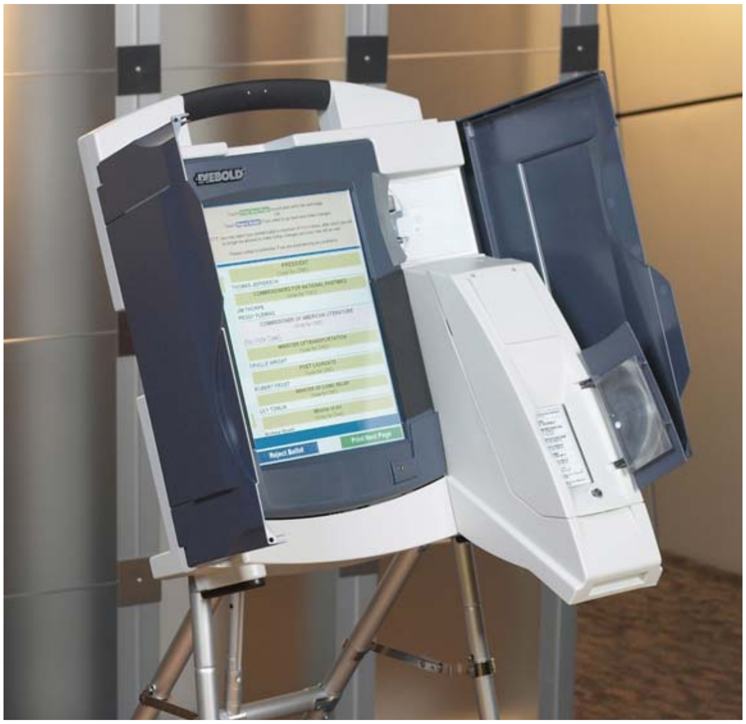

Equipment shown is Diebold's AccuVote TSX with printer, not the AccuVote TS that is used in Maryland (provided by Diebold).

The voter enters the voting booth just like he or she would in DRE voting system. The voter initiates the voting process in the same manner as would otherwise be required by the DRE system. When presented with the ballot, the voter touches the DRE touch screen to indicate his or her choices for each race or ballot question. When the voter is finished selecting his or her choices, the DRE displays a Summary Screen with the voter's selections and requires the voter to print the ballot in order to begin the verification process prior to recording the votes cast. Upon touching the "Print Ballot" button on the touch screen, the DRE sends the voter's selections to the printer. The printer then prints the voter's selections onto the printer paper in intervals no larger than the size of the window and displays them on the printer tape in the glass-covered window on the front of the printer unit. The DRE gives the voter the option to proceed to the next window of printed choices, or to reject the ballot and then modify his or her selections. When the voter has printed all selections and is satisfied that all selections displayed on the touch screen correspond to the selections printed on the printer paper, the voter can select the "Cast Ballot" button on the touch screen. The DRE then accepts and records the voter's choices and stores the data in the standard manner. The printer scrolls the paper up on the spool to prevent the next voter from seeing the previous voter's choices and stores the paper record in a sealed canister. At this point, the DRE is ready for the next voter.

At the end of the Election Day, election officials in each precinct open the devices and remove the paper rolls that contain the printed votes. The officials then secure the rolls (in an appropriate manner) and transport them (securely) to the county election office for secure storage. Each paper roll is associated by means of unique identifying numbers with the DRE from which it received and printed the voting results. In the event of the need to conduct an audit or recount, the votes recorded on the paper rolls can be hand tabulated.

VoteHere

The following description of the VoteHere product is lengthier than the descriptions of other products examined here. This is because the VoteHere product is the most complex product that we examined, and a more lengthy description is necessary to portray this product accurately. The length of the description is not, nor should it be seen as, any endorsement by the study team of this product. Appendix D provides some additional technical details and Appendix E discusses the experience of voters and election judges with the VoteHere system.

VoteHere offers a system that can be used in conjunction with the Diebold system to provide independent verification and audit capability. Specifically, the VoteHere system allows each voter (1) to verify in the voting booth that his vote was cast as intended, and (2) to verify after the election that his vote was correctly tallied in the total. As is also true for Pnyx.DRE and the VVAATT system, VoteHere can also be used with any other type of DRE. This description of the VoteHere system is based on VoteHere technical documents (Green 2003, Neff 2005, VoteHere 2005) and communications with VoteHere staff.

If the voting system did not have to protect voter privacy, these two goals could be easily achieved simply by posting each voter's name and vote on a public bulletin board for anyone to examine. VoteHere modifies this simply strategy to achieve privacy by posting the encrypted (sealed) ballot of each voter on a public web site without the voter's name. In the voting booth, the voter receives a printed copy of his encrypted ballot to take

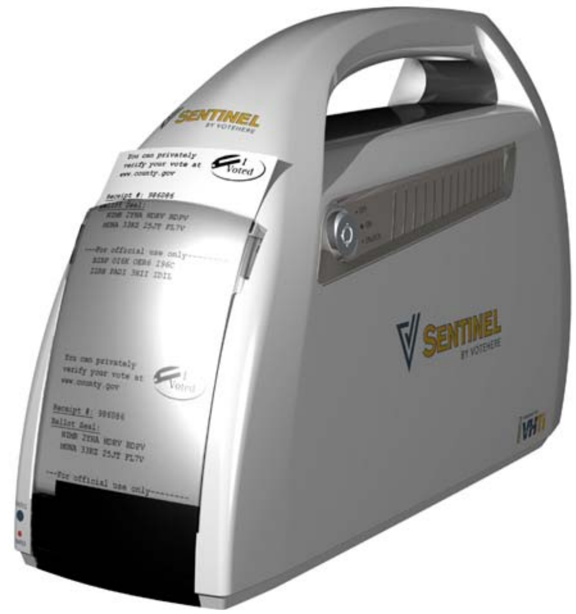

VoteHere's Sentinel verification module (provided by VoteHere).

home, so that he can later check that a copy of his encrypted ballot is listed on the web site. In the voting booth, the encrypted ballot is partially opened to convince the voter that it indeed represents his intended vote.

The voter experience in the voting booth works as follows. The voter receives a voting ticket from an election judge, selects ballot choices on the DRE, scans the ticket into the verifier, checks the DRE display and ticket against the printed receipt, and presses the accept-vote or reject-vote button on the verifier. If the voter accepts the vote, he takes his receipt and walks out of the booth (he may keep or throw away the ticket). If voter rejects the vote, he summons an election judge.

In the voting booth, the voter engages in a verification process (sometimes referred to as a "challenge-response protocol") producing a receipt that the voter takes home. This process intends to convince the voter that his vote was cast as intended. After the election, using his printed receipt and information posted on a public web site, the voter can easily verify that his vote was recorded in the official election data as cast in the voting booth (specifically, the voter checks that the web site has a copy of his receipt). The voter, or anyone, using trusted sophisticated mathematical software can also perform steps that intend to verify that the official tallies are consistent with the official election data (these steps check that these data are consistent with posted log files from the post-voting election computation performed by election officials). The receipt does not reveal how the voter voted. The voter experience in the voting booth is slightly more complicated than for other systems under study—about as complicated as an ATM banking transaction. What goes on internally, however, is moderately complex and based on some complicated mathematical cryptography. The figure on p. 31 describes the VoteHere voter experience.

