

Submission
to
Commission on Electronic Voting

Michael McMahon

25/03/04

Table of Contents

1	Overview.....	3
2	Why is e-voting different from other kinds of IT?.....	4
2.1	Issues which an e-voting system must deal with.....	4
2.2	Does the Nedap/Powervote system deal with these issues?.....	5
2.3	Relationship between secrecy and accuracy.....	5
3	Comparison with other trusted technology.....	6
4	Comparison with existing manual system.....	8
5	Evaluating the accuracy of the proposed system.....	9
5.1	Claim 1: Testing by the test institutions.....	9
5.2	Claim 2: Testing by the Department of Environment.....	11
5.3	Claim 3: Extrapolation of test results.....	11
5.4	Claim 4: Capability to print paper copies of ballots.....	13
5.5	Claim 5: 15 years of experience without apparent problems.....	14
6	Is more testing the answer?.....	15
7	What is the answer?.....	16
7.1	Voter-verified audit trail (VVAT).....	16
7.2	Chaum system.....	16
8	Recommendations.....	18
9	About the author.....	18
10	Appendix A: Description of Chaums system.....	19
10.1	Overview of Chaums system.....	19

1 Overview

The purpose of this submission is to show:

- that the Nedap/Powervote electronic voting and counting (e-voting¹) system cannot be proven to be accurate based on the nature and extent of the testing done so far.
- that such extra testing that would be required to prove the accuracy of the system, without making changes to the system, is impractical and infeasible.
- that the system needs to be modified in such a way that accuracy can be proven without depending on elaborate tests of the equipment itself. The submission shows how this can be done.

The “accuracy” of any device or machine depends on its user being able to check without reference to the machine itself, that what it produces is verifiably correct. For example, musicians may use a tuning fork and their sense of hearing to tune a musical instrument. They then rely on their hearing, in order to measure the continued musical “accuracy” of the instrument. For this reason, the notion of “accuracy” is normally closely tied to the notion of “verifiability”.

The opposing argument is that the accuracy of a system can be established exclusively by some amount of prior testing. This may be partially true in some situations. For example, there are some kinds of technology which we are obliged to trust to some extent. These range from devices like weighing-scales used in a super-market, to evidentiary devices which can be used in court to convict a person of a crime. However, it will be shown that there are fundamental differences between these trusted devices and the kind of electronic voting system provided by Nedap/Powervote.

Trust in the accuracy of the Nedap/Powervote system depends on both the quality of the prior testing and the presumption that each system will always behave the same way in live elections as the limited amount of equipment which was tested in the lab.

The submission will show that the quality of the testing was not as high as is widely believed and also the assumption about the repeatability of test conditions in live use is simply not provable.

The submission then shows what kind of testing would have to be applied to the Nedap/Powervote system to make it provably accurate, and that this kind of testing is in reality quite impractical. Finally, the submission outlines the changes that are necessary in order to make the Nedap/Powervote system trustworthy, verifiable and therefore accurate.

Note, that what is proposed here can be considered an enhancement to the proposed system rather than its replacement with something different.

¹ Commonly understood as an abbreviation for electronic-voting.

2 Why is e-voting different from other kinds of IT?

2.1 Issues which an e-voting system must deal with

E-voting refers to the registering and collection of votes by computer (software programmed voting machine). The collected votes are brought to a count centre in an electronic module where they are aggregated and then counted by another computer. There are two general problems to be considered which potentially affect the accuracy of an e-voting system.

2.1.1 Reduction in scrutiny.

It is an irrefutable fact that electronic voting results in a reduction in transparency, and the general visibility of an election. The report of the Independent Commission on Alternative Voting Methods in the UK, found that electronic voting (of the type being used in this country) results in a “significant reduction in scrutiny”². It is hard to imagine how there can be any dispute over this, since all the individual tasks involved in running an election were formerly conducted in full view of candidates, their agents, and a plethora of election officials. In an electronic system, all of these tasks are performed invisibly by the computer with a minimal number of people involved. To take one example: before the count can commence, all the votes collected have to be mixed randomly before being numbered and then counted. In the manual system, the mixing process can be observed and judgement made about the quality of the mixing. In the computerised system, no such observation can be made, and consequently no judgement can be made on the quality of the mix or even that any random mix actually took place.

