

Submission to the inquiry into iVote

Dr Vanessa Teague

School of Computing and Information Systems
The University of Melbourne
vjteague@unimelb.edu.au

Dr Chris Culnane

School of Computing and Information Systems
The University of Melbourne
christopher.culnane@unimelb.edu.au

Dr Aleksander Essex

Department of Electrical and Computer Engineering
Western University, Canada
aessex@uwo.ca

Prof Rajeev Goré

Leader, Logic and Computation Group
Research School of Computer Science
The Australian National University
rajeev.gore@anu.edu.au

Prof J. Alex Halderman

Director, Center for Computer Security and Society
University of Michigan, USA
jhalderm@eecs.umich.edu

December 5, 2017

Contents

| | |
|--|-----------|
| 1 Responses to the terms of reference | 3 |
| 2 Introduction | 4 |
| 2.1 We use internet banking all the time. Is secure internet voting harder? | 5 |
| 2.2 Does Internet voting increase turnout? | 6 |
| 2.3 Error and fraud are possible in paper-based elections too. Is Internet voting worse? | 6 |
| 3 Verifiability | 7 |
| 4 Is the iVote protocol verifiable? | 9 |
| 4.1 What were iVote’s verification failure rates? | 10 |
| 4.2 How many votes were cancelled or deleted? | 11 |
| 4.3 Are there significant differences between iVote returns and paper returns? | 11 |
| 4.4 What are the opportunities for undetectable fraud in the current design? | 12 |
| 5 Does the iVote protocol keep votes secret? | 13 |
| 5.1 Does the iVote protocol defend voters against coercion? | 14 |
| 6 If Internet voting did go ahead, what improvements should be made? | 15 |
| 6.1 What is end-to-end verifiability? | 16 |
| 6.1.1 Does the blockchain help? | 16 |
| 6.1.2 Should NSWEC run an end-to-end verifiable protocol? | 17 |
| 6.1.3 A comparison with Switzerland | 18 |
| 6.2 Opening the source code. Are there real IP issues or security issues? | 19 |
| 6.3 Summary | 20 |
| 7 If Internet voting is discontinued, what are the other options? | 20 |
| 7.1 Electronic delivery and paper returns | 20 |
| 7.2 Electronic voting in a polling place | 21 |
| 8 Conclusion | 22 |

1 Responses to the terms of reference

1. *Whether the security of the iVote system is appropriate and sufficient.*

It is not.

2. *Whether the transparency and provisions for auditing the iVote system are appropriate.*

They are not.

3. *Whether adequate opportunity for scrutineering of the iVote system is provided to candidates and political parties.*

It is not.

4. *What improvements to the iVote system would be appropriate before its use at the 2019 State General Election.*

Complete discontinuation. Some alternative ways of serving remote or disabled voters are described in Section 7.

We respectfully suggest that the terms of reference are asking the wrong questions. Rather than, “How could we improve Internet voting?” it would be better to ask, “How could we improve election conduct?”

Possible solutions could be assessed according to criteria including convenience, accessibility, privacy, verifiability, security, cost, *etc.* Different groups of voters might be better served by different solutions—for example, a disabled voter living in Sydney might need a different voting method from a student working overseas. There are many reasonable solutions for taking advantage of the efficiency and convenience of computers without sacrificing election integrity—for example, voters could vote (in a verifiable manner) on a computer in a polling place, or could gather candidate information online and return a paper vote. There is no reason to assume that paperless Internet voting is the right solution.

The rest of this submission first addresses the terms of reference, explaining why iVote is not appropriate or adequate, then sketches some answers to the more important question of how NSW election conduct could be improved.

2 Introduction

iVote is not secure and does not keep votes private, particularly from insiders. There are a number of inside parties, such as software providers and external service providers, who are not as trustworthy as the electoral commission, yet have privileged access that would allow them to read or alter votes.

There have been serious problems in almost every conceivable aspect of iVote’s operation: externally exploitable security vulnerabilities, errors in the user interface design, noticeable deviations from the paper returns (in the Legislative Council 2015), inconsistent public remarks about the rates of verification failure, *etc.* There is still very little public information about the basic facts of the 2015 run, and even less about the WA 2017 run.

Our technical analyses of the NSW 2015 run and the Western Australian 2017 run are both available online [HT15, CEET17]. In 2015 we demonstrated that it was possible for an internet-based man-in-the-middle attacker to expose and manipulate votes. In 2017 we found that the third-party host service, with servers all over the world, could do so too.

The real problem of iVote is not that it might be insecure (though every state election instance has been), nor that undetected software errors could change the outcome (though bugs have affected NSW counting), nor that the probability of deliberate attack is high (though other democracies have already been attacked), nor that the consequences could be serious (though they obviously could). The real risk of iVote, and of Internet voting generally, is that there is no reliable way to detect these failures, so there is no way to verify whether the announced outcome is what the voters truly chose.

Kenya’s supreme court recently annulled their presidential election, citing concerns over an unverifiable electronic system. The deputy chief justice criticised the electoral commission and emphasised verifiability.