Figure 1

VoteHere: The Voter Experience

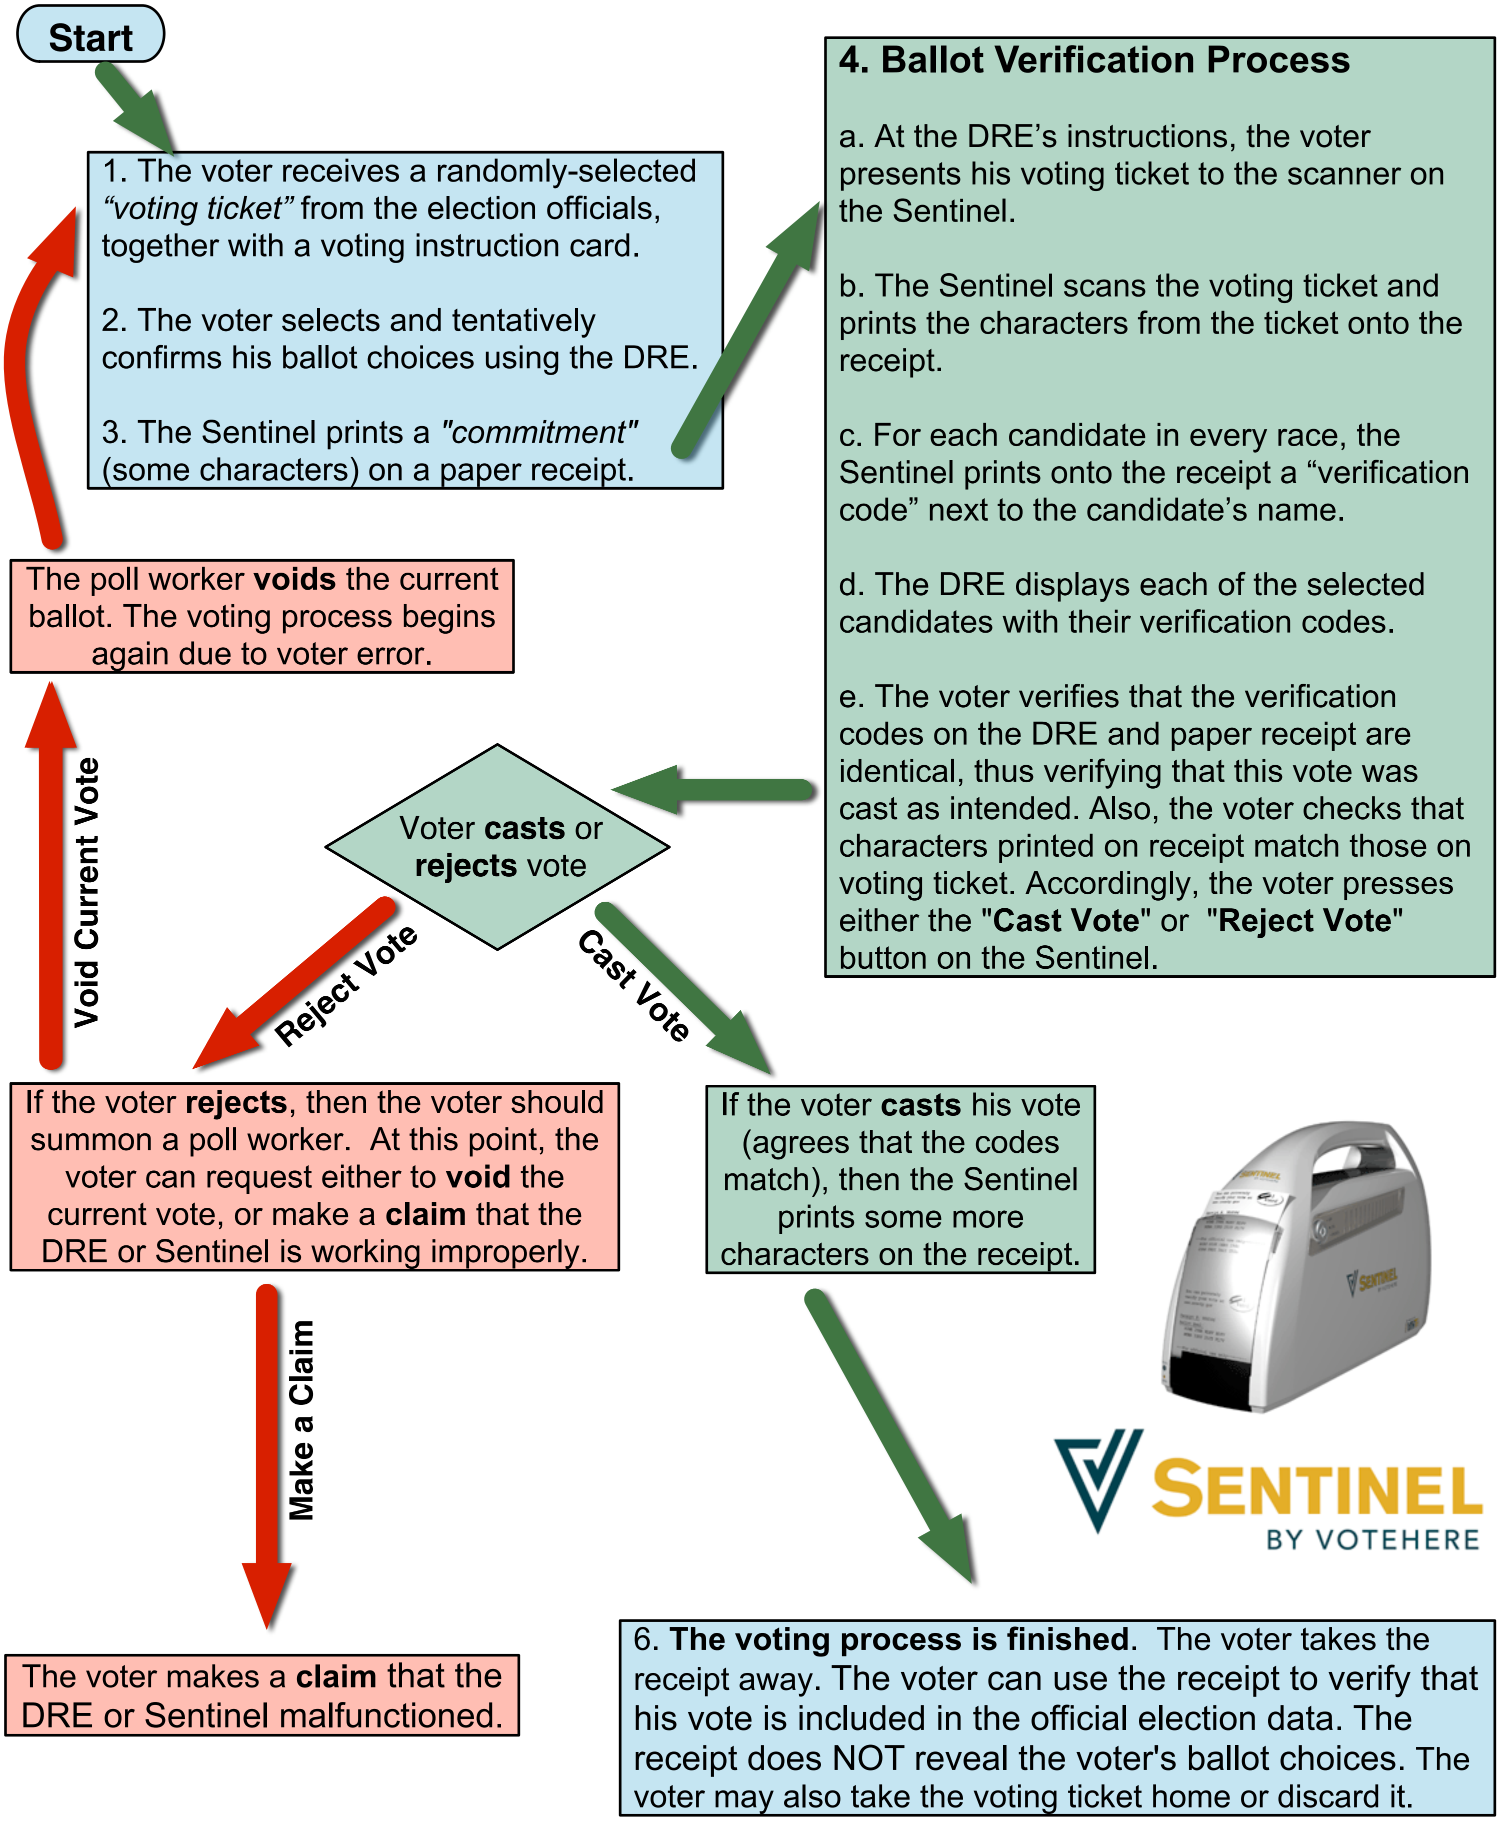

1. The voter receives a randomly-selected *“voting ticket”* from the election officials, together with a voting instruction card.

2. The voter selects and tentatively confirms his ballot choices using the DRE.

3. The Sentinel prints a *"commitment"* (some characters) on a paper receipt.

4. Ballot Verification Process

a. At the DRE’s instructions, the voter presents his voting ticket to the scanner on the Sentinel.

b. The Sentinel scans the voting ticket and prints the characters from the ticket onto the receipt.

c. For each candidate in every race, the Sentinel prints onto the receipt a “verification code” next to the candidate’s name.

d. The DRE displays each of the selected candidates with their verification codes.

e. The voter verifies that the verification codes on the DRE and paper receipt are identical, thus verifying that this vote was cast as intended. Also, the voter checks that characters printed on receipt match those on voting ticket. Accordingly, the voter presses either the "**Cast Vote**" or "**Reject Vote**" button on the Sentinel.

Voter **casts** or **rejects** vote

If the voter **rejects**, then the voter should summon a poll worker. At this point, the voter can request either to **void** the current vote, or make a **claim** that the DRE or Sentinel is working improperly.

The poll worker **voids** the current ballot. The voting process begins again due to voter error.

The voter makes a **claim** that the DRE or Sentinel malfunctioned.

If the voter **casts** his vote (agrees that the codes match), then the Sentinel prints some more characters on the receipt.

6. **The voting process is finished**. The voter takes the receipt away. The voter can use the receipt to verify that his vote is included in the official election data. The receipt does NOT reveal the voter's ballot choices. The voter may also take the voting ticket home or discard it.

The most distinctive feature of the VoteHere system is that, at the costs of higher voting complexity and complicated mathematical underpinnings, the system provides so-called “end-to-end voter verification of election integrity.” This means that, from the challenge-response protocol, receipt, and information posted on the public web site, the voter can verify for himself that his intended vote was correctly tallied in the total. This verification is based on the ballot verification process (challenge-response protocol), characters printed on the receipt, and information on the public web site; it does not require trusting VoteHere software or hardware. To the extent any type of trust is required, the nature of the required trust is very different and much less than that for any other of the examined systems. As such, VoteHere offers a different and stronger verification capability than do systems that depend on the voter simply trusting the verification system or its management of audio or paper tapes. No other system examined provides end-to-end voter verification of election integrity.

“End-to-end voter verification” of election integrity means that the voter can convince himself that his ballot is cast as intended in the voting booth, recorded as so cast in the official election data posted on the web site, and correctly tallied as so recorded in the official results. “End-to-end” means there is no “gap” of trust, such as trusting election officials to store paper records securely, or trusting the verifier to record in its memory the same vote displayed on a screen. Whereas many verification systems are based on trusting computer system security, VoteHere is based on trusting mathematical cryptography. Any voter who is not a cryptographic expert must trust that not all experts are wrong who assert that the cryptographic mechanisms of VoteHere work properly. Voters must also trust that they have some reliable way to examine the official published election data (e.g., by browsing the public web page). To preserve voter privacy, as is true for all verification systems under study, the voter must trust VoteHere software and hardware.

This section describes the version of VoteHere considered in the Maryland study. Defining this version is difficult for three reasons. First, the system is rather complicated. Second, the system is under development, undergoing changes, and can be specialized and configured in different ways. Third, the physical system delivered to UMBC for examination does not include all the functionality of the system defined for this study.

The VoteHere system is based on its *VHTi* validation and verification technology, as delivered through a configuration of the *Sentinel* physical device, with supporting *VHTi Management Tools.* The management software is used for initialization and post-voting processing, including computing information for the public web site. Although VoteHere currently supplies the Sentinel, VoteHere would welcome other vendors to become their hardware suppliers; this business strategy explains in part why VoteHere views the Sentinel as configurable in various ways.

All of the VHTi software is publicly available from the VoteHere web site for anyone to examine. The software currently publicly available is a general reference library. The software that would go into the product defined here that calls this reference library is not yet fully written. VoteHere states that it will make available all product software on its web site when it is completed.

The election process works in four stages: election initialization, ballot preparation and casting, ballot tabulation, and election verification. The VoteHere system assumes there is a so-called "Board of Election Trustees," to oversee the process, which in Maryland would be the Local Board of Elections (LBE). We now summarize the system components and how they are configured, vendor claims, and the four stages of the election process organized from the separate perspectives of the voter and election officials.

System Components and Configuration

The product comprises two parts: a Sentinel device that attaches to the Diebold AccuvoteTS DRE, and management software that runs on a dedicated election audit computer systems at the LBEs and Maryland State Board of Elections (SBE). The Sentinel includes a printer, scanner, and removable memory stick. One Sentinel is required for each DRE.

Although the version demonstrated at UMBC was not so equipped, we assume the system for the Maryland study has the following additional components: cast-vote/reject-vote physical buttons, USB port for specialized interfaces for voters who desire them (e.g., disabled voters), and an optional audio headset. Such specialized interfaces would communicate the same information printed on the receipt, and they could issue cast-vote/reject-vote commands. Although available as an option, we assume the Sentinel does not have a screen display.

As so configured, the VoteHere system makes no attempt to preserve consistency between the DRE and Sentinel. Instead, the DRE is viewed as an input device. A vote is not cast until and unless the voter presses the cast-vote button. The official tally is computed by VoteHere.

Integration with the DRE requires a change in the DRE's ballot display software. Thus, the SBE would have to negotiate with Diebold to make such changes. A unidirectional cable goes from the DRE (serial port) to the Sentinel (serial port). In the configuration described for the Maryland study, the DRE will not know if the voter presses the reject-vote button.

There are two types of outputs from the Sentinel: a paper receipt given to the voter to take home with him, and the audit log written on the removable memory stick. This audit log comprises basic election information, precinct identification, and a copy of each receipt. After voting, LBE officials download election data from the memory stick to their election audit computers. Voters who use specialized interfaces could choose to take home their receipt in an alternate form (e.g., Braille receipt).

Receipts

For the VoteHere system in our study, each voter will engage in the challenge-response protocol in the voting booth and receive a receipt.

Alternate configurations of the VoteHere system not considered here give the voter a choice between a so-called “short receipt” and “long receipt,” where the type of receipt significantly affects voting complexity and level of election integrity provided. We do not consider this alternative configuration for several reasons: (1) Giving each voter a choice complicates the system for voter, election officials, and policy makers, because all must understand the implications of each option to make an informed decision. (2) If all voters receive the short receipt, no voter can prove that his vote was correctly tallied as intended (other than trusting VoteHere). (3) There might be profound issues of user acceptance and constitutional law if different voters receive different types of receipts.

Vendor Claims

In our security evaluation we will examine the following vendor claims.

1. While in the voting booth, each voter can verify that his vote was cast as intended.

2. Using his receipt and the public web site, after the election, each voter can verify that his vote was recorded in the official election data as cast in the voting booth.

3. From the public web site, after the election, anyone can verify that all votes recorded in the official election data were correctly counted.

4. The system preserves voter privacy. In particular, the receipt does not show how the voter voted, even when the voter is untrustworthy.

5. All votes cast on each DRE are independently monitored and recorded in real time.

6. In comparison with other systems, this VoteHere product greatly reduces required trust in any system, software, and procedures. The information printed on the receipt and posted on the public web site is crucial to this claim: what matters is what information is printed or posted; it does not matter how it arrives there.

Our full analysis of vendor claims is in Appendix F.

Evaluation of Vote Verification Technologies

In the following section, we discuss each of the four vote verification technologies that we examined in this study per the following criteria: implementation, impact on elections, data management, functional completeness, reliability, accessibility, election integrity and privacy. We discuss the systems in the following order. First, while not a vote verification technology, *per se*, we provide an evaluation of the SBE’s process of parallel testing as a baseline. Next in order, we discuss: Prof. Ted Selker's VAATT; SCYTL's Pnyx.DRE; Diebold's VVPAT; and VoteHere.