This reduction in scrutiny is a common feature of all computerised systems, when compared to their manual counterparts. So, what is the problem that distinguishes e-voting systems from other types of computer systems?

2.1.2 Requirement for Secrecy

The answer is explained by the second factor which is the *requirement for secrecy*. Normally, computer systems compensate for their lack of transparency by providing information (audit trails) which can be used to check their operation externally. The classic example, is banking systems which provide bank statements, ATM receipts etc. which allow customers to verify their banking transactions.

In the case of e-voting systems however, it is an important requirement whether they be manual or electronic, that the secrecy of the ballot be maintained. This is to eliminate the risk of vote selling or voter coercion. Secrecy of the ballot makes it not allowable to link any voters ballot with the identity of the voter. This makes it difficult to provide an end to end audit, equivalent to that of a banking system. When you combine this fact with the

² Pg10 “From paper ballot to e-voting”. Report of Independent Commission on Alternative Voting Methods

reduction in scrutiny, it becomes clear that when a voter makes their choices on a voting machine, there is no way for the voter to be sure that their vote is ever counted (or counted without being modified in some way). This is simply because the voter cannot *see* the vote and no information can be associated with it which would allow them to check later that the vote was counted, as this would compromise ballot secrecy.

Neither can the voter make any personal judgement on the security of the election process. The voter cannot look at the equivalent of the steel ballot box and make a judgement as to how secure it is, because the electronic equivalent of the ballot box, which is software running in the voting machine, cannot be seen. The only assurances the voter actually has are the opinions of others, which as will be shown later in this document, are not sufficient to deserve trust, not least because the people whose opinions are being depended upon and therefore we have to trust (including the designers and testers of the system), are not in any way legally accountable for the correct conduct of elections. As will be shown, it is even *difficult to identify* all the individuals who may be in a position to influence electronic elections.

2.2 Does the Nedap/Powervote system deal with these issues?

If an e-voting system is to be provably accurate and trustworthy, it must compensate for the two issues described above. The Nedap/Powervote system does not in any way compensate for these issues. Although, the system has some internal checks and balances, no *external* evidence (i.e. evidence not produced by the system itself) can be provided to the voter to prove that his vote was registered correctly and counted.

When a voter uses the Nedap/Powervote system, he chooses his selections by pressing buttons on a console, and when the selections are correct he presses a “CAST VOTE” button to record the vote. From this point on, the voter has no independent evidence that his vote makes it to the count centre, and ends up being counted correctly. Instead, trust in the system is predicated on whatever prior testing is done, together with certain assumptions about how those tests can be extrapolated into the future. Electronic voting systems such as the Nedap/Powervote system, are unique in this respect. I am not aware of any significant IT system which has absolutely no way for the user of the system to check externally that it is functioning correctly.

2.3 Relationship between secrecy and accuracy

It is easy for an e-voting system to be either “secret” or “accurate” because a normal computer audit trail which would link voters to their votes, could be provably accurate, but would not be secret. And conversely, a system such as this one, which has no external verifiability, may be secret, but because it has no external verifiability, cannot be proven to be accurate. In other words, achieving *both* secrecy and accuracy is difficult, and the first generation of e-voting systems have simply chosen not to provide both secrecy and accuracy because of this difficulty. However, as will be seen in section 7, some novel solutions have been developed recently which can provide both secrecy and accuracy.

3 Comparison with other trusted technology

Inevitably, comparisons will be made between the accuracy of the Nedap/Powervote system and the accuracy of other kinds of trusted technology. While, I do not claim any particular legal expertise, some observations and comparisons can be made, which illustrate the fundamental differences between e-voting systems and technology that is trusted in the legal system.

Two examples are the “Intoximeter” alcohol breath testing machines and speed-cameras for detecting vehicles exceeding the speed limit. Both of these devices produce prima-facie evidence which can convict a defendant in court and to some defined extent our legal system trusts these devices to generate evidence which cannot be tested in court in the normal way.

There are three significant differences between the trust that we place in the evidentiary devices mentioned and the trust demanded by the Nedap/Powervote election system.

