

INQUIRY INTO THE 2015 NSW STATE ELECTION

Name: Dr Vanessa Teague and Prof. Rajeev Goré

Date Received: 31/07/2015

SUBMISSION TO NSW JSCEM INQUIRY INTO THE CONDUCT OF THE 2015 NSW STATE ELECTION

Vanessa Teague,
Senior Lecturer,
Department of Computing and Information Systems,
University of Melbourne
(03) 8344 1274
vjteague@unimelb.edu.au

Rajeev Goré,
Professor,
Research School of Computer Science,
ANU, Canberra.
(02) 6125 8603
rajeev.gore@anu.edu.au

This submission addresses the security, privacy, verifiability and accuracy of electronic voting and counting in the recent NSW state election. We would be happy to discuss or expand upon these issues.

We are endorsed by the executive of CORE as Experts for the purposes of this submission. The Computing Research and Education Association of Australasia, CORE, is an association of university departments of computer science in Australia and New Zealand.

Vanessa Teague is an expert in cryptographic protocols for electronic voting. She worked (on a voluntary basis) with the Victorian Electoral Commission on their e-voting project based on Prêt à Voter.

She is proud to be on the advisory board of Verified Voting. Concluded advisory roles also include the University of Luxembourg's Secure and Trustworthy Voting Systems project and the University of Surrey's Trustworthy Voting Systems Project. She recently departed the advisory council of the US Overseas Vote Foundation's end-to-end verifiable Internet Voting project. She is an editorial board member of the USENIX Journal of Election Technologies and Systems, and a program committee member of various cryptography and electronic voting conferences including EVOTE ('14), Vote-ID ('15), RSA-cryptographers' track ('14 and '15), Applied Cryptography and Network Security ('15), Principles of Security and Trust (POST '16) and the European Symposium on Research in Computer Security ('15).

She receives funding only from The University of Melbourne and the Australian Research Council (apart from occasional one-off travel reimbursements *etc.*), including a current grant on electronic voting privacy.

Rajeev Goré is an expert in formal methods applied to vote-counting programs for complex counting schemes such as single transferable voting (STV). He is on the programming committee of the 2015 Vote-ID conference to be held in Bern in September.

Goré currently holds a DP14 grant on electronic vote-counting and the project aims to produce formally verified vote-counting programs for use in Australian elections.

Goré and Teague, with other colleagues, have applied for a DP16 grant from the Australian Research Council to explore methods of improving electronic voting and counting in Australian elections.

We are scheduled to give a seminar on our research at the Parliament House library in November 2015.

SUMMARY OF RECOMMENDATIONS

Recommendation 1: *Discontinue Internet and telephone voting.*

Recommendation 2: *For each election, each voter should get good evidence that their vote is cast in the way that they intended, and scrutineers and the public should get good evidence that all the votes are properly input and accurately tallied. Electronic voting should have a voter verifiable permanent paper record, either manually counted or publicly reconciled with the electronic data. (Though in future an end-to-end verifiable voting system like the VEC's vVote might be considered.)*

Recommendation 3: *As much as possible of the system's technical details (including source code) and documentation must be publicly available.*

Recommendation 4: *There should be a public process for initializing the randomness used in the Legislative Council count.*

Recommendation 5: *Complete electronic preference data should be made available as soon as the count is complete, and there should be a public process for reconciling it with paper evidence.*

INTRODUCTION

New South Wales has two distinct electronic processes that are trusted to handle votes: the Internet voting system, iVote, and the electronic Legislative Council counting system. Each one of these in its present form represents a significant risk of assigning a seat in the Parliament of New South Wales to someone other than the voters' true choice.

The legislative council count is relatively easily fixed by introducing ordinary transparent processes for making the source code and vote data available for public inspection, along with the possibility to audit the paper evidence. This is discussed on p.6.

The Internet voting system, iVote, was subject to serious security problems including opportunities to expose and manipulate votes. We recommend discontinuing Internet voting and insisting that all electronic voting have a voter verifiable permanent paper record.