Parallel Testing

Implementation. We assume that the State of Maryland will continue to conduct some form of parallel testing in the future based on the experience from the 2004 election. As a result, costs are assumed to be in the baseline budget for Maryland SBE. These costs are minimal, though, consisting of the cost of transporting the selected DREs from the LBE site to SBE and back. Storage costs for the parallel test units are already incurred in the budgets of the LBEs. The LWV provides the majority of the human resources to administer the parallel test, with SBE incurring some cost to set up the parallel test location prior to the election. Additional costs for each election include purchasing videotapes to record the parallel test.

There is no integration impact on or cost to the state.

Therefore, we awarded parallel testing a 5.0 for implementation.

Impact on Elections. While there is clearly some impact on the SBE to administer the parallel test process and this impact may become greater if parallel testing is expanded, there is no impact on an election, per se. This is because while the parallel test is conducted on Election Day, it is conducted completely independently of the election. Therefore we award it a 4.0 for impact on elections.

Data Management. In the case of parallel testing, the issue of single point failure does not arise because the two systems are not connected in any way. Hence the score for parallel testing has been assigned NA, which stands for "not applicable".

On atomicity, we have assigned a score of NA because there is no electronic communication between the DRE and the Verifier. For the same reason we assign the same score on tamper-proof data entry. Regarding event recording it is theoretically possible to record the date and time of normal and abnormal events by comparing the two systems on a vote-by-vote basis, hence we assign a score of 5.

Parallel testing precludes exception recording because the error can take place in only one of the machines, which will not be recorded in the other. The functionality required for self management does not exist in parallel testing because the DRE and the verification system are not connected. Hence there is no automated manner in which consistency checks can be performed across the two systems and it receives a score of 1.2.

For parallel testing and the four vote verification systems we examined, please refer to Appendix G for more detailed evaluation scores for functionality and reliability.

Reliability and Functional Completeness. Although we did not examine parallel testing for reliability or functional completeness, it would have the same properties in this regard as the Diebold DRE system alone. Moreover, parallel testing does not affect elections and, therefore, does not impact voters. As a result, we gave it a 5.0 for both reliability and functional completeness.

Accessibility. Not applicable. We did not examine parallel testing for accessibility because it does not affect elections and, therefore, does not impact voters in terms of accessibility of voting technology. As a result, we gave it an NA (not applicable) for accessibility.

Election Integrity (Security). Parallel testing aims to detect widespread improper operation of DREs by carefully testing a sample of DREs randomly selected immediately before the election. The testing is performed on Election Day in a fashion that attempts to simulate exactly the true election experience. This testing adds significant value, provided the sample is indeed selected at random, the selected machines are not modified, and the selected machines cannot detect (e.g., through signaling or perception of differences between the true and simulated election conditions) that they have been selected for parallel testing.

As noted in the following section, additional procedures involving the loading of software and storage of election system equipment also aim to reduce the chance of DRE corruption.

Unless parallel testing can be subverted, testing a sufficient number of DREs gives high confidence that there is not widespread corruption of DREs. See Appendix H for a mathematical formula that computes the probability that parallel testing will detect a bad DRE. For example, if 50% of the 19,000 DREs in Maryland were corrupt, then testing only 10 DREs would detect corruption with 99.9% confidence. However, if 19% of DREs were corrupt, then testing 10 DREs would detect corruption with only 9.6% confidence; testing 100 units would increase this confidence to 63.5%.

The main limitation of parallel testing is the danger that the bad DREs might be able to detect that they have been selected for testing. There are many subtle ways in which this might happen. For example, the DRE might be able to detect a difference in the voting rates or procedures, or an adversary might be able to signal the selected DRE, for example, by touching the screen in a special way. A few conspirators who are authorized to vote on the selected machines might be able to touch the bad DRE screens in simple special ways that would be undetected by the videotape log. With access to key cards or voter access cards, an insider might be able to signal bad DREs via choice of data on those cards (e.g., if the 7th number of the key is 7, behave properly). Consequently, great care is needed to secure the selected DREs, to ensure that conditions of parallel testing exactly match those of the real election, and that the officials carrying out parallel testing are trustworthy. Although it might be difficult to signal bad DREs without detection, this threat must be considered.

Another limitation of parallel testing is that it will not detect compromised DREs that behave badly only when signaled to do so. Signaling a large number of DREs to act badly, however, would require a large conspiracy, increasing the chance of detection. Therefore, a more likely threat would be to compromise all DREs (e.g., through a slight modification of the operating system) and then to signal only those DREs selected for parallel testing to behave properly

.

For these reasons we gave parallel testing a 3.0 for election integrity.

Voter Privacy. Parallel testing does not affect voter privacy because no official votes are cast on the tested DREs. Although parallel testing perfectly safeguards voter privacy, parallel testing does not correct any potential vulnerabilities to voter privacy created by Diebold. This point of view explains why, in the summary evaluation of voter privacy, parallel testing earns the same reference score we assign to the basic Diebold system. Thus, we awarded a 3.0 to parallel testing for voter privacy.

Additional Observations. Parallel testing, in combination with the usual election procedures, is the "baseline" for comparison of the verification-audit systems under review. We offer a few selected observations about how elections are currently conducted in Maryland.

1. At the end of the election, each DRE prints a Totals Report showing the total number of votes for each candidate cast on that DRE. (This report has nothing to do with the paper roll printed by the Diebold VVPT verification system.) These receipts are tallied by hand at the local election board and compared with the tallies computed by the GEMS server at the local board. This procedure provides a check against manipulation of the GEMS server or memory devices that transmit data from the DRE to the GEMS server. This procedure makes undetected attacks against the GEMS server difficult. A limitation of this procedure, however, is that it does not detect a compromised DRE, which might print a false receipt consistent with false data written on the memory device.

2. The integrity of the election depends crucially on trust in the DRE hardware, operating system, and software. To check that the proper software is loaded on the DRE, a "cryptographic hash" value (unique fingerprint) is computed from the software (any modification of the software would likely result in a different hash value). The precise procedure for loading an checking the OS and DRE software is extremely critical to the security of all DREs. Prior to installation, the DRE installation executable code is hashed on a secure machine, and then the installation executable is loaded onto the DRE. In addition, the SBE states that it can hash the executable that is stored on the DRE by copying this executable from the DRE to a secure machine. Performing such a copy, however, likely requires trust that the OS performs the copy command correctly. Once the DRE executable is loaded it is not routinely checked again unless and until a new version is loaded.

   When a new DRE enters the equipment pool, the OS is loaded from a trusted storage device. This OS software is not hashed. Once the OS has been loaded on a DRE, it is not checked again.

The SBE must pay great attention to its process of loading and checking DRE software and OS to ensure that the proper software and OS is installed on all DREs. To avoid a single human point of failure, the person in charge of loading software should be different from the person in charge of parallel testing.

3. Between elections, the DREs and GEMS servers are stored at the local election boards, in a locked, access-controlled storage facility. Each machine has a locked case and is sealed with special tamper-evident security tape with a unique serial number on the tape. The case locks are of relatively low quality and all 19,000 DREs in Maryland share the same physical key. The security tape provides a layer of protection against physical tampering. The security tape provides a deterrent against unauthorized access to DREs by unsophisticated attackers. A limitation of the security tape is that a sufficiently capable adversary could replicate it with a mobile tape manufacturing device in a motor vehicle. As a rough conjecture, we speculate that such a capability could be purchased for less than $100,000. Use of security tape also comes at the minor cost of some increased disruption caused by legitimate accidental breakage of the tape.

VVAATT

Implementation. This system has two possible scenarios for deployment and each has different implications for costs and implementation. One scenario is to leave a tape recorder plugged into each DRE set up in a polling place. Another scenario is to have a central “bin” or basket for tape recorders and headsets where voters would pick up a tape recorder from the bin for use during the vote and return it to the bin after voting.

One-time costs for acquiring the tape recorders, cords and associated media are estimated at $100 per DRE. Ongoing cost would include storage of tape recorders and media and transportation from precincts and LBEs. The state would have to decide whether to reuse media between elections or to purchase new media for each election. The tape recorders are battery-operated so there is an expectation that each election would require a new set of batteries, which involves both logistical and cost issues.

Integration is not an issue since the tape recorder is totally independent of the DRE. There needs to be enough space in the stand for the DRE to also hold the tape recorder, cord and headset.

We awarded the VVAATT a score of 3.5 for implementation.

Impact on Elections. This system would impact elections in Maryland in most if not all of the ways indicated in our discussion of the impact of any independent verification system on Maryland election administration, i.e., at least double the complexity of election administration. However, The VVAATT has probably the least impact because it involves the least amount of integration. It plugs into an audio port on the Diebold AccuVote TS and receives audio output. Nevertheless, the audio recorders and the audiotapes must be securely transported and stored. Moreover, election officials and judges will have to be trained to deal with a range of issues and problems that will predictably result from incorporating the VVAATT into the Maryland election system. These include but are not limited to: recorder batteries dying; recorders dying; recorders running out of tape; malfunctions in the interface between the DRE and the recorder; users being confused about or failing to understand instructions form the recorder.

Finally, we would be remiss if we did not note a personal hygiene and public health issue associated with the deployment of the VVAATT. Not all persons are "hygienically correct." This means that some fraction of voters who use the headphones for the VVAAT system will present with personal sanitation, hygiene or medical issues that will render the head phones susceptible to the transmission of those conditions to subsequent users of the headphones (e.g., lice, bacterial or viral infection). Either fresh, sanitized head phones will have to be provided for each voter, or some other method (e.g., disposable covers) will need to be developed to ensure that conditions resulting from sanitation, hygiene or medical issues are not transmitted via the head phones.

Therefore, we awarded the VVAAT a 2.5 for impact on elections.

Data Management. In this system, regarding single point failure, the voting events are recorded and hence it is similar to parallel testing in the sense that there is no rollback function implemented because the DRE and the verification systems are not connected in any way. Hence our score is NA.

On atomicity, we have assigned a score of NA because there is no two-way communication with the DRE. For the same reason we assign the same score on tamper-proof data entry. Regarding event recording this system by recording every voting event also creates a record of every event that took place with a date and time stamp; hence our score is 5. For the same reason this system records every error condition and receives a score of 5. Regarding self-management, this system is a passive recorder of events, thus it has no self-management functionality and receives a score of 1 in this category.

Reliability and Functional Completeness. Our test results for VVAATT can be seen in Appendix J. This solution was a functionally simple audio-trail solution. It implemented all of its promised functionality. Therefore, the functional completeness assessment is 3.5. There were some problems with the operation of the voice-operated recorder that can lead to infrequent problems. Therefore, the reliability assessment for this system is 4.0.

Accessibility. We assigned a rating of 1 to the Selker system indicating that it was hardly accessible at all. The Subpart B Criterion is not satisfied, as the user does not have the ability to interrupt, pause or restart the audio. The Subpart C Criterion a requires "At least one mode of operation and information retrieval that does not require user vision shall be provided" since the assistive technology clause is not appropriate. This clause is not satisfied. The demo system provided for evaluation purposes did not support use of the system without vision. The system provides an auditory confirmation for visually capable users, but does not provide an interaction method suitable for users without vision. Users without vision would not be aware of this system without modified instructions or guidance from another individual. Further, if the initial audio verification was not heard properly, users are not given the capability to repeat the audio.

The Subpart C Criterion b requires “At least one mode of operation and information retrieval that does not require visual acuity greater than 20/70 shall be provided in audio and enlarged print output working together or independently”. This clause is not satisfied since the speech generated and recorded by this system is not shown visually (e.g., with captions).

The Subpart C Criterion c requires “At least one mode of operation and information retrieval that does not require user hearing shall be provided.” There is no feedback given to users who cannot hear in this system, so this criterion is not satisfied.

Seventeen criteria listed in Subparts B, C, or D were either not applicable for this solution or were satisfied. Four criteria may or may not be satisfied depending on the final implementation. Four criteria were not satisfied.