The first difference is that the devices mentioned can be (and are required to be) *calibrated regularly*, using a method that is independent of the device itself. For example, the Intoximeter can be tested using a known mixture of alcohol vapour and air, and its accuracy is regularly evaluated in this way. Similarly, the evidence produced by speed-cameras (such as time lapsed photographs) is dependent on the accuracy of the clock in the camera. The accuracy of a clock is something that can be tested on a regular basis³.

This is not the case with the Nedap/Powervote election system. Apart from the initial (limited) testing done, there appears to be no requirement for regular testing of the system after it is put into service. This specific point is discussed in more detail in section 5.

Second, and more fundamentally, when a citizen is confronted with evidence produced by either an Intoximeter or a speed-camera, he/she can at that time, make a judgement on the accuracy of the evidence, and can conceivably look for independent evidence (such as a traditional blood-test) which might be usable as a defence in court. The fact that most people don’t do this, means it can be argued that the technology does work, and over time confidence in it is likely to increase. In the case of speed-cameras, accused motorists can make a judgement on whether they really were exceeding the speed-limit at the time, and they have some (however limited) scope to defend themselves if they decide to challenge the evidence. Therefore, if a belief exists that this kind of technology can actually be trusted, there is some amount of justification, due to the ability of the accused person to make an independent assessment on the accuracy of whatever was being measured.

In the case of e-voting with the Nedap/Powervote system, the voter has no independent proof that their vote was recorded correctly, and consequently, the trust in the system can never increase over time, because there is no rational basis (ie. evidence) for it.

³ a legal requirement in some jurisdictions.

Third, evidentiary devices are less prone to tampering than IT equipment based on PCs⁴. For example, evidentiary breath testing machines are small, self-contained, single-purpose units which are kept under secure conditions by the Gardai, and we can be reasonably confident that they are not likely to be tampered with. On the other hand, PCs are general purpose computing devices. They are designed to be easy to write software for, and in particular easy to *install* and *replace* software, and therefore easy to tamper with. This problem is well understood by the main players in the PC industry⁵, and various initiatives are underway towards solving the problem⁶. The critical point to understand, is that these initiatives recognise that secure systems cannot rely exclusively on the trust-worthiness of the PC software and hardware which is available today.

According to Microsoft, their technology which will deal with this problem, together with the enhanced hardware to support it, will be available within a few years. Therefore, computer applications, built *today* with PCs, must have some external verifiability built in, in order to be confident about tampering. The claim by Nedap/Powervote that their PC system is “security hardened” flies in the face of the understanding in the PC industry, as well as the everyday experience of computer users.

Furthermore, while it can be argued that the use of evidentiary devices has certain risks possibly impacting individual liberty, the benefits to society resulting from their efficiency may possibly outweigh these risks. However, no such claim can be made for an untrustworthy e-voting system. This is because whatever benefits accrue from the Nedap/Powervote system can also be gained from a trustworthy system⁷.

⁴ The IES (counting system) is a regular Personal Computer (PC).

⁵ Including Microsoft and Intel

⁶ e.g. Microsoft's “Next Generation Secure Computing Base” project: <http://www.microsoft.com/ngscb>

⁷ e.g. the one described in this submission

4 Comparison with existing manual system

Accuracy in the manual voting and counting system is derived from two general properties. First, the visibility and transparency of the operation itself gives confidence to voters and other parties because most if not all of the process can be externally observed and therefore mistakes can be noticed and corrected. For example, the recording of votes on a ballot paper is a function performed directly by the voter, and once a ballot is deposited in the ballot-box, the voter has confidence in its security. Similarly, the opening of ballot boxes at the count centre can also be observed, demonstrating that ballots make it in to the count.

The second property is the accountability and identifiability of the people who conduct the voting and counting process. What this means, is that anyone in a position to possibly influence the result or conduct of an election can be identified and held legally accountable for their actions, and this is the case whereby everyone involved, from voters to polling clerks, up to the returning officer him/herself has defined legal obligations. The threat of legal sanction for misbehaviour can only increase accuracy. Conversely, the lack of legal accountability on the individual (and unidentified) programmers of the Nedap/Powervote system raises a question-mark over the imperative for accuracy.

Of course, human error has a tendency to reduce accuracy, but it has to be pointed out that there is no known case, of the wrong result being called in an election due to human error. The checks and balances in the system ensure that results which appear to be close get the extra scrutiny which they deserve.