“Elections are not just about the numbers,” said Ms Mwilu. “You only get points for the answers if you show your working.” The contempt shown by the commission, she went on, could be explained only by accepting, as Mr Odinga alleged, either that the electoral IT system “was infiltrated and compromised and the data therein interfered with,” or else simply that the commission, “refused to accept that it had bungled the transmission system and were unable to verify the data.”¹

¹www.economist.com/news/middle-east-and-africa/21729376-organising-new-

Elections are perhaps the most difficult cybersecurity problem of all. It is not an ordinary cybersecurity problem, with external attackers and perfectly trusted insiders. Nor is it an ordinary risk-management problem, because the most serious risk is that a problem might remain undetected, or that the public (and the losing candidates) might refuse to accept the result if there is no evidence supporting the announced outcome.

There is no known solution for returning ballots over the Internet that is adequate for Australian government elections.

2.1 We use internet banking all the time. Is secure internet voting harder?

The two have completely different privacy requirements, allowing different risk mitigation strategies and different opportunities for detecting errors. In internet voting, the electoral commission is supposed to carefully verify your eligibility, but not know how you voted.

In internet banking, the first thing that the bank asks you to do is to verify your identity using a customer number and password: thus they know exactly who you are. The second thing that happens is that the bank gives you a receipt to show you how much money you transferred, from which account, and to whom. Thus you can check and prove that you did such and such a transaction.

If there is fraud against your account, you will eventually notice (when you run out of money), even if the bank doesn't notice immediately. If fraud or error alters your private vote, there is no immediate way to detect this.

The also have completely different notions of correctness: a bank must demonstrate to each individual that it has properly accounted for his money; the electoral commission must demonstrate to scrutineers and the public that it has properly dealt with *all* the votes.

Thus, the comparison with internet banking is totally false.

Even given their much easier privacy constraints, banking and electronic commerce lose a certain amount of money to fraud. This is an acceptable risk for them, because they save money by putting their services online. There is no equivalent calculation for elections: a small amount of fraud, or a reasonable belief that fraud happens, could do tremendous damage.

[election-will-not-be-easy-kenyas-supreme-court-explains-why-it](#)

2.2 Does Internet voting increase turnout?

There is interesting research on turnout from Switzerland, where researchers collected evidence from similar cantons in Geneva and Zurich to assess whether the turnout was higher in those with Internet voting [GS17]. They found that it was not.

The claim that iVote increased turnout in NSW rely entirely on most of iVote’s users being in the category of voters they claimed to be—specifically, those who would not be near a polling place on Election Day. Since this is impossible for the electoral commission to test, this constitutes no evidence that turnout actually improved among those groups. The apparent increase could instead have been caused by a few Russian teenagers, or a large number of Sydney residents who didn’t want to join a queue.

2.3 Error and fraud are possible in paper-based elections too. Is Internet voting worse?

Any remote voting runs a risk of coercion and manipulation—it would be much better if more voters were encouraged to vote in a polling place.

Postal voting has certainly produced documented cases of electoral fraud in the UK and strong and plausible accusations of coercion in the USA. This is a strong argument for less postal voting, but not a good argument for Internet voting. Although fraud is certainly possible with paper, the Internet makes things worse for at least 3 reasons.

1. Fraud can be automated, so one person might be able to manipulate a very large number of votes very quickly in an electronic setting.
2. An Internet system can be attacked from anywhere in the world.
3. Electronic fraud could be completely undetectable.

This applies to vote manipulation and voter impersonation.

Paper processes are not perfectly secure or reliable, but neither are computers. For example, the lost vote rate in the 2013 West Australian Senate race (1370 out of 1,348,797, slightly over 0.1%) was 100 times smaller than the verification failure rate in iVote 2015, (627 verification failures out of approximately 5000 attempts, about 10%). The WA Senate incident received much more attention because the AEC immediately told the public about the issue, and reran the election. The NSW Electoral Commission does not even seem to have understood the problem they observed—they assured the public that no verifiers had reported any anomalies.

3 Verifiability

Our existing paper-based electoral processes are carefully designed to protect the secret ballot while allowing scrutineers to observe that the election was properly conducted. Electronic processes are much harder to scrutinize. Being in the room and watching a report on a screen does not tell the observer anything meaningful about what the computer is doing. “Verifiability” has various technical definitions, but its intention is to replicate this opportunity for meaningful scrutiny of the accuracy of the election result.

This is best expressed in “This is not an urne,” Andrew Appel’s analysis of the French Internet voting system. Although he is a trained security researcher, Appel focuses not on security analysis but instead on the distinction between watching a physical voting urne and watching a picture of one on a computer screen. He explains that he does not see evidence that the system returns the correct election outcome.