Election Integrity (Security). An audio recording provides a check against the threat that a malicious DRE might record a vote differently from the vote cast. Any voter, who chooses to listen, could notice in the voting booth if the vote spoken in a headset is the same as the vote cast. A cassette tape provides an analog audio record of the votes that eliminates required trust in any electronic device, but introduces required trust in managing the recordings. There is evidence that voters are more likely to catch mistakes when checking with headsets than by looking on a separate screen or printed receipt.

The voter cannot, however, check if the vote recorded in the DRE is the vote cast. Only after the election, by listening to the recordings, would election officials notice a discrepancy between the audio logs and the electronic totals.

Auditing the audiotapes could be done in two ways: by human listening or by automatic electronic voice processing. If humans listening and tallying perform auditing, then the process will have an inherent inaccuracy from human error. If the auditing is performed automatically, then the hardware and software that performs this task must be verified.

This system depends on the security and accuracy of counting votes from the audiotapes. The system provides no cryptographic protection of the tapes, and managing tapes is fraught with significant security challenges (e.g., forging, adding, removing, or modifying tapes).

As configured for the study, the VVAATT system has no software and therefore requires no trust in any software. It does require trust in the simple hardware to ensure that the audio signal being recorded is the audio signal played in the headsets; the voter can visually inspect this hardware while in the voting booth.

Therefore we awarded the VVPAT a score of 3.5 for election integrity.

Voter Privacy. As does the VVPAT, the VVAATT system raises some threats to voter privacy. First, the tape preserves the order of the votes cast. Moreover, any voter could cause the device to record a unique "signature" by selecting and deselecting candidates in a special order. By observing at the precinct who voted before or after the person creating the signature, anyone with access to the tapes and knowledge of the signature could determine how such people voted.

Second, if the system permits the voter to hear the ballot choices in a language other than English, then it might be possible to identify the voter from the recorded ballot choices on the audiotapes, especially if the foreign language is unusual for that precinct. For example, in some precincts there might be only one or a very few people who choose to use Spanish. If the voter registration information indicates the language, such identification would be facilitated by that information.

Third, as with all vote verification systems, if a voter requires assistance with the system, there arises an increased vulnerability to the voter's privacy: the poll worker who provides the assistance might be able to observe how the voter's ballot choices. Therefore, any system such as this one that increases complexity (in this case by requiring a voter to wear headsets) slightly increases the risk of loss of voter privacy.

Therefore, we awarded the VVAAT a score of 2.0 for voter privacy.

Pnyx.DRE

Implementation. The vendor estimates the purchase cost for the verification unit at $500 per unit. It is not clear whether this includes the cost for USB sticks needed for each unit. It is presumed this does not include the cost for the additional memory for the DREs and additional storage on the DRE if the state decides to store the encrypted vote records on both the verification module and the DRE.

Each DRE would need an additional 1 megabyte of memory according to the vendor to allow for the interaction between the DRE and the verification module. There are options for plugging the verification module into the DRE, including a USB port or serial port. Given the relatively short battery life for the verification units of several hours, there is a need for a separate power source. There needs to be enough space in the stand for the DRE to also hold the verification unit. Because the verification unit has to be upright for the voter to interact with it during the voting process, the stand for the DRE may require modification.

Because this solution maintains a relationship between the vote tally on the DRE and the verification module, software integration is expected to be somewhat burdensome. Without having access to the Diebold software, SCYTL estimated that it might take two hours of software development on the DRE software that manages the voter confirmation process to achieve integration between the verification module and the DRE. Additionally, we note that the researchers at UMCP spent over 100 programming hours with this system to make it work in their field tests.

We awarded the Pnyx.DRE a score of 2.2 for implementation.

Impact on Elections. Pnyx.DRE will produce all of the impacts on elections as previously discussed. In addition, because the system involves cryptography that is required for this verification solution, there are significant key management requirements placed on both the SBE and the LBEs. There is an open issue about whether there would need to be any key management at the precinct level, with the presumption that it would occur at the LBE as the lowest level. Each verification unit needs to be set up and hooked up to its corresponding DRE. Because the solution establishes a relationship between the DRE and verification module, the setup for each machine might be somewhat burdensome for local election judges. Because of the size of the verification unit, there may be an issue of whether it will fit in the current stand for the DREs.

Finally, during the vote counting process after the election, the LBE must calculate an integrity record from the votes stored in the DRE. Then, the LBE compares the integrity record obtained from the DRE to the integrity record of the verification module. If the two integrity records match, this means that the votes in the DRE were not changed from what the voters voted on the DRE. If the two records do not match, then a parallel recount can be performed by decrypting the votes stored in the verification module and comparing them to those stored in the DRE one by one.

Therefore, we awarded the Pnyx.DRE a 2.0 for impact on elections.

Data Management. Regarding single point failure, Pnyx.DRE does provide a two-way communication between the DRE and the verification system. Hence a failure caused by inconsistent recording of a vote in the two systems will be flagged. While recovery from such a failure is possible, it has not been implemented in the current system. Hence we assign a score of 2.

Pnyx.DRE received a score of 4 for atomicity because it implements a two-way communication between the DRE and the verification system hence a commit protocol could be implemented. However, such implementation does not exist in the system provided by the vendor at the time of this study. On tamper-proof data entry we assign the same score for the same reason. Regarding event recording, Pnyx.DRE gets a score of 1 since it does not record events. For the same reason it does not record the error conditions in the DRE, and in exception recording it receives a score of 1. On self-management, this system contains some level of diagnostic capability but does not have a system monitor, and receives a score of 3.

Reliability and Functionality. This solution was a software intensive solution. In our testing, it did not suffer from mechanical problems. In our examination, we tested each item in this list of functions. Our test results are given in Appendix J.

Pnyx.DRE did not implement some of its functions. These are noted in our test results using the phrase “Functionality not included.” Some of the functions that were not implemented can be justified because implementing them requires the knowledge regarding the operation of the voting terminals and the file formats in which the votes are stored. However, SCYTL provided a virtual voting machine to simulate integration with the Diebold DRE. These functions could be also implemented to work with the virtual voting machine.

During the configuration and operation of Pnyx.DRE, we experienced a number of failures. We reported them to SCYTL and communicated back and forth to solve them. SCYTL was cooperative and helpful in trying to resolve the issues. Note that our test report in Appendix J mentions the feedback provided by SCYTL in appropriate places. However, some of the problems still persisted.

The functional completeness assessment for this system is 3.5. It provided more than half of its promised functionality. The reliability assessment is 2.0. There were frequent failures, some of which were related to the important system functions.

Accessibility. Of the systems evaluated, the Pnyx.DRE system received the most favorable rating. With a rating of 3.0, this system was considered somewhat accessible but there were still concerns. The Subpart B Criterion c(1) requires that “Controls and keys shall be tactilely discernible without activating the controls or keys.” Keys on the keypad provided were labeled with Braille, but there were concerns that individuals reading these Braille labels may accidentally activate the keys since the Braille is on the key tips rather than to the side of the keys. Further, this implementation requires non-visual users to be able to read Braille, but only a small fraction of non-visual / legally blind individuals read Braille. In addition, the "ent" embossing for the "ENT" (enter) key is presented vertically, which is not a standard method for presenting Braille.

The Subpart B Criterion e requires “When products provide auditory output, the audio signal shall be provided at a standard signal level through an industry standard connector that will allow for private listening. The product must provide the ability to interrupt, pause, and restart the audio at anytime.” The solution evaluated allows users to interrupt and restart the audio, but they cannot pause it.

The Subpart C Criterion b requires “At least one mode of operation and information retrieval that does not require visual acuity greater than 20/70 shall be provided in audio and enlarged print output working together or independently.” The Pnyx.DRE system provided solutions that worked adequately independently, and the user can use either the audio or the screen with large print. However, the two cannot be used together effectively. For example, throughout the transaction, the audio prompts do not match the visual instructions. In addition, near the end of the transaction (confirmation screens) the combined instructions are particularly confusing, making it easy for the voter to walk away before completing the final step of the process.

Thirteen criteria listed in Subparts B, C, or D were either not applicable for this solution or were satisfied. Eight criteria may or may not be satisfied depending on the final implementation. Three criteria were not satisfied. One criterion could not be evaluated given the information and equipment provided.

Election Integrity (Security). The Pnyx.DRE System provides an independent electronic record of the votes cast, providing a check against the threat that a malicious DRE might record a different vote from the one cast. Also, by displaying the ballot choice communicated by the DRE to the Pnyx.DRE, the system increases the chance that the voter will cast his ballot as intended. The Pnyx.DRE applies standard cryptographic protections to increase assurance that its electronic record cannot be tampered.

Within its goals, the Pnyx.DRE System is well engineered from a security perspective. Standard cryptography is applied appropriately. However, the system does not provide the ability of any voter to verify that his vote has been correctly tallied. To believe the tally computed by Pnyx.DRE, the voter must trust Pnyx.DRE.

As with VoteHere, Pnyx.DRE uses threshold cryptography to reduce the chance that any one election official can misuse his keys.

In summary, Pnyx.DRE provides a well-engineered independent electronic audit of the votes cast. To affect an election without detection, an adversary would have to compromise both the DRE and Pnyx.DRE. However, for an adversary who is able to compromise the DRE (e.g., by subverting its operating system), the marginal difficulty of compromising a similar independent electronic system (e.g., by subverting its operating system) might be relatively low. For example, it would likely help such an adversary if both the DRE and Pnyx.DRE units were stored in the same storage facility under the same security procedures and access controls. If this system is used, care should be taken to make this marginal attack difficulty as high as possible.

The Pnyx.DRE system requires complete trust in the software to ensure that the vote recorded in the verifier is the vote displayed on the verification screen.

Therefore, we awarded the Pnyx.DRE system a score of 4.0 for election integrity.

Voter Privacy. The addition of this verification-audit system increases the risk of invasion of voter privacy because this second system must also protect privacy. Pnyx.DRE’s design strategy does provide privacy protection through the use of encryption and reordering of encrypted ballots. We awarded the Pnyx.DRE the score of 2.5 for privacy.

VVPAT

Implementation. Diebold’s estimate for each verification unit required for each DRE is $1,500, not including case or consumables, and this estimate may be low. There will be on-going costs for transportation, storage of the printer units, and for associated supplies like paper.

Per Diebold documentation, the Accuview Printer Module would require a "retrofit on existing Accuvote-TS machines." This has not yet been done in production, and we have no hard estimate of the cost for the retrofit.

We gave the VVPAT a score of 1.5 for implementation.

Impact on Elections. The Diebold VVPAT will have all of the impacts on elections that we discussed earlier. Furthermore, voters may not use it well or effectively. For example, in the precincts that were observed in Las Vegas in the Nevada election in 2004 that used VVPATs, something less than 40 percent of voters actually did the verification. That is, less than 40 percent actually looked at the printer screen, compared what it displayed to the DRE, and touched the DRE to indicate that they had verified. The rest did not look at the printer at all and simply touched the screen to say that they had verified (Los Angeles County, CA, Registrar/Recorder, 2004). This same phenomenon will probably occur with other vote verification technologies – that is, only a fraction of voters will really verify their votes. Additionally, we would expect that, over time and as voters become more accustomed to the technology, actual vote verification will drop off even further.

The fact that the printer tested in this study had a significant failure rate strongly suggests yet another impact. LBE officials and volunteers (or on-site technical staff) will have to fix printer malfunctions and failures, assist voters when these problems occur, ensure that voters' voting decisions are not interfered with by printer problems, and ensure that printer problems do not permit election officials or subsequent voters to see the votes cast by previous voters.

Finally, if a recount were to occur, it would rely on a mix of manual tally and automated scans of the paper receipts and would likely be quite burdensome. Data presented to the California Assembly on the potential impacts of VVPAT implementation estimated that a manual recount of paper receipts could be quite substantial ($47 million over 8 years for the over 10 million registered voters in California under certain assumptions.)

Based on these impacts, we awarded the Diebold VVPAT a score of 2.0 for impact on election administration.