The old election system which has hardly changed since the foundation of the State is performed by people, using tools no more complicated than paper, pencils, elastic-bands and steel ballot boxes. It is possible to make use of technology (such as electronic calculators) but fundamentally, the counting job is performed by people who rely on their eye-sight to read the ballot papers, and who rely on their own counting ability. The ability to do recounts, tedious as that may be, does help to catch mistakes, and also ensures the honesty of the election staff by requiring different staff to recheck material which was counted by others. The voting process relies on the simple but effective security provided by private polling booths and locked ballot boxes. The presence of multiple election officials for each task helps to keep each individual honest. Overall, the process is transparent, and has stood the test of time.

The new system could not be more different. Everything mentioned above has been replaced, the people doing the counting, the paper ballots, the pencils used to write on the ballots, the ballot boxes, the elastic bands used to group ballots, are all replaced by the new voting machines, the electronic ballot modules and the counting system computer.

5 Evaluating the accuracy of the proposed system

How can the accuracy of the Nedap/Powervote system be evaluated ?

The main claims made on behalf of this system are:

1. the system has been tested by a number of reputable test institutions
2. the system has been tested internally by the Dept. of Environment
3. the test results achieved above can be extrapolated forward for all time.
4. the ability of the voting machines to print out copies of ballots cast, is a sufficient safety net, in case of concerns about the rest of the system.
5. the system has been in use over 15 years without apparent problems

Let us examine each of these claims in turn.

5.1 *Claim 1: Testing by the test institutions*

5.1.1 **Physikalisch-Technische Bundesanstalt (PTB)**

From the two published reports of the PTB, it appears that the testing they performed on the voting machine was somewhat superficial. By this I mean, the actual testing was limited to comparing the visually observable behaviour and construction of the machine with how it was expected to behave (and be constructed). Unfortunately, the real output of the voting machines, the data on the ballot modules, was not tested. Presumably, this was due to the lack of suitable tools or devices which can be used to perform such tests.

To compensate for not performing functional tests on the output of the voting machines, the PTB conducted a detailed source-code review of the voting machine software. They also used a number of software analysis tools to evaluate the software. While such analysis is useful, it is *not testing* and in no way compensates for the lack of functional tests on the actual output of the voting machines. To argue otherwise, would be like checking the accuracy of a clock, by simply taking it apart and examining all of its components instead of checking it with an externally verified time source, or like calibrating a musical instrument without using a tuning-fork or equivalent independent device.

As a result, in no way can the PTB tests be regarded as any kind of calibration or accuracy test of the voting machine and it is *highly questionable* whether PTB were entitled to sign off on many requirements in the DoE specification, e.g. Requirement #45 “a vote recorded in the primary ballot module must be the vote that the voter has cast”. No action or investigation done by PTB actually tested this requirement.

5.1.2 Testing by Electoral Reform Services

The testing by ERS, appears to have been a fairly extensive test of the election counting software, i.e. that the counting software complies with relevant electoral rules. However, the test report does not describe the individual tests so it is not possible for independent third parties to really evaluate the quality of the tests, and in particular to identify if some scenarios which should have been tested were in fact tested.

Nevertheless, publishing the complete set of ballots, after the election, can be used to audit the counting software (so long as the published set of ballots is provably accurate).

However, the one significant limitation of the ERS test, is that it was concerned exclusively with the counting part of the IES computer system (the PC used for counting of votes). One critical function of the IES, is the collection and aggregation of ballots from the ballot modules. This function was not tested by ERS.

5.1.3 Testing by other test institutions

The tests undertaken by Kema and PNO were clearly within a limited scope (environmental testing) which does not relate to the main functions of the system. Therefore, these test reports will not be discussed here.

5.1.4 Other consultants reports

None of the consultants reports considers an “end-to-end” view of the accuracy of the whole election system. Instead, they focus on specific issues in specific parts of the system.

The report written by Zerflow, is an assessment of the security of the voting machines and it makes a number of recommendations in that regard. The document is not an assessment of the accuracy of the machines or the system as a whole.