In a normal French polling place (*bureau de vote*), there are many safeguards, and every safeguard is there because in the past, without the safeguard, there was cheating in elections. Many countries around the world—not just France—have experienced cheating in elections, and many countries have very similar safeguards. Therefore, it is important that the *assesseurs* can see with their own eyes that the ballot box (*urne*) is empty at the beginning of the day—because there was ballot-box stuffing in the past. They can see with their own eyes that the voter enters the voting booth (*isoloir*) alone—because in the past there was vote-selling and coercion of voters. The *assesseurs* can see that the voter deposits just one ballot in the ballot box—in fact, the ballot box is even transparent to make it easier to monitor—because in the past there was cheating. ... Therefore, when the poll workers and *assesseurs* report results at the end of the day, these results are accepted as legitimate because everyone can see and understand every part of the process. There are many safeguards in this process, every safeguard is there because without it there was cheating in the past, and every safeguard is one in which the *assesseur* participates directly.

In contrast, the process of an Internet election—this Internet election for the *Assemblée*—has no safeguards that the *assesseurs* can assess directly.

Andrew Appel

An identical comparison would apply between the scrutineering of an Australian polling place and iVote.

The Internet voting system for the French Assemblée continued for many years despite scientific criticism, but was discontinued this year ahead of the Presidential Election. Officials from the French National Cybersecurity Agency said there was an “extremely high risk” of cyber attacks. “In that light, it was decided that it would be better to take no risk that might jeopardize the legislative vote...”²

They were right: the Macron campaign was, subsequently, deliberately targeted. Fortunately the electoral system itself was no longer online.

The Kenyan presidential election has already been mentioned. Unverifiable electronic processes have contributed to serious, and unresolvable, electoral disputes in many other democracies.

In India, independent security analysis demonstrated serious weaknesses in the electronic voting machines [WWH⁺10]. In 2013 the Supreme Court ruled that all Electronic Voting machines in Indian general elections must carry a Voter-Verifiable Paper Record (VVPAT), dismissing a counterargument by the Electoral Commission of India who claimed that the machines were tamper-proof. The judgement stated:

...we are satisfied that the “paper trail” is an indispensable requirement of free and fair elections. The confidence of the voters in the EVMs can be achieved only with the introduction of the “paper trail.” [oI13]

In elections in Punjab this year, Delhi Chief Minister Arvind Kejriwal “alleged that large-scale irregularities in electronic voting machines (EVMs) might have led to transfer of about 20-25% of Aam Aadmi Partys votes...”³. He demanded a count or audit of the VVPATs, which were present in some but not all booths.

Most of the existing controversies relate to unverifiable polling-place electronic voting. This is because Internet voting is less common, not because it is more secure or more deserving of trust.

The most famous example is, of course, the US 2016 presidential election, where there remain unanswered questions over whether unverifiable

²<https://www.reuters.com/article/us-france-election-cyber/france-drops-electronic-voting-for-citizens-abroad-over-cybersecurity-fears-idUSKBN16D233>

³<http://www.hindustantimes.com/delhi-news/arvind-kejriwal-rakes-up-evm-tampering-again-seeks-poll-panel-intervention/story-3xRoonFvb2RS9K48DwLWzI.html>

voting systems in Michigan, Wisconsin, or Georgia might have produced a Presidential election outcome different from the one voters chose.

This is exactly the point: there is no evidence (as far as we know) that Russian manipulation changed the US election outcome, but nor is there evidence that the outcome is right. The same is true of the 280,000 votes that were received over iVote in 2015—there is neither evidence that they were manipulated, nor evidence that they were not. No Australian election should end up in this undetermined state.

In particular, the standard defence of the system from the NSWEC that “there is no evidence to indicate that anything went awry” is totally flawed: any hackers worth their salt would ensure that they left no trail, and many manipulations would not require breaking in to NSWEC hardware.

4 Is the iVote protocol verifiable?

The iVote protocol is not end-to-end verifiable. Although the telephone-based verification mechanism might allow detection of some vote manipulations at the voter’s end, the processes within the electoral commission are completely unavailable to independent scrutiny. There is no meaningful opportunity for independent verification that the votes that were verified by voters are the ones that go into the count, nor is there any opportunity for scrutineers or the public to verify that only the right votes are deleted. The most detailed explanation of the iVote protocol doesn’t mention vote deletion in the technical section at all [BCGG15].

This protocol places tremendous trust in a small number of officials and their computers. If an error, a piece of malware, or dishonest behavior by the official altered votes on a NSWEC computer, we do not see sound evidence that anyone else would be able to detect the problem.

We could attempt a detailed and technical explanation of how the process for reconciling, deleting and auditing votes after the close of polls occurs, and attempt to understand and explain it well enough to show that there are opportunities for undetectable error or fraud. However, the clearest demonstration that iVote is not publicly verifiable is by example: there is now overwhelming evidence that in the 2015 state election there was a high rate of verification failure, but it does not seem to have been noticed and understood, and was certainly not disclosed accurately to the public, at the time. This is exactly what a verifiable system should not allow.

4.1 What were iVote’s verification failure rates?

Even a genuinely verifiable system would not prevent cheating—it would merely provide an opportunity to estimate the rate. If iVote’s verification mechanism was sound, we could calculate what fraction of verification attempts failed, and hence estimate the overall rate of wrong votes. We could then multiply that by the total number of iVotes to estimate the size of the problem and compare it to the election margin. We would need the answer to one crucial question:

How many voters tried to verify but failed?