Data Management. On single point failure, Diebold VVPAT has a one-way communication between the DRE and the verification system and hence inconsistent recording of votes may be flagged but there is no way to recover from such an event. Hence we assign a score of 1. On atomicity, we have assigned a score of NA because there is no electronic communication between the two systems. On tamper-proof data entry we assign the same score for the same reason. Regarding event recording, Diebold VVPAT by recording every voting event also creates a record of every event that took place with a date and time stamp; hence our score is 5. For the same reason it does not record the error conditions in the DRE, and in exception recording it receives a score of 1. The functionality required for self management does not exist in Diebold VVPAT because the DRE and the verification system are not connected. Hence there is no automated manner in which consistency checks can be performed across the two systems and it receives a score of 1.

Reliability and Functional Completeness. We voted 350 times using this system to receive the paper receipt. We experienced paper jams for a total of 29 times. These failures can affect the operation of an election considerably. Printed papers sometimes do not roll-up which would allow the next voter to see the votes of the previous voter. In our testing, we corrected paper-jams before testing, which may not be the case in an election setting.

There were no differences between the votes that we entered and those printed out by the Diebold's system. We did not observe any problem in this respect. The problems were mechanical and appeared in a random fashion.

The functional completeness assessment for this system is 5.0. Diebold was able to supply a fully functioning system to fulfill its promises and mission. The reliability assessment is 2.0, "frequent failures." Considering the functional simplicity of the system, better reliability could be expected.

Accessibility. The Diebold system, as provided, could not be effectively evaluated for accessibility since controls for users with low or no vision were not provided. The following describes the part of the evaluation that we could conduct with the available equipment.

The Subpart B Criterion h requires "When a product permits a user to adjust color and contrast settings, a range of color selections capable of producing a variety of contrast levels shall be provided." The Diebold solution does provide contrast settings, but the controls provided do not provide a range of settings as required by the criteria.

Eight criteria listed in Subparts B, C, or D were either not applicable for this solution or were satisfied. Nine criteria may or may not be satisfied depending on the final implementation. One criterion was not satisfied. However, an additional seven criteria could not be evaluated given the information and equipment provided.

Election Integrity (Security). In the Diebold VVPAT system, a printout on a roll of paper provides a check against the threat that a malicious DRE might record a vote differently from the vote cast. The voter may view the printout of his ballot choices through a glass window. Any voter who chooses to look could notice in the voting booth if the vote printed on the paper roll is the same as the vote cast. The paper printouts provide a non-electronic record of the votes that eliminates required trust in any electronic device, but introduces required trust in managing the paper.

The voter cannot, however, check if the vote recorded in the DRE is the vote cast. Only after the election, by counting the printed votes, would election officials notice a discrepancy between the paper and the electronic totals. The voter can verify only what goes into the paper storage unit. The voter cannot verify what comes out of the paper storage unit.

This system depends on the security and accuracy of counting the paper printouts. The system provides no cryptographic protection of the printed votes, and managing paper is fraught with significant security challenges (forging, modification, stuffing, destroying, accuracy, inefficiency).

Each printout also contains an optical scan barcode that claims to encode the same information printed on the roll in plain text. Election officials could use these barcodes in counting the paper votes. Realistically, these bar codes cannot be checked in the voting booth. Moreover, they are generated from information sent from the DRE to the printer. As such, the barcodes cannot be trusted without trusting the DRE and the printer. Therefore, if the barcodes are used, they must be verified (it might suffice to verify a sufficiently large random sample of the barcodes). It might be better not to print the barcodes.

On general principles, it is problematic that the same vendor, Diebold, which produces the DREs, provides the printer device. Thus, this system does not provide vendor independence in the verification-audit component. On the other hand, this concern is reduced somewhat given the simplicity of the printer unit.

Except for the bar code, the VVPAT system requires no trust in the VVPAT software concerning the integrity of what is printed on the paper roll (the voter can visually inspect the printout).

We gave the VVPAT a score of 3.5 for election integrity.

Voter Privacy. The Diebold VVPAT system raises some threats to voter privacy. First, the roll of paper preserves the order of the votes cast. Moreover, any voter could cause the device to print a unique “signature” by selecting and deselecting candidates in a special order. By observing at the precinct who voted before or after the person creating the signature, anyone with access to the roll and knowledge of the signature could determine how such people voted.

Second, the barcode, which cannot reasonably be checked by the voter, might secretly encode information helpful in identifying a voter. For example, it could encode the time of day, or how a previous person voted. Similarly, other covert channels may be possible by having the printer make apparently innocuous small marks on the paper.

Third, if a paper jam occurs, a poll worker or the next voter might be able to see the vote of the current voter.

Fourth, if the printer prints the ballot choices in a language other than English, then it might be possible to identify the voter from the printed record, especially if the foreign language is unusual for that precinct. For example, in some precincts there might be only one or a very few people who choose to use Spanish. If the voter registration information indicates the language, such identification would be facilitated by that information.

Finally, we note that the Diebold VVPAT is not an independent vote verification technology – independent that is, from the same company that manufactures and supports the DREs. This presents a substantial issue of trust for Maryland voters. If they lack trust in the outputs of the DREs, why would they have any greater trust in the outputs of the paper produced by the same company.

We gave the VVPAT a score of 2.0 for voter privacy.

VoteHere

Implementation. In a briefing with the UMBC study team, VoteHere estimated the cost of a verification unit (i.e., Sentinel) to be approximately $500 and an additional $50 if the state wanted to add an LED screen for each DRE. On-going costs would include paper for printers. Additionally, there are also the potential one-time and recurring costs to develop the web site that VoteHere advocates the state uses to allow voters to check whether their vote had been recorded.

Integration is expected to be significant between the Sentinel and DRE as the vote verification process require integration between the VoteHere system and each DREs, the election management system on GEM services and the web site that voters would use to check whether their votes had been recorded. The DRE integration would be an interface issue because voters have to move between the DRE and verification unit to record their vote in the DRE, check it in the Sentinel and then submit it finally on the DRE. Audit capabilities designed for this solution require integration with the election management system at the GEM server.

We awarded the VoteHere system a 2.0 for implementation.

Impact on Elections. Because VoteHere is an election management system, not just a vote verification system, it is far more complex than other products examined here. To begin with, considerably more work is required to configure the VoteHere Sentinels in advance of an election. Because of the cryptography required for this product, there are significant key management requirements placed on both the SBE and the potentially the LBEs. There is an open issue about whether there would need to be any key management at the precinct level, with the presumption that it would occur at the LBE as the lowest level. VoteHere maintains that physical set up would require an independent power source for the Sentinel and the DRE and a connection for the Sentinel to the DRE through a serial port.

Due to the complexity of this solution, state and local election officials will have to undertake a significant public information campaign to explain to voters how and why this solution helps ensure the integrity of the voting process. This need for public information and education likely extends to election judges who would have the responsibility to explain the significance of the receipt and the entire vote verification process to voters and to intervene with problems arise (e.g., equipment malfunctions, voter confusion). The VoteHere solution records the votes a second time in the verification unit, allowing for a parallel recount. Because there is no two-way data integration between the DRE and Sentinel by design, it is possible that recounts could be complicated by each method producing different tallies.

There is an open issue about whether printing a receipt for voters will produce issues with printers and replenishing papers supplies during an election. There is also the possibility of printer malfunctions.

There will also be significantly more training required for LBE officials and volunteers to be able to understand this system and do at least the following: initialize the Sentinels; create and share cryptographic keys; explain to voters the differences between the long and short receipts; explain to voters the meaning of the numbers printed on the receipts and what to do with the receipts; and tally the votes from the Sentinels and compare them to the results from the DREs after the election. Additionally, voters will have to engage in a complex interaction with the Sentinel in the voting booth, choose and accept a receipt from the Sentinel, and acquire trusted software with which to access information on a public web site to tell if their receipt numbers match the numbers printed on the web site.

For these reasons we gave VoteHere a 1.0 for impact on elections.

Data Management. On single point failure, VoteHere has a one-way communication between the DRE and the verification system and hence inconsistent recording of votes may be flagged but there is no way to recover from such an event. Hence we assign a score of 1. On atomicity, VoteHere implements a one-way communication and hence receives a score of 1. On tamper-proof data entry we assign the same score for the same reason. Regarding event recording, VoteHere gets a score of 1 since it does not record events. For the same reason it does not record the error conditions in the DRE, and it receives a score of 1 in exception recording. On self-management, this system does have some diagnostic capability but does not have a system monitor, and we assign a score of 3.

Reliability and Functional Completeness. VoteHere was the most complex system that we examined. The company provided us with a prototype system that does not receive, encrypt, store, or verify any vote, whatsoever for vote verification purposes. The device provided by VoteHere merely acted as a printer that produced some paper receipts to simulate the functionality provided by the system. In addition, there was no implementation of a web site or phone system that can be used for the verification purposes. Because we could not evaluate its reliability, we were unable to provide it with a reliability score. VoteHere received a 1.5 for functional completeness.

Accessibility. We evaluated the accessibility of the VoteHere solution based on the prototype and operational descriptions provided by the vendor. We assigned a rating of 2.0 to the VoteHere solution. The Subpart C Criterion a requires “At least one mode of operation and information retrieval that does not require user vision shall be provided.” The receipt, as presented, is not usable without vision. The printed receipt shows various pieces of information related to the transaction including a ballot lookup number and tracking code. This information is not provided in audio-only form. In an alternate implementation that was not demonstrated to the researchers, the user with an embossing of the receipt can stamp the receipt. Specifications for this stamping system were requested but were not provided for the evaluation. Only samples of stamped receipts were supplied and, therefore, the mechanism for user interaction for this stamping process could not be evaluated. However, we can say that on those samples, the words "I Voted!" are printed in large print. The words "I voted" (small caps, no exclamation point) are embossed in raised type. And the words "I voted" are embossed in Grade 2 Braille (the '-ed' at the end of voted is contracted). Braille is read by only a small fraction of the non-visual and the legally blind population; and Grade 1 Braille readers do not generally read contractions. Grade 2 readers can read Grade 1 (uncontracted) Braille. Neither the embossed information nor the Braille included the lookup number or tracking code. In addition, there are interface concerns related to this criteria that may or may not be satisfied depending on the final implementation.

Fifteen criteria listed in Subparts B, C, or D were either not applicable for this solution or were satisfied. Nine criteria may or may not be satisfied depending on the final implementation. One criterion was not satisfied.

Election Integrity (Security). Of the systems examined, VoteHere is the only one to provide end-to-end voter-verifiable election integrity. Each voter receives a paper long receipt, which he may take home. After the voting process closes, using his receipt and information posted on a public web site, the voter can verify that his vote was recorded in the official election data as cast in the voting booth, and that all recorded votes were correctly tallied. This verification ability, based on complicated mathematical cryptography, offers a revolutionary improvement in the level of election integrity—albeit at a slightly greater procedural complexity to the voter and election officials, slightly greater conceptual complexity of the system’s informal high-level properties, and significantly greater mathematical complexity of the underlying details, mechanisms, and proofs.

In the rest of this section, we consider use of the cast-vote button, use of optional display screen, use of optional special user-interface devices, risks from complexity, voter verification from the public web site, and the nature of required trust. In addition, we carefully examine each of the vendor claims on election integrity and voter privacy.

By issuing the cast-vote or reject-vote command on the Sentinel, the voter does not have to trust the DRE to process them properly. These buttons improve upon alternative configurations of the Sentinel for which the voter issued these commands on the DRE. On the other hand, in this configuration, the DRE does not learn if the voter presses a button or which one is pressed. As a result, even if the DRE and Sentinel are honest, there is no guarantee that the units will record the same number of votes.

Although, for some voters, a display screen might possibly be easier to read than a printed receipt, there are several reasons why having a screen on the Sentinel might be a bad choice. First, the end-to-end voter verification of election integrity requires the voter in the voting booth to verify the receipt he takes home (whatever form that receipt takes, whether it be a standard printed receipt, Braille receipt, or audio recording). As such, the screen might be a distraction from the important task of verifying the receipt. Second, screen displays add complexity to the software, thereby decreasing software assurance. Third, there is no compelling reason to add a screen. Checking the receipt in the voting booth enables the voter to verify that he has cast his ballot as intended. Voting instructions can be communicated through an instruction card. The command to cast (or reject) a vote is issued via a physical button on the Sentinel. In addition, voters still have the option of inserting a specialized user-interface device with screen into the Sentinel's USB port. For these reasons, we assume the VoteHere system for the Maryland study does not have a display screen.