One report (of two) written by Nathean is a code-review of the counting system software. As with the PTB code-review, such analysis is a useful piece of the jigsaw, but is not in itself conclusive proof that the system is accurate. The other Nathean report makes a number of recommendations relating to specific aspects of the system, but also does not consider the overall accuracy or trustworthiness of the system.

5.1.5 Conclusion relating to independent testing

The partial testing done by the independent test institutions left a number of critical functions of the system completely untested.

5.2 Claim 2: Testing by the Department of Environment

The Department of the Environment (DoE) apparently did some tests of individual voting machines and specific IES count PCs in the form of mock elections. This can be considered a “system test” in the normal sense. However, no documentation has been made public as to the details and extent of these tests, how realistic the scenarios were, and how much equipment was actually tested.

As will be shown later, system tests only prove that the system functioned at a particular time, but do not provide proof that the system will always and forever more behave the same way.

5.3 Claim 3: Extrapolation of test results

The biggest problem for the Nedap/Powervote system is that because the main function of the system cannot be tested when it is in live use, its accuracy depends on identical conditions applying when in live use as compared to when tested in the lab. This is problematic from a number of perspectives:

5.3.1 Risk of unpredicted hardware failure modes

The designers of the system have clearly anticipated some hardware failure modes, for example the failure of part of the ballot storage area. For this reason, they store each ballot in four different places in the ballot module hardware. However, it is an entirely different matter to claim that if the hardware fails, it is *guaranteed* to do so, in a way that will be detected, and will not cause incorrect information to be written to the ballot modules. Such claims are not provable, and there has not been any attempt to prove this claim with respect to this system. As will be discussed later, the claims about years of trouble-free use do not help in this regard either.

5.3.2 Risk of intentional software error

Software is written in a human-readable computer language, but before it can be distributed, installed, tested or used on a computer it must be translated into “binary” form. In binary form, software is *unreadable* by people. Furthermore, once installed on a system, software is not only unreadable, but it is *invisible* also.

This leads to a large class of risks relating to malicious tampering. In the absence of external verifiability of a system, it can not be proven that a system has not been tampered with.

Since the IES PC is built using widely used tools and technologies (Microsoft Windows, Access, Delphi etc.), there are literally tens of thousands of programmers around the world, who have the ability to write software, which once installed would not be seen and would be *highly likely* to be able to interfere with the operation of the IES.

The only challenge for hackers is how to get such malicious software installed on to the PC. This implies that the accuracy of the counting system depends on a complicated chain of trust, which is used to ensure that no unintended software is loaded onto the IES PC.

The links in this chain include items like ensuring no unauthorised access to a PC:

- before and during an election
- during transportation to/from storage
- during storage between elections

For example, if someone gains access to one of these PCs while they are being stored in between elections, it would be very easy to install a program which would not be noticed and would activate itself whenever the PC is next used in an election. One very simple attack, would be to replace the Windows serial-port driver⁸ with a new version that could modify the incoming data (votes) from the PRU in any arbitrary way. None of the consultants have made an attempt to evaluate this risk.

The unreadability of software leads to the following kinds of risks. The returning officers (and the DoE) cannot be sure that the software they receive to be installed on the PC, corresponds in any way with the software source-code which was reviewed by the consultants. This exact problem has occurred recently in California, where the company Diebold is being investigated for installing uncertified software on some voting machines. Since software cannot be seen once it is installed, the presence of uncertified software might never be detected. This is a potential risk on any software controlled devices including both the Nedap voting machines and the IES count PC.

Software tampering would probably take the form of subtle changes in behaviour triggered only when a system is in live use, or else triggered by a specific date in the calendar (timed to coincide with a particular election). None of the testing done, including system-tests could deal with this risk adequately.

While it is clearly against the interests of the owners and senior management of the vendors to engage in deliberate tampering of the software which they deliver for the count system (or the voting machines), this is not necessarily the case with employees who do not have the same stake in the reputation of the company, and it is the employees who are more in a position to tamper with the software than the senior management. Incidents where disgruntled employees have deliberately sabotaged a company's products have been documented⁹ and while it may be unlikely to happen, the exact risk cannot be quantified.

The much smaller number of election staff at the count centre, would be in a position to replace part or all of the correct software with implementations they had developed or

⁸ Or whichever driver is responsible for accessing the PRU.