NSWEC announced after the election (March 2015) that “1.7% of electors who voted using iVote also used the verification service and none of them identified any anomalies with their vote.”

However, successfully retrieving an unexpected vote is only one form of verification failure. If an attacker had exploited the vulnerability we discovered (or some other vulnerability) to control the voter’s web browser and substitute their vote, then the attacker could also return to the voter an untruthful receipt number, or no receipt number at all. When the voter attempted to verify, he would be unable to retrieve any vote. At a seminar at the University of Melbourne in October 2015 (at which the VEC leadership was also present) V Teague asked a NSWEC official, “How many people tried to verify but failed to retrieve any vote?” He replied, “It was in the tens.”

A member of the NSW JSCEM asked the same official in August 2016 “Of that 1.7 per cent, how many people failed to verify?” He replied, “We had the situation where we had seven people through the course of the election hit the button on the Interactive Voice Response [IVR] system after verifying their vote to say it was not as they cast it... we were able through caller ID to identify a couple of phone numbers.” [Par16] He went on to say that two of the individuals had said they pressed the wrong button by mistake.

Later in the same hearing, the committee chair asked, “Of those who called the verification service how many failed to retrieve any vote?” The same official answered, “We looked at it and found 627 callers to the verification service out of the total of 5,300 calls had entered their credentials wrongly in some way...”. This is not an answer to the question that was asked, because voters may have failed to retrieve a vote for other reasons.

PwC’s post-implementation report mentions an incident affecting the verification service: “Verification service changes: Fix signature file, which

was preventing verification.”⁴ There is no date. It is unclear how many verification attempts were prevented before it was fixed. A careful reading of Hansard suggests these failures may not have been included in the 627.

The 627 acknowledged failures out of 5,300 attempts represent a verification failure rate of about 10%. This rate, extrapolated to all 280,000 iVotes, would have been quite enough to call into question the accuracy of the disputed Legislative Council seat, but the disputing candidate does not seem to have been informed of the failures at the time.

We are unable to find any data on verification rates, verification failure rates, or even iVote preference data from the 2017 WA iVote run.

4.2 How many votes were cancelled or deleted?

The iVote protocol allows for officials to remove some iVotes. This might happen for several reasons, such as if the person has voted at a polling place, or if they have telephoned the NSWEC, cancelled their vote, and re-registered. PwC’s post-implementation report lists as an incident, “EMA interface for multi vote removals. EMA failure when processing file of electors who had voted using or registered for the iVote system”

It doesn’t say what the failure was or how it was resolved. Nowhere does the PwC report, or any other official publication we have found, disclose the number of votes that were removed.

4.3 Are there significant differences between iVote returns and paper returns?

In the NSW 2015 state election there were noticeable differences between the first-preference results in the Legislative Council from iVote compared to paper methods. (Antony Green wrote a blog about it at the time but unfortunately the archive no longer seems to be online.) For example, parties on the leftmost side of the ballot received a much higher percentage of votes than they did via paper. The ALP, which was in a column that would not have fitted onto most voters’ first screen, received 5% less through iVote (25% vs over 30% by paper). This discrepancy has never been convincingly explained: it may have been the result of a user interface design error, a software glitch, a processing error, deliberate manipulation or external attack, demographic differences between the voter populations, or something else entirely.

⁴https://www.elections.nsw.gov.au/_data/assets/pdf_file/0020/220484/NSWEC_iVote_Post-Implementation_Report_FINAL.pdf

It would be interesting to know whether similar discrepancies were observed in WA in 2017.

4.4 What are the opportunities for undetectable fraud in the current design?

It is not only NSWEC officials who have privileged access to votes. Many other third party corporations, consultants and suppliers do too. Both of the times that the current iVote protocol has run in state elections (NSW 2015 and WA 2017) we have demonstrated an opportunity for third parties to read and manipulate others' votes.

In 2015 iVote incorporated a small program from a third-party provider called piwikpro. It was intended to provide NSWEC with harmless analytics information, such as data about how many voters scrolled all the way to the bottom of the screen, how often they backtracked, *etc.* We demonstrated that it was possible for any external attacker to interfere with the code and use it to expose and manipulate votes [HT15]. The piwikpro service actually suffered from two separate vulnerabilities, one (FREAK) already public and one (logjam) still under embargo, but known to A Halderman who was a co-discoverer. The opportunity for external attack was the result of poor security practices at piwikpro, but it is important to understand that even if their server had been secure against external attack, legitimate administrators at the company would have been able to expose and manipulate votes anyway. The link to piwikpro was removed when we notified CERT of the issue, but by then (according to the ABC) 66,000 votes had already been cast.

In 2017 iVote used a cloud-based Internet service delivery company called Incapsula to serve the iVote code. We assume this decision was made to limit the risk of distributed denial of service attacks like the one that affected the census. Unfortunately, this puts that company in a position of trust between voters and the WAEC/NSWEC. The voter's experience seemed to be a direct connection to the WA electoral commission, but was in fact running through Incapsula.