The security properties of the receipts depend only on the information on the receipt and not on how that information is represented on the paper receipt. Therefore, voters who would like to use special user-interfaces can enjoy all receipt properties by downloading the receipt information onto their special devices (e.g., Braille printer) connected to the Sentinel's USB port. In this sense, VoteHere is well engineered for the potential to support end-to-end voter verification of election integrity for disabled voters. Care must be taken that the voter cannot adversely affect the Sentinel through the USB port. Since this USB port feature has not yet been implemented, we cannot evaluate it further.

Because the procedure carried out by the voter in the voting booth for the receipt is slightly complicated, there is a risk that a malicious Sentinel could try to alter this protocol in a way that might go undetected by the voter. For example, a malicious Sentinel could print a ballot commitment after reading the voting ticket. Similarly, a corrupt election judge could pass out malicious instruction cards. It is unlikely, however, that such an attack could affect many votes without detection by some observant voter. Nevertheless, this possible attack underscores the need for voter education.

Each voter can check that his vote is correctly listed in the official election data on the public web site. He can also check for any other voter, if the voter gives him her receipt and an assurance that the challenge-response worked properly in the voting booth. Anyone can verify that the posted tallies are consistent with the official election data.

A limitation of the VoteHere System, however, is that no voter can detect if any of the other votes posted on the public web site are bogus (except as noted above). This is a fundamental property of all end-to-end voter verification systems: a sufficient number of voters must participate. In order to detect any significant tampering with high probability, however, it would suffice for a relatively small percent of the voters to verify that their receipts are correctly listed in the official election data.

The verifiability of votes also depends on the integrity and authenticity of the information on the public web site, and of the tools (e.g., browser) the voter uses to view the site. Voters, however, can check the web site from multiple computers or browsers of their choice. Also, election officials could also redundantly publish the web site information through multiple channels.

It is a benefit of the VoteHere System that, for election integrity, very little trust needs to be placed in the software. What matters is what is printed on the receipts and what is posted on the web site, not how these computations were performed.

Although a large number of corrupt election judges can collude to read encrypted ballots (in an attempt toward violating voter privacy), the soundness of the verifiable mix-net computation is robust even when all election judges are corrupt and collude. To some extent, this threat can be addressed by adjusting the required threshold defined in the election configuration.

During the verification process using the public web site, there are two separate items to verify. First, the voter can easily verify that his encrypted ballot is included in the official posted election data. All that is required is to check that the numbers on this receipt exactly match those in the image of this receipt on the web site. The ballot pledges (which are posted on the web site) together with the voter's receipt encode all of the voter's ballot choices; therefore, if her receipt is on the web site, then her vote is included in the official election data.

Second, anyone can verify that the posted election results are consistent with the official election data. No receipt is needed to perform this verification step, but the checker does need to use some sophisticated mathematical software. The checker can use any trusted software of his choice. He does not have to use software supplied by VoteHere or SBE. It is not necessary for every voter to perform this step. It is sufficient for enough independent checkers to carry out the work.

Now let us examine the nature of the required trust for a voter to believe that the posted results are consistent with the official election data. Few voters would be able to understand how to perform this verification step or to write a computer program to perform this step. Therefore, nearly all voters would have to place trust in some expert, or trusted organization, in much the way someone might place trust in a lawyer or consumer group to explain the terms of a contract to sign. The voter can choose his own experts, and the voter and experts can scrutinize all data needed to perform the verification. The voter can run any number of trusted programs of his choice. Nearly all voters must trust a cryptographic expert, but the nature of this trust is very different from that required to trust any of the other verification system under study. For example, to believe the correctness of results certified by SCYTL requires trust in a system that cannot be examined by the voter or by experts of his choice.

For these reasons, we awarded the VoteHere system a 5 for election integrity.

Voter Privacy. As with all verification-audit systems, VoteHere increases the risk of invasions to voter privacy, because more system components must protect voter privacy. Moreover, the use of a public web page further increases that risk, for example through a covert channel that might signal how a particular ballot was cast (e.g., encrypted ballots ending in a zero might signal a Republican vote).

On the other hand, the VoteHere protocols do properly protect voter privacy. For example, the use of voter tickets in the challenge-response protocol eliminates certain covert channels since the Sentinel makes no random choices. Also, human identity is not directly tied to ballots. Instead, ballots are identified by their hash values and voting ticket numbers. Without product source code (in addition to the VHTi library source code which is available), however, we cannot assess if the protocols are properly implemented.

VoteHere's use of threshold cryptography[14] reduces the risk that a small number of officials at the SBE could misuse knowledge of secret keys to violate voter privacy. Each election judges knows only a "share" of the master key; reconstructing the entire master key requires a certain number of election judges to contribute their shares.

The greatest risk to voter privacy stems from the complexity of the voting experience in the voting booth. If a voter summons a poll worker for help, the poll worker might be able to see the ballot choices displayed on the DRE. Therefore, voter education is vital. There is nothing on receipt printed by the Sentinel, however, that could compromise voter privacy.

Another possible risk is the possibility for a malicious Sentinel to print covert messages on the receipts (e.g., via tiny dot codes), revealing secret keys or ballot choices. Even worse, the Sentinel could digitally sign such messages, proving their authenticity. This risk exists for any device with printer, and for practically any user interface. Finally, use of foreign languages might help identify the voter, but the receipt is encrypted and does not reveal the ballot choices.[15]

[14] Threshold cryptography means that a key is split into two or more shares, which each share held by one of the election judges. At least some threshold (e.g., 2/3) of the election judges must collaborate in order to reconstruct the key.

[15] *Static Analysis of VHTi Source Code*. The underlying source code—the VHTi library—is publicly available from the VoteHere web site. Markus Dale (Dale, 2005), a UMBC graduate student working with Dr. Sherman, performed a "static analysis" of the VHTi library using the free open-source tools Rough Auditing Tool for Security (RATS) and Flawfinder. These tools seek many common software security errors, such as buffer overflows, without executing the tested software. Although imperfect, these tools help check that the source code is free of certain serious software errors. The U.S. Department of Homeland Security lists static analysis as a best coding practice.

In examining about 10,000 lines of software, these tools found (only) nineteen unique potential vulnerabilities (nine by RATS, sixteen by Flawfinder, with an overlap of six potential vulnerabilities found by both tools). Upon further examination, none of these potential vulnerabilities appear to be true vulnerabilities. Overall, the quality of the VHTi source code seems to be very high.

Limitations of this static code analysis include the following. First, static analysis does not check for

We awarded the VoteHere system a 2.5 for voter privacy. Please also see Appendix K for a summary table of detailed scores for privacy, security and disruption for all solutions examined.

Conclusions and Recommendations

In this study, we have examined four vote verification systems as well as the SBE's method of parallel testing. We examined them using the following criteria: implementation, impact on voters and election administration, data management, functional completeness, reliability, accessibility, election integrity (security), and privacy. Our summary rankings are found in Table 6.

Table 6: Vote Verification System Summary

	Parallel Testing	VVAATT	SCTYL	Diebold VVPAT	VoteHere
Implementation	5.0	3.5	2.2	1.5	2.0
Election Administration	4.0	2.5	2.0	2.0	1.0
Data Management	1.2	1.8	2.5	1.3	1.3
Functional Completeness	5.0	2.5	4.0	5.0	1.0
Reliability	5.0	4.0	2.0	2.0	N/E
Accessibility	N/A	1.0	3.0	N/E	2.0
Election Integrity	3.0	3.5	4.0	3.5	5.0
Privacy	3.0	1.5	2.5	1.5	2.0

N/A means not applicable because voters do not participate
N/E means not evaluated because the vendor did not provide the necessary equipment
Note: Although parallel testing perfectly preserves voter privacy because it does not use real voter data, it does not correct the threats to voter privacy created by Diebold AccuvoteTS. Therefore, parallel testing receives the same privacy score as the baseline score for Diebold.

As readers will note, Table 6 presents rankings from 1 to 5 for each criterion for each system including parallel testing. A ranking of 1 means that the system (or parallel testing) does not meet the criterion at all, and a ranking of 5 means that a system meets a

errors of logic in the algorithm or implementation of it. Second, no static tool can find all vulnerabilities, and one vulnerability might be sufficient to compromise an entire system. Third, because we were not provided source code for any particular product, we could test only the source code for the VHTi library. We could not test the software of any product that called this library. The security of any product depends both on the library and the source code that calls the library functions. Most of the nineteen potential vulnerabilities detected were situations where it is necessary for the calling software to avoid certain pitfalls. In each case, the VHTi source code clearly documented requirements on the calling program.

This analysis was for security purposes and was not related to the evaluation of reliability and functional completeness.

criterion completely. Table 6 does not contain average scores for each system. This is because, first, different persons and organizations will assign different values, weights or importance to each of these criteria. For example, for some, security may have the greatest importance, while for others, the impact on voters and election administration may be in the forefront. To still others, accessibility may have the greatest value. We did not want to assign our subjective judgments here, preferring to leave it to readers to assign their own values.

Second, for some of the systems, some of criteria did not apply. For example, accessibility did not apply to parallel testing because parallel testing was not an add-on technology to be used by voters in elections. Therefore, we could not examine it as such and could not rank for accessibility. Third, for other systems, we were unable to conduct some tests because vendors did not provide the full range of equipment necessary to conduct the tests (e.g., Diebold and VoteHere).

Finally, we believe that the narrative discussion and tabular rankings provided herein should be sufficient to enable policy-makers to understand what these systems do or purport to do, to understand their relative strengths and limitations and to make decisions about them – particularly decisions about whether the State of Maryland should adopt and implement any one of them as part of its election system. In this process, policy-makers will understandably apply their own subjective judgments about the relative importance of each criterion.

The following is an overview that briefly describes parallel testing and each of the vote verification systems and discusses their principal strengths and weaknesses.

Parallel Testing

Parallel testing is a method of testing an electronic voting unit by producing an independent set of results that can be compared against the results produced by the voting unit. On the day of statewide elections, SBE partners with the Maryland League of Women Voters (LWV) to produce this independent set of results. Parallel testing represents the baseline verification system for this analysis since the state used this process as part of the 2004 general election.

Since the SBE has applied this verification system already, continuing to implement should not require many, if any, additional state resources. The partnership with the LWV also helps keeps on-going costs to a minimum. Since the parallel test is independent of actual polling, the impact on election administration is minimal. This also means that no one’s real vote is used during the test, minimizing any potential privacy issues.

The relatively small number of units selected from just one local board of elections means that election integrity might be compromised.

VVAATT

The Voter Verified Audio Audit Transcript Trail (VVAATT) system includes a voice-operated recorder that is connected to a voting machine (DRE). The voter puts on headphones in the voting booth. Subsequently, every activity of the voter is spoken to the voter using an audio feedback mechanism and recorded in an audiotape. If the election results are challenged or need to be audited, the tape can be played back and the votes can be manually counted and compared with those of the DRE.

The vendor for the VVAATT has a relatively well-developed prototype. The system is simple and easy to install. Since the tape recorder is independent of the DRE, integration with the DRE is not an issue. Being an analog audio record, this system eliminates the possibility of any mistrust on electronic devices.

The vendor lacks a business plan for producing and marketing large number of units that would be required if this product is selected. The system has significant accessibility problems, poses threats to voter privacy, and is not resistant to disruption. The system only provides manual operations for playbacks and recounts of votes, making it very labor intensive, which might introduce significant impacts on election administration (i.e., in the event of recounts).

Pnyx.DRE

The Pnyx.DRE solution is a device attached to a DRE and provides voters with a secure and reliable environment in which they can verify that their votes are cast correctly. Once the voter has selected the voting options in the DRE voting terminal, these choices are sent to the verification module where the voter can verify them. The cryptographic protocol implemented by the verification module protects the privacy and integrity of every single verified vote. By means of this protocol, the electoral board members can automatically check whether the tally of the votes is consistent with the recorded votes.