⁹IT Sabotage, by Marcie Terman: Commercial Risk/Financial Systems Autumn 2002:
http://www.datafort.co.uk/commercial_risk.pdf

acquired themselves. As a relevant example here, a programmer who is involved in the Irish Citizens for Trustworthy E-voting (ICTE), used the same tools that were used by Powervote, to develop a simple election software application which shows how easy it is for software to appear to behave one way, but in reality behave totally differently. According to the developer himself, this sample application took 2 hours to write and test.

5.3.3 Risk of unintentional software error

The main difference between intentional and unintentional software error is the motive (or lack of motive) of the person or persons responsible. Therefore, the scenarios above could still happen but would be the result of accidental error rather than deliberate sabotage.

Other unintended risks could be software errors introduced in follow-up versions of the system. It is highly likely that future updates to the system software will be scrutinised (and tested) less than the initial version. Therefore, the possibility that errors may be introduced in a future update have to be considered. It is impossible to quantify this risk exactly, even though it could be significant.

5.3.4 Accountability risks

It should also be clear that formerly, trust in the election system was vested in identifiable and accountable individuals who are directly involved in the election. Now with this electronic system, trust is vested in the programmers (who apparently cannot be identified) and the testers, both of whose jobs were done many months if not years before a particular election takes place. This makes it impossible to hold these people accountable in case something goes wrong.

Indeed, there is no way that the supplier of this system, which is a relatively small company, could possibly indemnify the State against the effect of a botched election.

These are only a flavour of the kind of risks the Nedap/Powervote system will be subject to. It should be clear that these risks are really not quantifiable since they depend ultimately on knowing the minds of the people who make the mistakes (deliberately or otherwise) and therefore the assumption that the system will behave the same way in live use as compared with during the limited amount of system testing, is impossible to prove.

5.4 Claim 4: Capability to print paper copies of ballots

Nedap/Powervote claim that the capability for paper copies of the electronic ballots to be printed out is a sufficient form of auditing of the electronic system. There are a number of problems with this.

First, this form of audit would only cover the aggregation and counting of votes. It does not cover the voting procedure itself. The individual voter would still have no proof that

his vote was recorded correctly. In other words, the software on the voting machines would still be free to modify any votes that it sees fit. The paper ballots printed out are only a reflection of the voting machines view of the votes.

Second, a court-order will be required to make this happen, and no rational criteria exist for predicting when such a court-order ought to be granted. It is a mistake to assume that a close result should be one such scenario. In the old system, mistakes were more likely to be small than large, and a close result could easily have resulted incorrectly from a small counting error. Conversely, a large margin of victory is less likely to be wrong because larger errors are less likely than small ones. However, with computers there is no way to predict the magnitude of errors. They can just as easily be small or large. Some argue that they are more likely to be large. Therefore, a mistake may be just as likely to affect an apparently large margin of victory as a small one.

Furthermore, since a court order will be required to trigger such a manual count, it will certainly not be a routine occurrence.

5.5 Claim 5: 15 years of experience without apparent problems

It should be clear that this claim does not stand up to scrutiny. This is because every occasion that the system is used live, is an occasion where it could not be verified to be correct. If the system has counted 70 million votes in real elections, then no-one can prove what, if any percentage of those votes were incorrectly counted.

Since the Irish election system is (very nearly) unique, special software was written for both the voting machines and the IES PC. Therefore, this is effectively a totally new system as far as the risks and concerns relating to software are concerned.

6 Is more testing the answer?

One possible solution would be to commission more testing. In order to get a similar level of trust as exists with evidentiary technology then every voting machine and every ballot module would have to be tested regularly using a test device designed specially for the purpose. This is certainly conceivable, but since over 6000 voting machines and modules exist, there would be considerable cost and effort involved. Also, suitable test devices would have to be designed and purchased from the supplier.

More difficult is how to calibrate the IES and test the system as a whole. In reality, the only completely realistic test is to do a large scale system test, involving all voting machines, where a large number of votes would be entered (manually). These votes would then be transferred to the IES and counted. Because in this situation, the election would be a test, the results could be compared with the original data. This kind of test would have to be done shortly before every election to ensure all equipment is functioning correctly and tampering would not be feasible subsequently. Clearly, the time and effort required to do this makes it quite impractical.