A compromised or maliciously programmed Incapsula server could easily alter the iVote code going into the voter's browser, in order to read or change votes. We also demonstrated that votes could be read without actively changing the iVote code, merely by decrypting the copy of the vote that was temporarily stored on the server.

This arrangement also allows Incapsula to link a voter's registration step (which obviously requires their name) with their subsequent voting step, if

they do these from the same computer, thus breaking one of iVote’s basic privacy protections.

Initially we thought that the Incapsula server was located in California, but we later realised that there were Incapsula servers with valid security certificates for WAEC all over the world, including some in China, North America and Eastern and Western Europe. Although there were some in Australia, there is no guarantee that all votes would be routed through those.

Ironically, we also noticed that the system was not properly configured to prevent distributed denial of service attacks against the main iVote server, which was hosted by NSWEC. This was corrected when we notified WAEC of the problem.

A Essex has posted a short talk explaining our findings: <https://www.youtube.com/watch?v=tfxzp2SuBso>

This demonstrates a general fact about security: when we trust someone we are trusting them not only for their honesty but also for their competence. With such a large number of (often foreign) companies with the ability to read or change votes, we are trusting that they are all honest, all competent and have all secured their services against external attack. We do not believe this assumption is justified.

In Estonia, the smart cards used for voter authentication in their Internet voting system seem to have been vulnerable to the ROCA attack [NSS⁺17], though there is no evidence that it was exploited. The attack would allow large-scale voter impersonation, would not require physical access to the card, and would not be detectable unless individual voters somehow noticed someone else had voted on their behalf. This vulnerability affected the voter authentication system, not the Internet voting system directly, but it has serious implications for the integrity of elections.

Many of the most serious security breaches in Australia and overseas are the result of unwisely-trusted third parties. The Red Cross breach was caused by a contractor. The sale of Medicare numbers on the Dark Web is probably another example. It is not entirely clear (to us) how this happened, but it was presumably either bad behaviour or bad security practices by someone who had legitimate access to the HPOS system.

5 Does the iVote protocol keep votes secret?

iVote’s verification service exposes the content of the vote to whoever administers that service. Note that it is not true, though it has sometimes

been claimed, that the Verification service needs the voter to enter their receipt number before vote decryption: we have looked carefully at how the votes are encrypted on the client—they can be immediately decrypted with the Verification service’s private key. Remember that NSWEC were able to telephone people who had called the verification service, suggesting that anyone who calls the service from an identifiable phone number immediately links their identity to their decrypted vote. Note also that NSWEC must have some way of linking a vote cancellation request (from a particular ID) to that person’s encrypted vote.

There is a tradeoff between complete privacy breach and impairment of the independence necessary for any semblance of real verification. If the verification service is hosted within the NSWEC, then the same officials (or the same malware) may have the opportunity to manipulate both the core voting server and the verification server; if the service is hosted outside NSWEC, some third party learns the contents of the votes in a way that is easily linked to (at a minimum) those who choose to verify. This is not easily solved. Indeed, this sort of tradeoff between genuine distribution of trust, and vote privacy, is at the heart of all good election design. This is a particularly weak solution, which neither protects vote privacy nor offers genuine verification.

5.1 Does the iVote protocol defend voters against coercion?

The telephone-based verification service could allow a voter to prove how they had voted by giving their login credentials and receipt number to a coercer, who could then call the verification service to check the vote.

The option to cancel a vote and re-vote has been used to justify the privacy invasion associated with the verification mechanism. A NSWEC submission to the Parliament of Victoria in 2014 claims that the 2015 design will, “limit voter coercion—this allows voters the ability to re-register and re-vote...”⁵ Similar assurances were given to NSW voters at election time.

There has been much controversy on this point, primarily because of the conflation of two separate issues. One side has argued (correctly) that the iVote protocol’s coercion-resistance mechanisms are not technically sound. The NSWEC and some other commentators have replied that vote coercion is not a serious social problem in NSW. We are not social scientists, so we make no comment on this last point, except that we are unaware of any supporting evidence, and that it begs the question of why the revoting

⁵https://www.parliament.vic.gov.au/images/stories/committees/emc/ifvea/Submissions/Submissions_2013/No_6_Electoral_Commission_New_South_Wales.pdf

mechanism was introduced into the protocol at all. Even if coercion does not exist now, it might appear if we make the opportunity available to overseas entities—even with postal voting, it is difficult to coerce voters from a distance.

On the technical question of whether iVote’s revoting mechanism securely allows voters to defend themselves from coercion (if it exists), the answer is clearly ‘no.’ When a voter cancels a vote and requests another, she receives a new iVote ID and PIN number, then when she casts a second vote she receives a new receipt number. The first (cancelled) vote is not deleted from the verification service, but it is supposed to be deleted from the list of votes that are entered into the count. iVote includes a second verification step, in which voters can query a NSWEC website to learn whether their receipt number corresponds to a vote that was included. In the 2015 run this website did not require login credentials such as iVote ID and PIN—it simply allowed anyone to enter a receipt number and learn whether it was included. If the vote was cancelled, and a coercer knew the receipt number, then entering the receipt number into this website would immediately expose that the coerced voter had cancelled their vote and revoted.