Pnyx.DRE automates the post-election audit process. The electoral board members can automatically check whether the tally of the votes is consistent with the recorded votes. The accessibility of Pnyx.DRE received the most favorable rating among the systems examined. However, there are still important improvements needed to make it fully accessible as explained in our report. The system is well documented and the vendor provided us with the necessary documentation, e.g. use-cases, to improve our understanding of the system. The cryptographic protocol digitally signs the votes before storing. During audit and recounting, a mixing protocol ensures privacy by shuffling the decrypted votes. It is engineered well from security point of view. It implements the standard security protocols well.

Pnyx.DRE is a software intensive solution. It lacked some of the functionality given in its specifications, however, more than half of the system was implemented. There were frequent failures in operation. The vendor was helpful and made some recommendations to solve some of the problems. However, some important failures persisted. The implementation of this solution will require additional development effort because Pnyx.DRE has to communicate with Diebold's software on DREs to receive votes in appropriate file formats. In addition, the cryptographic protocols in this solution will put some burden on the officials at SBEs and LBEs. Keys should be generated and distributed among election officials because they are needed for audit and recount purposes.

VVPAT

The Diebold Voter Verified Paper Audit Trail (VVPAT) system consists of a small printer encased in a sealed take-up unit housing that attaches to the side of the DRE. After the voter selects his or her choices on the DRE, the DRE displays the voter's selections and requires the voter to print the ballot to begin the verification process prior to recording the votes. The voter's selections will be directly sent from the DRE to the printer through an integrated serial port and printed out. If the selections displayed on DRE correspond to the selections printed on the printer paper, the voter can accept and record his or her selections and to store the data in a standard manner. Otherwise, the voter can reject the ballot and then modify his or her selections. At the end of an election day, the paper rolls that contain the printed votes will be removed from the printer and securely stored. The votes recorded on the paper rolls can be hand tabulated when there is a need to conduct an audit or recount.

The Diebold VVPAT has a relatively fully functioning system. The system is simple. Because it is produced from the same vendor as Diebold DRE, the system integration effort between the two could be minimal.

Due to the high failure rate of the printer observed during our test, reliability of the Diebold VVPAT is low, and the implementation of this system could be difficult and costly (more than $1,500 per unit not including cases and consumables). Ensuring the minimal effect of printer problems during an election would be challenging. Although all the votes are recorded on the printed paper, an audit or recount rely on a mix of manual tally and automated scans of the printed paper could be labor- and time-consuming. Its election integrity is just so-so due to the fact that the system does not provide cryptographic protection of the printed votes. The system raises some threats to voter privacy as well because the roll of paper preserves the order of the votes cast and can potentially be used to identify a voter. We cannot address whether accessibility is strength or weakness of the Diebold VVPAT because the vendor did not provide the equipment needed for accessibility testing.

Finally, because the same vendor as DRE manufactures VVPAT, it is not an independent system, which reduces its effectiveness and reliability as a verification solution.

VoteHere

VoteHere provides a method of voter verification of election integrity, based on receipts and complicated mathematical cryptography. In the voting booth, the voter enters his ballot choices on the DRE and verifies those selections on a printed receipt, which he may take home. The receipt defines an encrypted ballot; the receipt does not reveal how the voter voted. After the voting process, the voter may check that a copy of this receipt (and thus his ballot as cast in the voting booth) is included in the official election data posted on a public web site. Anyone, using trusted complex mathematical software of his choice, can verify that the official results are consistent with the posted data.

All VoteHere software is open source, and election integrity depends little on this software. VoteHere provides very high election integrity, provided that enough voters verify their votes in the voting booth, enough voters check that their receipts are in the official data, and enough verifiers check the tallies against the data – regardless of whether voters understand VoteHere's complicated underpinnings.

Disadvantages include the following. First, the product is not functionally complete, existing only as a reference library without application software. Second, the voter's experience in the voting booth is slightly complicated. Third, because the system is complicated to understand, election officials will have to be educated in it and will also have to be able to educate voters, and some voters might not have confidence in a system they do not understand. Fourth, voters with limited eyesight might have difficulty reading the receipt, and the planned functionality for alternative user interfaces is not yet available. Fifth, election officials must set up and maintain an authenticated web site. Sixth, as configured, there is no attempt to maintain consistency between the Diebold and VoteHere systems, even when both units are honest. Seventh, as is true for all systems under study, the system requires integration with the DRE display software.

Additional Considerations

In terms of election integrity, each system under study requires a different type of trust. For example: parallel testing requires the voter to trust that the selection of units to be tested is random and that the selected units cannot be signaled to behave properly. VAATT and VVPAT require the voter to trust that the audio or paper tapes are stored securely and counted accurately. Pnyx.DRE requires the voter to trust that the system works as claimed. The voter cannot examine or verify system software for himself. VoteHere requires voters to trust cryptographers of their choice that certain security properties are true (mainly, that trusted software of their choice can verify that the posted results are consistent with the official election data).

Verification systems that preserve voting order (VVAATT and VVPAT) notably degrade voter privacy. All verification systems necessarily degrade voter privacy by increasing risks without correcting existing vulnerabilities. Similarly, all verification systems necessarily increase the risk of disruption.

Data management issues related to recording the votes are also very important. The DRE and verifier comprise a two-part distributed system, where each of the two units contains a repository of votes cast. In such system, casting a vote in the booth is an atomic transaction spanning the DRE and the Verifier. The two repositories should remain consistent, i.e., should record the same votes. If atomicity is not enforced, the two repositories may not record the same votes, leading to a difference in the vote counts between the DRE and the Verifier, i.e., exhibiting the phenomenon of not recorded votes. This is a data management issue independent from the security of the system.

Findings and Recommendations

In the following section, we present the principal findings of our study. These are followed by our recommendations.

First, each of the systems that we examined – only one of which provides for a pure paper solution – may have something to offer the State of Maryland in terms of vote verification. But this would be true only if the system were fully developed, fully integrated with the Diebold DREs and effectively implemented. If this were the case, then each system in its own way could provide a degree of vote verification beyond what is available through the Diebold system as currently implemented.

However, and importantly – and this is our second principal finding – none of these systems is a fully developed, commercially ready product. None of these products had been used in an election in the U.S. (The Pnyx.DRE system has been used outside the U.S., and a different version of the Diebold VVPAT has been used in the U.S.).

Were the State of Maryland to decide to acquire any of these products, anywhere from a relatively small to a considerably large amount of money and effort would be required on the part of the vendor to produce an actual product and make the product ready for use in actual elections. Indeed, nearly all of these vendors are looking for someone (e.g., a state government) to contract with them so that they can fully develop and commercialize their products.

In our expert opinion, it is a very bad idea for governments to buy products that are not functionally compete and that either do not have positive records in the market place or that cannot be fully and effectively tested in simulated elections to ascertain their performance characteristics.

We also note several specific additional issues of concern around these products. These often involve trade-offs that would be implicit in the introduction of any vote verification product -- trade-offs in the sense that were these products fully developed and thus able to provide some of their promised benefits to the state, these benefits would not arrive without costs, sometimes substantial costs. These costs include at least the following:

- All of these products would impose significant one-time implementation and on-going management burdens (cost, effort, security, etc.) on the SBE and the LBEs.
- To a greater or lesser extent, all would increase the complexity of the act of voting.
- To a greater or lesser extent, all would increase the amount of time required to vote.
- All would at least double the amount of effort required to administer elections.
- All would adversely affect voter privacy.
- These products would have both potentially positive and potentially negative impacts on security and election integrity (i.e., increase the possibility that votes will be recorded and counted as cast, and increase the possibility of disruption).
- None can be considered as fully accessible to individuals with visual or hearing impairments and none of them fully meets the accessibility standards of Section 508 of the Rehabilitation Act.
- Integration of these systems will require the cooperation of Diebold to develop and/or ensure the viability of a working interface between the vendors' products and the Diebold system. It is unclear whether this would be in Diebold (or any other DRE system vendor's) self-interests. As such, there is the clear potential here for added difficulties when implementing any of these solutions.

Based on the evidence from this study, we cannot recommend that the State of Maryland adopt any of the vote verification products that we examined at this time.

No election system, regardless of the technology involved, is foolproof nor is any election system completely immune or secure from fraud and attack. Indeed, there is a long and inglorious history of election fraud in the U.S. that dates back to the founding of the country or before and involves nearly all methods and technologies of voting. It would be prohibitively costly to make any election system – or an information system for that matter -- totally secure.

Having said this, though, and having recommended against adopting any of the four vote verification technologies we examined, we would be remiss if we did not make a further recommendation to the State of Maryland. Regardless of what the state does in the near term with regard to vote verification and vote verification systems, in future elections, it should expand the use of parallel testing. The state should also undertake a full-scale assessment of the security procedures and practices around its current voting system. We say this even with the knowledge that the SBE's security procedures are reasonable and prudent and that the SBE's system of parallel testing reduces considerably the possibility of widespread fraud and attack on the system. These additional measures might include: randomly selecting DREs for the test the day before the election; ensuring that the persons responsible for parallel testing are not the persons who loaded the software; selecting a larger number of DREs, possibly from more than one jurisdiction for testing; and ensuring that conditions for the test are as nearly identical to a real election as possible. The SBE should continue to carefully monitor and record the parallel test.

To summarize, each of the products we examined could, if fully developed and properly implemented and managed, offer some value in the area of vote verification. However, none is fully developed. Additionally, (potentially significant) trade-offs exist with all of them (e.g., greater security and degraded privacy), and all would require considerable cost and effort to implement and to manage during elections. For these and other reasons, we cannot recommend that the state of Maryland acquire and implement any of them at this time.

References

Beiler, David. (1989a). A Short in the Ballot Box. *Campaigns & Elections. 10*(2), 39-41.

Beiler, David. (1989b). Shortfall in the Sunshine State. *Campaigns & Elections.* 10(2), 40.

Brooks, F. P. (1987). No Silver Bullet, Essence and Accidents of Software Engineering. *IEEE Computer*, 10-19.

Brooks, J. F. P. (1995). *The mythical man-month (anniversary ed.)*: Addison-Wesley Longman Publishing Co., Inc.

Brunvard, Erik, John Carter, Alan Dechert, David L. Dill, Kathy Dopp, Gensh C Gopalakrishnan, David Hanscom, Michael Jones, Arthur Lee, Jay Lepreau, Kent Seamons, Peter Shirley, Barbara Simons, Association for Computing Machinery, Pamela Smith, and Phillip Windley. (2004). Response to 'American Attitudes about Electronic Voting' Survey and Advice for Utah's Voting Equipment Selection. Memo. Retreived from UtahCountVotes.org/Voting_systems.pdf on 11/2/2005.

Burmester, Mike and Emmanouil Magkos. (2003). Towards Secure and Practical E-Elections in the New Era. In Dimitris Grizalis (Ed.). *Secure Electronic Voting.* (p. 63-76). Boston: Kluwer.

Cal Tech/MIT. (2001). Cal Tech/MIT Voting Technology Report: What is, What could be, Fast Facts. Retrieved from www.vote.caltech.edu/media/documents/july01/fast_facts.pdf on 11/3/2005.

Cranor, Lorrie Faith. (2003). In Search of the Perfect Voting Technology: No Easy Answers. In Dimitris Gritzalis (Ed.). *Secure Electronic Voting*. (p. 17-30). Boston: Kluwer.

Election Assistance Commission. (2005). About the EAC. Retrieved from http://www.eac.gov/about.asp?format=none.

Election Data Services Press Release from 2/12/04 called New Study Shows 50 Million Voters will use Electronic Voting Systems, 32 Million Still with Punch Cards in 2004. Retrieved 1/10/06 from www.electiondataservices.com.

Green, R., and J. Adler, "Threat analysis," VoteHere (2003), unpublished manuscript, 24 pages. http://www.votehere.net/downloads.php

Hall, Thad E. and R. Michael Alvarez. (2004). *American Attitudes about Electronic Voting: Results of a National Survey.* Salt Lake City, UT: Center for Public Policy & Administration.