7 What is the answer?

Two different approaches are described here which can solve the problem. The first is the conventional (and widely promoted) voter-verified audit trail (VVAT). The second is an ingenious system recently devised by the US mathematician Dr. David Chaum.¹⁰

7.1 *Voter-verified audit trail (VVAT)*

With a VVAT, the voting machine prints a paper copy of the ballot before the voter confirms her vote electronically. She verifies that the paper copy is a true reflection of her choices. She then electronically confirms her vote, and posts the paper copy of the ballot into a traditional ballot box. If, for some reason the printed ballot is wrong, the voter can cancel it, and repeat the process. The machine operator will have an additional role in ensuring that the voter does not take the paper ballot away from the polling booth.

In this scenario, the paper ballot is the definitive record of the voters choices (because it has been personally verified by the voter) and in the case of any discrepancy, the paper record takes precedence over the electronic one.

The auditing process could be done manually, or with machines which optically scan the printed ballots, or more likely a combination of both, with machine scanning used mostly for speed and accuracy and some manual checks done as an audit of the scanning machine. Manual checks can also be done quickly with computer assistance.

In this situation, you would have at least two (or possibly three) independent systems checking against each other, and the likelihood is that discrepancies will be found quickly. All the audit has to do, is to compare the entire set of paper ballots against the electronic data and verify that there are no discrepancies. Auditing of the count itself is a separate process, which will be done more or less the same way in any electronic system.

The disadvantage of this system is the cost associated with always auditing. Compared to the Nedap/Powervote system, it also has to deal with the paper ballots which have to be transported securely in their ballot boxes and delivered to the count centre where they are audited. Nevertheless, even accounting for the administrative overhead involved in such an audit, this system would still be substantially less labour intensive and more efficient than the old manual count process, while being provably accurate and trustworthy.

7.2 *Chaum system*

The system devised by David Chaum is radically different from conventional VVATs, but interestingly it uses a model of verifiability that the general public is already very familiar with. It also involves the voters in the process, in an interesting way. A detailed description of the system can be found in Appendix A.

¹⁰ Similar systems are commercially available e.g. <http://www.votehere.com>

The essence of the system is that the voter is given a receipt, which she must consciously decide to verify after leaving the polling station (when polling closes). The receipt is encrypted in a way that does not compromise ballot secrecy, i.e. when the receipt is taken out of the polling station it does not reveal the voters choices.

The system also depends on a mechanism for converting the encrypted receipts, copies of which are kept by the voting machines in electronic form, back to readable (and therefore countable) ballots. This process is under the control of a number of people called trustees. The system is entirely electronic and therefore does not require the election authorities to keep any paper based audit trail. The system is also mathematically provable, trustworthy and accurate subject to certain well-understood constraints.

Checking is done partially by the voters themselves by verifying their own votes, but also the specific function performed by the trustees is audited to ensure they have done the job correctly.

The main difficulty with the Chaum system is understanding it and explaining how it works to the election stake-holders. Also, due to its novelty no commercial implementations exist yet.

8 Recommendations

I would suggest that the Commission recommends:

- to postpone the introduction of the chosen e-voting system because it will not be possible to prove adequately that the chosen system is trustworthy without voter-verifiability built in.
- to add voter verifiability to the chosen system before it is used in another election.
- to study enhancements such as the verification system proposed by David Chaum (and other similar systems).
- To test in pilot scenarios, different verification systems before any decision on a particular system is made.

9 About the author

The author holds a BA BAI degree in Computer Engineering (TCD, 1986), a MSc. in Computer Science (TCD, 1988), and has been working in the IT industry in Ireland and Germany since 1987.

10 Appendix A: Description of Chaums system

This is a short overview of Dr. David Chaums e-voting system. The system is based on a voter-verifiable audit trail, which differs from voter-verified audit-trails, in that the voter is given a receipt, which she must consciously decide to verify after leaving the polling booth, rather than the VVAT system which requires the voter to verify a paper copy of the ballot at voting time.