It would have been better to inform voters honestly and clearly that if they tried to use this mechanism to revote after being coerced, the person who coerced them could detect the update. The protocol-level weakness puts vulnerable voters at risk (if coercion exists in NSW) and the inaccurate advice that they can avoid coercion by revoting puts them at even greater risk.

6 If Internet voting did go ahead, what improvements should be made?

The simplest improvement is to make it smaller. The fewer votes it inputs into the tally, the less likely that the state election outcome is close enough for iVote problems to matter. (We would say zero is optimal.) A restriction to genuinely disabled voters would greatly reduce the risk—they are a tiny fraction of iVote users.

More transparency and honesty about the system and its properties would also be a great improvement. Disabled voters should be given the most accurate possible information about the risks to integrity and the number of people or corporations who could read their vote. Opening the source code would help.

If Internet voting does go ahead, a requirement for end-to-end verifiabil-

ity should be added.

6.1 What is end-to-end verifiability?

End-to-end verifiability is one way of allowing voters (and scrutineers) to check the accuracy of the election outcome without trusting the computers. Our nontechnical explanation is here: <https://arxiv.org/pdf/1504.03778.pdf>

And end-to-end verifiable protocol has three main opportunities for verification [BRR⁺15]:

Cast As Intended: voters, at the time of vote casting, can get convincing evidence that their encrypted votes accurately reflect their choices;

Recorded As Cast: voters or their designees can check that their encrypted votes have been correctly included, by finding exactly the encrypted value they cast on a public list of encrypted cast votes; and

Tallied As Recorded: any member of the public can check that all the published encrypted votes are correctly included in the tally, without knowing how any individual voted.

The last step is often called *universal verifiability* or *public verifiability*, because anyone can verify the proper tallying of the published votes. By contrast, only the individual voter can verify that her vote is cast as she intended without sacrificing privacy.

6.1.1 Does the blockchain help?

Most end-to-end verifiable systems, including Helios and vVote, have some sort of electronic public ledger for recording encrypted votes. It is commonly called a “bulletin board” in the e-voting literature, but the companies advocating voting on the blockchain mean roughly the same thing.⁶ All the (encrypted) votes for the election are published on the bulletin board, so everyone can check that their votes are included and everyone can verify the tally.

A public ledger of encrypted votes allows voters to check that their vote is recorded as cast (*i.e.* properly included). However, it doesn’t ensure

⁶Technically, there is a good argument that a fully distributed blockchain is not the right design for the sort of public ledger needed for elections—the vVote bulletin board design arguably represents a much better tradeoff between robustness and defence against manipulation—but this level of technical detail is probably beyond the interest of this inquiry.

that only eligible voters vote, that the encrypted vote accurately reflects the voter’s intention, or that the votes are accurately tallied.

6.1.2 Should NSWEC run an end-to-end verifiable protocol?

End-to-end verifiability is certainly necessary for Internet voting—if paperless Internet voting continues, a requirement for end-to-end verifiability should be added. These tools are useful, and do address some of the vulnerabilities of Internet voting.

However, there is no solution that protects privacy adequately and provides genuine end-to-end verifiability in a way that ordinary voters can use. Australia’s complex preferential voting system makes the problem significantly harder—this has been the subject of a substantial research effort over many years.

There are some techniques for providing mathematical proofs of proper shuffling, decryption and tallying of preferential votes—the vVote project used a method called “randomised partial checking” to prove that all the included votes were properly shuffled, followed by zero knowledge proofs that all the votes were properly decrypted.

The difficulty lies in providing evidence to voters that their vote is cast as they intended and recorded as they cast it. There are some solutions in the academic literature, but they are hard to use and do not adequately protect privacy.

One possibility is Helios, an open source end-to-end verifiable voting system developed from Ben Adida’s MIT PhD thesis.⁷ The system is used by the International Association of Cryptologic Research (IACR) for our board elections.

Although it is an end-to-end verifiable system, highly resistant to fraud given adequate verifying, there are several things that make it inappropriate for Australian elections.

- The current version doesn’t support preferential voting. I have thought a lot about how it could be adapted. I think it could be done, but would require a fair bit of work.
- Although it protects privacy well, even against a few corrupt authorities, it has the unfortunate property that you can prove how you voted if you choose to. This exposes voters to coercion. This is a well-known issue that cannot be solved in any simple way.

⁷<https://heliosvoting.org/>

- The user’s verification process is quite difficult to understand. If it isn’t done properly, it doesn’t work.

Some more general concerns about the limitations of end-to-end verifiability include the difficulty of knowing how many voters have verified, the possibility of manipulating the verification process so that some voters can’t verify, the difficulty of attributing blame if a problem occurs, coercion and vote buying, and the complexity of the resulting proof methods⁸.