We emphasise that we are making no claim that the security problems in iVote were exploited, that software errors affected the election outcome, or that anyone took advantage of any of the opportunities for manipulation that we describe here. But nor is there any evidence that such exploits did not occur. Our purpose is to explain how NSW electoral processes must be improved to close off these possibilities, so that future electoral outcomes are supported by verifiable evidence of their correctness.

IVOTE

Secure and usable remote electronic voting, *i.e.* Internet voting, remains an unsolved problem, and we have argued in the past that the risks outweigh the perceived benefits. The 2011 version, provided by Everyone Counts, did not even reliably produce valid ballots. The 2015 version, produced by ScytI, had serious problems relating to reliability, privacy, security and verifiability.

The abstract of Teague's security analysis with Alex Halderman is included here. The full paper is available at <http://arxiv.org/abs/1504.05646>

In the world's largest-ever deployment of online voting, the iVote Internet voting system was trusted for the return of 280,000 ballots in the 2015 state election in New South Wales, Australia. During the election, we performed an independent security analysis of parts of the live iVote system and uncovered severe vulnerabilities that could be leveraged to manipulate votes, violate ballot privacy, and subvert the verification mechanism. These vulnerabilities do not seem to have been detected by the election authorities before we disclosed them, despite a pre-election security review and despite the system having

run in a live state election for five days. One vulnerability, the result of including analytics software from an insecure external server, exposed some votes to complete compromise of privacy and integrity. At least one parliamentary seat was decided by a margin much smaller than the number of votes taken while the system was vulnerable. We also found protocol flaws, including vote verification that was itself susceptible to manipulation. This incident underscores the difficulty of conducting secure elections online and carries lessons for voters, election officials, and the e-voting research community.

Approximately 66,000 votes were cast in the days before the security problem was identified and fixed. Compare this with the difference in tallies between the last eliminated candidate and the lowest seated candidate at the last stage of the Legislative Council count, which was 3177. Successful manipulation of half that many votes would have altered the outcome of that seat.

The vulnerability was disclosed to the Australian Computer Emergency Response Team (CERT), who notified the NSWEC. In keeping with standard practices of responsible disclosure, we also notified the ABC and other trusted media outlets, under strict embargo. The story became public only after the vulnerability had been removed.

Some first-preference NSW Legislative Council vote tallies produced by iVote differed notably from those received via paper-based methods. For example, the ALP received more than 30% of the vote from every other method, but only 25% of iVotes¹. The reason for this discrepancy is unclear. We know of no way to discern whether this was a result of a donkey vote, a user interface problem, a software error, a security breach involving deliberate vote manipulation, or something else altogether.

Publicly available information about iVote is very limited, with no source code and only vague overviews of the system's structure available. There is some opportunity for voters to query a verification system to ask how their vote was recorded, but only a very poorly-described process for a limited number of participants to verify the subsequent vote processing. There could be other serious problems in the large part of the system that is unavailable to public scrutiny.

The security problems that were discovered might have had even more serious implications, but for two items of good luck:

1. The method NSWEC chose to defeat the identified attack also happened to remove a vulnerability to logjam, a similar but distinct attack that was not public until 20 May. (Fortunately Halderman was one of the discoverers of logjam and could check.) But for this,

¹ <http://blogs.abc.net.au/antonygreen/2015/04/does-electronic-voting-increase-the-donkey-vote.html>

all Internet votes would have been susceptible to manipulation and privacy breach from an attack announced after the return of the writs.

2. The overall margin in the 2015 state election was comfortable, and it is very unlikely that iVote problems could have affected which party won government. However, no-one who has watched Australian politics over the past few years could expect comfortable margins to be guaranteed.

It is entirely plausible that a serious security problem affecting hundreds of thousands of Internet votes could compromise the integrity of a future state election outcome. Security problems may be present, and exploited, without necessarily being noticed.