What the system claims to achieve is a way for voters to check afterwards, by using a receipt that their vote has been registered and included in the count without being altered. This is done in a way that does not compromise ballot secrecy. The receipt is what makes the system verifiable, rather than verified. i.e. if a voter throws her receipt away after leaving the polling station, then she can never be sure her vote was registered. On the other hand, if she keeps the receipt and takes the trouble to check, she can satisfy herself that her vote was registered and included in the count. This is a familiar model of trust, in the sense that it is very similar to how we trust banks with our money. If we take the trouble to reconcile our bank statements with transaction records, then we can have trust in the system. It seems that even though some people do not do this, having the possibility to do it is good enough reason to trust the system.

10.1 Overview of Chaums system

An election consists of three stages, a voting stage, a tally stage and a counting stage. Chaums technique is concerned with the first two stages, the combined output of which is the so-called "tally batch" i.e. a complete set of publicly viewable (but anonymised) ballots, which are then given to the counting stage, which in turn produces the election results.

10.1.1 The voting stage

The voter makes his choices on the voting machine and a receipt is printed out which consists of two layers of translucent material. When held together the two layers form a visible and readable record of the voters preferences. The voter must check that the combined receipt is a faithful record of his voting choices. Before leaving the polling booth, the voter has to separate the two layers and surrender one of them. At this point, it can be seen that each individual layer is an unreadable jumble of dots (although some additional information is printed readably at the bottom) and it is only when the two different layers are held together can anything be seen. The layer that the voter takes out of the polling station is also remembered by the voting machine and is eventually passed on to the next stage (in digital form) along with all other voters votes. To the human eye, the receipt looks like a random jumble of dots, but in fact the voters preferences are encoded (encrypted) in the dots. This is important because at a later stage, the receipt has to be converted back to its readable form before it can be counted as a vote.

The voter leaves the polling station with his verifiable receipt and can do various checks to verify that his receipt is genuine, and in particular that it is passed on to the tally stage of the election. The latter check is possible because when polling is closed, the authorities are required to post on a web site, the entire batch of (encrypted) receipts (this is called the receipt batch). Therefore, on the web, there must be an identical copy of the receipt that each voter takes out of the polling station. The presence of the digital copy on the web is proof that the vote was registered, and was not modified in any way. Absence of the copy, is proof that the vote was not registered, or was somehow modified.

10.1.2 The Tally stage

The tally stage takes the receipt-batch, which contains an encrypted version of every voters vote, and transforms each of these in a number of steps back to a readable (to the eye) version of the ballot for counting. Each of the steps in this process is under the control of a trustee. A trustee is a human-being, who is trusted to the extent that he/she will not conspire with the other trustees to subvert the election. The system guarantees that the election cannot be subverted unless all trustees conspire together. The exact role of the trustee is to provide for the election a pair of encryption keys (public & private). Each must publicise her public key, and keep her private key secret.

Each step of the tally stage is performed in the same way, but by a different trustee, one after the other. In other words, a receipt (actually the entire batch of receipts) is first passed to trustee A, who does a partial transformation of the jumble of dots in each receipt, and passes the transformed receipts on to the next trustee who does the same thing. At the end of the process, the output of the last trustee is a batch of readable ballots (the tally batch). Of course, it is not the trustee herself who does this. Rather, it is software and hardware acting on her behalf. The job of the trustee is to provide (and keep secret) the encryption key, which is used to do the conversion. Ballot secrecy is preserved by each trustee randomly mixing the receipts before passing them on to the next trustee. By doing this, nobody can make a connection between an arbitrary receipt and the corresponding transformed (readable) ballot.

Auditing of the tally stage is done by requiring each trustee to reveal a random sample of transformations, after they have passed their batch onto the next trustee. In particular, what the trustee has to reveal is how a particular identifiable input was transformed into a particular identified output. The samples chosen are random, but they are also chosen in such a way so that no individual readable ballot produced at the end of the process, can be traced back to its corresponding receipt, thus preserving secrecy of the process. The output of the tally stage is a batch of readable ballots (the tally-batch) which contains the same number of ballots as the receipt batch, but which can be counted electronically (or manually if they are printed out). Interestingly, because the audit procedure does not (and cannot) check every stage of the transformation, there is a theoretical possibility that some incorrect transformation may go undetected. Chaum (and other mathematicians subsequently) have proven what the exact probability of this happening is, and it is extremely low.