Although these schemes hold great promise, there is no usable design that would work for Australian elections now available. The vVote project (on which C Culnane and V Teague were coauthors) was widely criticised for being too complex to administer and too difficult to use. An end-to-end verifiable remote voting system would be at least as much so, without the advantage of polling-place assistance to voters.

So end-to-end verifiability is necessary for preventing undetectable electoral fraud, but may not be feasible and is arguably not sufficient.

6.1.3 A comparison with Switzerland

Switzerland had three different e-voting systems, with varying levels of controversy. One has been discontinued. The other two, belonging to Swiss-Post and the canton of Geneva, are being adopted more widely throughout Switzerland. Initially many Swiss academics, notably the group at Berner Fachhochschule (<https://e-voting.bfh.ch/>), expressed strong opposition and explained that the systems, as they were then, were highly vulnerable to fraud. In some inimitable Swiss way a consensus has been reached - the current leader of the e-voting project at the Federal Chancellery is a graduate of that group. There is some very good documentation about the requirements at <https://www.egovernment.ch/fr/umsetzung/schwerpunktplan/vote-electronique/>

Some highlights from the “Déclaration d’intention pour l’introduction du vote électronique”:

- “les systèmes de vote électronique bénéficient de la vérifiabilité complète et sont certifiés”
- “des mesures particulières dans le domaine de la transparence sont mises en œuvre pour les systèmes certifiés proposant la vérifiabilité

⁸See for example the report of the US Overseas Vote Foundation: <https://www.overseasvotefoundation.org/E2E-Verifiable-Internet-Voting-Project/News>

universelle ; En particulier, la documentation et le code source des systèmes sont accessibles sur Internet...”

The canton of Geneva has engaged the Bern group to design and prototype a new system. Their design and source code are online: <https://github.com/republique-et-canton-de-geneve/chvote-protocol-poc>, <https://eprint.iacr.org/2017/325>. (The other system does not seem to have online source code, despite the apparent requirement quoted above.)

Although there is much about the Swiss process that is admirable, it would not address Australia’s technical problems. The mathematics of verification works only when the tally is the simple sum votes, not for our preferential system which entails distributing preferences in complex ways. This is inherent in the design and not easily changed.

Also, it is based on a paper mailout of codes, which are subsequently used to verify that the vote has been accurately recorded. This defeats most of the reasons for Australians to use Internet voting. It would of course be possible to send the codes by electronic means, but this breaks the security of the verification process. If the malware that tries to manipulate your vote somehow learns the verification codes, it can commit undetectable fraud that appears to verify correctly.

Although they use the term “end-to-end verifiable”, it is debatable whether the system truly has that property. The verification is dependent on a secure (i.e. secret) process for generating and mailing the codes.

6.2 Opening the source code. Are there real IP issues or security issues?

We have no expertise in IP law. However, if there are IP issues, it is only because the project was not well structured. The source code for the Norwegian and Estonian systems is available online. Swiss rules also specify open source code (though only one system seems to have published it). The ACTEC and the VEC post their counting code online.

The security issues are often cited as a reason to keep the code secret, but this is not valid. We are not claiming that opening the code will make the system perfectly secure it won’t but keeping the details secret won’t either. V Teague and some US colleagues wrote the argument out for the American case: <https://www.lawfareblog.com/open-source-software-wont-ensure-election-security>

The key phrase is “won’t ensure”—open source code is helpful, but not sufficient, for a secure electoral process.

The NSW JSCEM discussed briefly whether opening the code to trusted independent experts would be better. It would be better than nothing, but not as good as making it truly open. It would introduce genuine IP issues, because it would require people working in the field to promise not to expose technical details that were not public. “Independent” is also a very important word—most electoral commissions, including NSWEC, already hire consultants to certify their software for counting and voting. The standards are low. We have found numerous serious problems in code that was certified to be secure or correct.

6.3 Summary

Any Internet voting system should be end-to-end verifiable.

After years of research, the e-voting community has been unable to design an end-to-end verifiable system that is usable, works for preferential voting, and provides reasonable vote privacy. Such a design may be possible in the future, but significant further scientific breakthroughs are required first.

Anyone claiming to sell such a system now should be treated with great scepticism because they are claiming to deliver something that years of worldwide scientific research has been unable to produce.

No Internet voting system should be trusted for public elections if the code and the details of the design are kept secret. Opening the source code and the design to the public allows independent experts to test the vendor’s claims about the system’s security, privacy and verifiability. As has been amply demonstrated, the current design of iVote is deeply flawed and does not match the assurances that were given to the public.

Thus we strongly suggest investigating other options.

7 If Internet voting is discontinued, what are the other options?

7.1 Electronic delivery and paper returns

One way to improve access for remote voters is allowing them to download a blank ballot online, or even fill it in on their computer, and then mail or deliver it to the electoral commission. In Los Angeles, voters can download an e-sample ballot⁹, print it out, and then either mail it back or drop it in

⁹<https://www.lavote.net/home/voting-elections/voting-options/e-sample-ballot/sample-ballot>

to one of numerous secure drop-off boxes.