ALTERNATIVES:

At present no Internet voting solution exists that provides a degree of security and verifiability as good as postal voting for those who can fill in their own postal vote. Telephone voting is no better - indeed the distinction between telephones and internet-connected computers is increasingly blurred. For voters who need assistance filling in their own paper vote, the two verifiable but not Internet-based solutions below provide superior security and verifiability to any Internet voting solution now available, or likely to be available in the near future. Disabled voters' democratic rights are not improved by providing an accessible remote voting solution that does not protect the integrity or privacy of their vote as well as alternative methods.

SOFTWARE-INDEPENDENT POLL SITE ELECTRONICALLY ASSISTED VOTING

Computer-assisted voting at the polling place with a human-readable paper trail allows voters to check a permanent paper record of their vote. This simple and verifiable solution is offered in Tasmania and WA to voters who have difficulty using paper and pencil. The voter interacts with a computer, which then prints out their vote for insertion into an ordinary ballot box alongside all the other (paper) votes. This allows each voter to verify that the printout matches their intentions, then scrutineers observe the counting process just as they observe all the other paper ballots being counted.

The Victorian electoral commission runs an end-to-end verifiable electronic voting system called vVote, developed by an international team of researchers including Teague and led by the University of Surrey. This system provides the voter opportunity to gather evidence that their vote is recorded as they intended and properly included, with a public proof of the system's correct output of votes. (Scrutineers must still check that those votes are correctly input into the count.) However, the vVote verification and privacy processes are quite complex, requiring significant effort and commitment by the commission to supporting a correct and transparent process. We would not recommend vVote for the NSWEC at present.

ELECTRONIC DELIVERY AND PAPER RETURNS

We have previously suggested electronic delivery of ballot information and paper (postal) voting returns, especially for local government elections which are otherwise a significant burden on the postal service. Although this remains subject to some of the same challenges as postal voting, it at least gives voters the opportunity to verify that they send the vote they intended to send. This might help to address some of the problems that seem to need Internet voting.

RECOMMENDATIONS

This is the second time that NSW voters have been offered an electronic voting solution and told that their votes were guaranteed to be private and immune to tampering. In 2011 this guarantee was not supported by evidence; in 2015 it was shown to be false (though of course we do not know whether anyone exploited the vulnerability). Verifiability is complex and difficult to achieve even when there is a genuine will on the part of developers and administrators to achieve it. We see no reason to expect that future deployments of iVote will be any more secure, private, verifiable or even correct than prior ones.

Recommendation 1: Discontinue Internet and telephone voting.

Recommendation 2: For each election, each voter should get good evidence that their vote is cast in the way that they intended, and scrutineers and the public should get good evidence that all the votes are properly input and accurately tallied. Electronic voting should have a voter verifiable permanent paper record, either manually counted or publicly reconciled with the electronic data. (Though in future an end-to-end verifiable voting system like the VEC's vVote might be considered.)

ELECTRONIC COUNTING

The NSW electronic STV count conducts an extremely complex randomized counting process for the Legislative Council. All of its important components are unavailable to public scrutiny, including the source code and the process for initializing the electronic source of randomness.

At the time of writing, 4 months after the election, voting data for second and later preferences remains secret, despite those preferences having been crucial determinants of the outcome of the last Legislative Council seat. NSWEC has released only first-preference tallies and their own distribution-of-preferences count², refusing repeated requests from our colleagues, ourselves and others to provide the full preference data. This means that there is no opportunity for an independent assessment of whether the announced outcome of that last seat was correctly computed, what the true margin was, or with what probability

2 http://vtr.elections.nsw.gov.au/lc-home.htm#lc/state/dop/dop_cnt_001

different randomised preference distributions might have produced a different outcome. Nor is it possible to examine the iVote returns separately to understand whether security or other problems might have caused them to differ from the paper returns in a way that could have affected this outcome.

Would candidates or voters trust a paper-based vote count if scrutineers were prevented from observing either the ballots themselves or the process of counting them?