Hanging Bytes, Pregnant Bits. (2002). *The Economist Technology Quarterly*. *364*(8291), 8.

Herrnson, Paul S., Benjamin B. Bederson, Bongshin Lee, Peter L. Francia, Robert M. Sherman, Frederick G. Conrad, Michael Traugott, and Richard G. Niemi. (2005). Early Appraisals of Electronic Voting. *Social Science Computer Review*. *23*(3), 274-292.

Herrnson, Paul S., Benjamin B. Bederson, Charles D. Hadley, Richard G. Niemi, Michael J. Hanmer (with staff assistance). 2006. The Usability of Four Vote Verification Systems: A Study Conducted for the Maryland State Board of Elections. College Park: Center for American Citizenship and Politics, University of Maryland College Park.

Kurlantzick, Joshua. (2004). 2000, the sequel: in theory, the Help America Vote Act was Congress' attempt to prevent the catastrophes of the last election from happening again; in fact, it may have made things even worse. *American Prospect.* 15 (10), 22-5.

Los Angeles County, CA Register/Recorder. 2004. DVD of Video of Voting on VVPAT in Las Vegas, NV, 2004 General Election. Los Angeles: Author.

Machlis, Sharon. (2004). Public, Security Experts' E-Voting Views Differ Sharply: Experts Worry More about Errors in E-Voting than Does the Public. *Computerworld. 2004*(August 6), unknown. Retrieved from www.computerworld.com/printthis/2004/0,4814,95094,00.html on 11/2/2005.

Dale, Markus, "Static analysis of the VoteHere VHTi reference implementation source code using Flawfinder and RATS," CMSC-691 Information Assurance Project Report, University of Maryland, Baltimore County (December 28, 2005), 14 pages.

Maryland State Board of Elections. 2006. Description of roles of the SBE and the Local Boards of Elections. Email communication, January, 2006.

Musa, J. D. (1998). *Software Reliability Engineering*. New York: McGraw-Hill.

Neff, Andrew C., ``Practical high certainty intent verification for encrypted votes,'' technical document (October 14, 2004), unpublished manuscript, 24 pages.

Pynchon, Susan. (2005). Diebold touch screens don't meet disability requirements (FL). *News-journalonline.com*, June 28 2005. Accessed 9/8/2005 Available online at: http://www.verifiedvotingfoundation.org/article.php?id=6072

Rubin, Aviel D., Dan S. Wallach, Dan Boneh, Michael D. Byrne, Drew Dean, David L. Dill, Douglas W. Jones, Peter G. Neumann, Deidre Mulligan and David A. Wagner. (2005). *A Center for Correct, Usable, Reliable, and Transparent Elections (ACCURATE)*: *A Research Proposal for an NSF CyberTrust Center.*

Selker, Ted. (2004). Fixing the Vote. *Scientific American, 291*, 90-97.

Wang, Tova Andrea. (2004). *Understanding the Debate over Electronic Voting Machines.* New York: The Century Foundation.

About UMBC

Founded in 1966, the University of Maryland, Baltimore County (UMBC) is a public university located outside of Baltimore, Maryland. Fall 2005 enrollment of nearly 12,000 included 9,400 undergraduate and more than 2,000 graduate students. The University delivers an undergraduate educational experience characterized by a strong liberal arts and sciences core. Graduate programs emphasize selected areas of engineering, information technology, science, public policy, and human services. UMBC is one of 151 institutions in the Carnegie Foundation’s doctoral/research-extensive classification for major research universities.

About MIPAR

Established in 1982, the Maryland Institute for Policy Analysis and Research (MIPAR) is the premier center for applied scholarly research on significant issues of public policy at UMBC. MIPAR conducts policy studies, program evaluations, surveys, and conferences on a wide range of topics. MIPAR activities, which are supported by federal agencies, private foundations, and state and local governments, link the resources of the University with policy makers in the state and region. Within the past few years MIPAR has developed a special strength in the area of information technology and government and e-government and e-democracy. MIPAR is affiliated with the UMBC Department of Public Policy, an interdisciplinary graduate program that offers master’s and Ph.D. degrees, as well as advanced graduate certificates.

About NCSE

In cooperation with the Maryland State Board of Elections (SBE), MIPAR established the National Center for the Study of Elections (NCSE) in 2005. The goal of the NCSE is to utilize the intellectual resources of the University to address issues concerning elections, election technologies and election administration in Maryland and across the nation. Initially, NCSE will provide technical assistance and research support to the SBE in a variety of areas. UMBC faculty associated with NCSE, independently and in conjunction with the SBE, will pursue an active research agenda on a wide range of topics around elections, election technology and election administration, and will seek funding from a variety of sources to support this research. In this way, the work of the NCSE will have value and impact within the state of Maryland and nationally.

About the Authors

Donald F. Norris is Director of the Maryland Institute for Policy Analysis and Research (MIPAR) and Professor of Public Policy at the University of Maryland, Baltimore County (UMBC). He is a specialist in urban politics, public management, and the adoption, management and impacts of information technology (including e-electronic government) in public organizations. Dr. Norris has published four books and is under contract for two more (both about electronic government) due to be published in 2006 and 2007. He has published over 50 book chapters and articles in scholarly journals and nearly 100 monographs, reports and scholarly papers. Dr. Norris is editor-in-chief of the International Journal of Electronic Government Research. He holds a B.S. in history from the University of Memphis and an M.A. and a Ph.D. in government from the University of Virginia. He is the principal contact for this report and may be reached at norris@umbc.edu

Charles Nicholas is Professor and Chair of the Department of Computer Science and Electrical Engineering at UMBC, where he has been on the faculty since 1988. He received the B.S. degree from the University of Michigan-Flint and the M.S. and Ph.D. degrees from The Ohio State University. In addition to his appointment at UMBC, Dr. Nicholas has held appointments at the National Institute of Standards and Technology (NIST), and the NASA Goddard Space Flight Center. He spent academic year 1996-97 on sabbatical at the National Security Agency. Dr. Nicholas' research interests include electronic document processing, information retrieval, and software engineering. His work has been funded by a number of agencies, including NASA, Maryland Industrial Partnerships, DARPA, AFOSR, and the Department of Defense. Dr. Nicholas has served five times as the General Chair of the ACM Conference on Information and Knowledge Management (CIKM). He also twice chaired the Workshop on Digital Document Processing. Dr. Nicholas is a member of the Board of Directors of UMBC Training Centers.

Andrew Sears, co-principal investigator of this project, is Professor and Chair of the Information Systems Department at UMBC. His research explores issues related to human-computer interaction with an emphasis on accessibility. Dr. Sears' research has been supported by numerous companies, government agencies and foundations including IBM, Motorola, NSF, the U.S.Department of Education, and the Verizon Foundation. He is on the editorial board of several journals including the International Journal of Human-Computer Interaction and Universal Access in the Information Society. He served as Conference and Technical Program Co-Chair of the ACM Conference on Human Factors in Computing Systems (CHI 2001) and as Conference Chair for the ACM SIGACCESS Conference on Assistive Technology (Assets 2005). Dr Sears is co-editor of Human-Computer Interaction Handbook. He earned his B.S. in Computer Science from Rensselaer Polytechnic Institute and his Ph.D. in Computer Science with an emphasis on Human-Computer Interaction from the University of Maryland, College Park.

Aryya Gangopadhyay is an Associate Professor and the Graduate Program Director of the Department of Information Systems at UMBC. His expertise is in the area of database management. He has done research in data warehousing and mining as well as database applications in electronic commerce, GIS, and healthcare. Dr. Gangopadhyay has published three books on electronic commerce and database issues in GIS. He has published over 50 peer reviewed articles in journals, book chapters, and conference proceedings. He holds a B.Tech.(Honours) in Ocean Engineering and Naval Architecture from the Indian Institute of Technology, an M.S. in Computer Science from New Jersey Institute of Technology, and an M.B.A. and Ph.D. in Information Systems from Rutgers University.

Stephen H. Holden is an Assistant Professor in the Department of Information Systems at UMBC. His research interests include electronic government, the management of public sector information technology and information policy. He has published in The Information Society, IEEE Internet Computing, Public Performance and Management Review, the International Journal of Public Administration, and Government Information Quarterly. Holden was also a co-author of two reports by a National Academy of Science study team examining the privacy impacts of authentication. He holds a Ph.D. (Public Administration and Public Affairs) from Virginia Polytechnic and State University, an M.P.A. (Master of Public Administration) and a B.A. (Public Management) from the University of Maine.

George Karabatis is an Assistant Professor of Information Systems at UMBC. He holds degrees in Computer Science (Ph.D. and M.S.) and Mathematics (B.S.). His research interests are on various aspects of Information Technology related to database systems, including heterogeneous distributed databases, semantic information integration, and applications for mobile handheld devices. Prior to his current appointment he was a Research Scientist at Telcordia Technologies (formerly Bellcore) where he led several telecommunications projects involving database related technologies. His work has been published in journals, conference proceedings and book chapters.

A. Gunes Koru is an assistant professor in the Department of Information Systems at UMBC. He is a specialist in software engineering, especially in the areas of software measurement, quality, maintenance, testing, and reliability. He has published scholarly papers in the top peer-reviewed journals, conferences, and workshops of his area. He serves as a regular reviewer for top journals, such as, IEEE Software and Journal of Systems and Software and several conferences. He holds a B.S. in Computer Engineering from Ege University, Izmir, Turkey, an M.S. in Computer Engineering from Dokuz Eylul University, Izmir, Turkey, and M.S. in Software Engineering and Ph.D. in Computer Science from Southern Methodist University, Dallas, TX.

Chris Law, M.S., is a faculty research assistant in the Department of Information systems at UMBC. He is a human factors and universal design expert with 10 years of professional experience in the U.K., U.S. and Canada. At the Trace R&D Center, University of Wisconsin-Madison, he was instrumental in the development of innovative universally designed interface techniques which allow people with various types of disabilities to be able to access standard products, co-developing a set of "EZ Access" techniques which have been widely reported and integrated into academic and industry prototypes. He has taught innovative approaches to user interface design and user testing of people with disabilities at workshops, conference sessions, and directly to industry. At UMBC, Mr. Law continues his interests in furthering the concepts of universal design, studying the processes used, and necessary resources for, procurement officers, designers and others involved in the development of mainstream electronic products which are universally accessible.

John Pinkston is a Professor of Computer Science and Electrical Engineering at UMBC. He received the B.S.E. degree with highest honors from Princeton University in 1964, and the Ph.D. degree from MIT in 1967, both in Electrical Engineering. His areas of specialty include Information Theory, Coding Theory, Communications, and Information Security. He came to UMBC in 1997, following a career in government and industrial research.

Alan T. Sherman is an Associate Professor in the Department of Computer Science and Electrical Engineering at UMBC and Director of UMBC's Center for Information Security and Assurance. He is an expert in cryptology, having carried out research in algorithm design, cryptanalysis, theoretical foundations for cryptography, and applications of cryptography. He earned his M.S. and Ph.D. under Ronald L. Rivest at MIT and is a cryptographic consultant for private industry, performing security evaluations and DARPA- and other government-sponsored security research. He has testified as an expert witness in electronic voting, is an editor for Cryptologia, and is a member of Phi Beta Kappa and Sigma Xi.

Dongsong Zhang is an Assistant Professor in the Department of Information Systems at UMBC. He received his Ph.D. in Management Information Systems from the University of Arizona. His current research focuses on information personalization, mobile computing, multimedia-based e-Learning, computer-mediated communication, and intelligent systems. Dr. Zhang serves on the editorial board of the Journal of Database Management and International Journal of Mobile Communication. He is a recipient of Google Research Award. He has produced more than 50 refereed scholarly publications, including those in premier journals such as the Journal of Management Information Systems, Communications of the ACM, IEEE Transactions on Multimedia, IEEE Transactions on Systems, Man, and Cybernetics, IEEE Transactions on Professional Communication, Decision Support Systems, and Information & Management, and others.