Of course, the vulnerabilities of remote voting apply to this method too—there is still the option to impersonate other voters or to manipulate the paper returns, but at least there is a limited opportunity to manipulate a very large number of others' votes.

7.2 Electronic voting in a polling place

Electronic voting in a polling place addresses many of the needs of disabled voters, except of course transport. There are good ways to run secure electronic elections in a polling place.

- Disabled voters could print out a paper ballot, check that it matches their intention, and put it in an ordinary ballot box to be counted manually with all the others.
- Everyone could print out a paper ballot and check that it matches their intention. An electronic tally could be conducted, but would be supported by paper evidence. This could be randomly audited, or manually recounted in the case of disputes. LA's recent project is a good example: <http://vsap.lavote.net/>.
- End-to-end verifiable systems like the Victorian vVote project are well suited to sending votes securely from a controlled environment overseas back to Australia. The system collected 1121 votes from supervised polling places, most of them from the Australian High Commission in London.

The first two could be implemented with very little extra work, by simply attaching a printer to the existing user interfaces of iVote or vVote.

There is a vast literature on methods for verifying election outcomes based on computers in a polling place. Colorado uses a rigorous statistical method called Risk-Limiting Audits [LS12]. It is designed for first-past-the-post elections, but we have some initial papers on extending the techniques to preferential elections.

NIST recently produced some principles for electronic elections. They are sufficiently abstract to be relevant to us.¹⁰

¹⁰<http://collaborate.nist.gov/voting/bin/view/Voting/VVSGPrinciplesAndGuidelines>

8 Conclusion

Technology changes rapidly, but the fundamental principles of election conduct haven't changed since ancient Athenian jurors placed similar-looking bronze voting tokens into a public urn and stood around to watch the count. Regardless of the technology being used, elections must derive public evidence of a correct result from secret ballots.

The threat of deliberate manipulation of Australian elections is real. It could come from inside or outside Australia, it might or might not be detected, and it provides a losing candidate with valid justification for refusing to accept an election result that is not supported by convincing evidence verifiable by scrutineers.

There are sensible alternatives to paperless Internet voting, including verifiable polling-place e-voting and paper returns of electronically-delivered ballots. Although they are not perfect, they represent a much better tradeoff between convenience and risk than any Internet voting system likely to be available by the next election.

References

- [BCGG15] Ian Brightwell, Jordi Cucurull, David Galindo, and Sandra Guasch. An overview of the ivote 2015 voting system, 2015. https://www.elections.nsw.gov.au/__data/assets/pdf_file/0019/204058/An_overview_of_the_iVote_2015_voting_system_v4.pdf.
- [BRR⁺15] Josh Benaloh, Ronald Rivest, Peter YA Ryan, Philip Stark, Vanessa Teague, and Poorvi Vora. End-to-end verifiability. *arXiv preprint arXiv:1504.03778*, 2015.
- [CEET17] Chris Culnane, Mark Eldridge, Aleksander Essex, and Vanessa Teague. Trust implications of ddos protection in online elections. In *International Joint Conference on Electronic Voting*, pages 127–145. Springer, 2017. Preprint on ArXiv: <https://arxiv.org/abs/1708.00991>.
- [GS17] Micha Germann and Uwe Serdült. Internet voting and turnout: Evidence from switzerland. *Electoral Studies*, 47:1–12, 2017.
- [HT15] J Alex Halderman and Vanessa Teague. The new south wales ivote system: Security failures and verification flaws in a live

- online election. In *International Conference on E-Voting and Identity*, pages 35–53. Springer, 2015. Preprint on ArXiv: <https://arxiv.org/abs/1504.05646>.
- [LS12] Mark Lindeman and Philip B Stark. A gentle introduction to risk-limiting audits. *IEEE Security & Privacy*, 10(5):42–49, 2012.
- [NSS⁺17] Matus Nemeč, Marek Šys, Petr Svenda, Dusan Klinec, and Vashek Matyas. The return of coppersmith’s attack: Practical factorization of widely used rsa moduli. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1631–1648. ACM, 2017.
- [oI13] Supreme Court of India. Civil appeal no.9093 of 2013—Dr. Subramanian Swamy vs. Election Commission of India, October 2013. <http://sci.gov.in/jonew/judis/40874.pdf>.
- [Par16] Parliament of New South Wales. Hearings before the joint standing committee on electoral matters—inquiry into the conduct of the 2015 state election., 2016. <https://www.parliament.nsw.gov.au/committees/DBAssets/InquiryEventTranscript/Transcript/9731/Hearing\%20Transcript.PDF>.
- [WWH⁺10] Scott Wolchok, Eric Wustrow, J Alex Halderman, Hari K Prasad, Arun Kankipati, Sai Krishna Sakhamuri, Vasavya Yagati, and Rop Gonggrijp. Security analysis of india’s electronic voting machines. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 1–14. ACM, 2010.