There are three important items of data or process that must be available to the public: the source code and other details of the counting software, the process for initializing the randomness used to distribute excess votes of seated candidates, and the full preference data itself. For each of these, we explain below why it should be public and why secrecy constitutes an evidence gap that represents an opportunity for error or fraud. Again we emphasise that we are not alleging that fraud or error occurred in this particular election. We are recommending improvements to the process so that it provides verifiable evidence of a correct election outcome, rather like a genuine opportunity for scrutineers to observe transparent paper-based processes.

SOURCE CODE OPENNESS

Undetected errors in STV counting code are a genuine risk, which is ameliorated by making the code available for public scrutiny. For example, Goré and the ANU Logic and Computation Group have found three bugs in the the ACT Electoral Commission's publicly available vote counting code. The ACTEC have acknowledged and fixed all bugs. For each bug, it was possible to design an election scenario in which the bug would have led to an incorrect count, though none of those scenarios had yet arisen in an actual election.

The VEC's recent decision to make their counting code available permits the same sort of analysis, which should have the same positive implications for reducing the number of bugs.

The ultimate aim should be to replace all of this in-house code with formally verified code which is provably correct with respect to its specification. This is the topic of Goré's current DP14 grant.

The correct implementation of random ballot selection is particularly challenging. This needs to be unbiased in order to be fair to all candidates, but correct electronic implementation of unbiased pseudorandom number generation is very difficult.

Since the NSW counting code has never had any kind of public scrutiny, let alone formal verification of its correctness, it is overwhelming likely to have errors.

Recommendation 3: *As much as possible of the system's technical details (including source code) and documentation must be publicly available.*

PUBLIC SEEDING OF RANDOMIZED PREFERENCE DISTRIBUTION

One subtle but important issue is in the initialization of the random process for selecting ballots to be transferred when a candidate is elected with more than a quota. Like all electronic random processes, the computation must start with some random seed values input to the program from outside. In principle, even if the code itself was a valid (pseudo-)random ballot selection, a person with knowledge or control of this random seed could use it to direct the counting computation to a desired outcome — it could be as simple as running the program multiple times, with different random seed values, until a desired outcome was produced.

The simple solution is to ensure that this initial random seed is derived in a public way from a genuinely random process, for example by throwing dice under observation by the candidates or scrutineers. The resulting random value would then serve as the seed for the electronic random process. This would then allow anyone to reuse the publicly available code and vote data to check the computation of the correct election outcome. A similar process is used in the USA for risk-limiting audits of paper ballots.³

Recommendation 4: *There should be a public process for initializing the randomness used in the Legislative Council count.*

THE IMPORTANCE OF PUBLICLY-AVAILABLE PREFERENCE DATA RECONCILED WITH PAPER EVIDENCE

This year's final elimination, of the No Land Tax candidate Peter Jones, was decided by a difference of 3177 votes after nearly 400 rounds of preference distribution. The seat was thus decided finally by preferences that flowed from Liberal and Green candidates after 9 and 2 rounds respectively of randomized distribution of those parties' excess votes. The margin represents less than 0.1% of the total votes.

A procedural mistake similar to the accidental omission of the Animal Justice Party from the iVote ballot could easily have changed such a close result. So could unnoticed software errors, such as a small systematic error in the distribution-of-preferences code, repeated nearly 400 times, or a small systematic bias in the random ballot selection.

Appendix A examines preference flows from Liberal and Green candidates that passed to Animal Justice and No Land Tax. Statistical analysis would require knowledge of the real preference data, which remains unavailable. We argue that *some* choice of random ballots would have produced a different outcome, given at least *some* possible values of that hidden data. Hence the available information, even if perfectly accurate, does not imply a

3 See Philip Stark's page at <http://www.stat.berkeley.edu/~stark/Java/Html/auditTools.htm>

unique outcome. This justifies insistence on opening the full preference data to public scrutiny for double-checking of the count.

It is particularly important that this occurs in time for the electronic data to be reconciled with the evidence of paper ballots.

Recommendation 5: *Complete electronic preference data should be made available as soon as the count is complete, and there should be a public process for reconciling it with paper evidence.*

Of course, none of this proves that the announced outcome is wrong. However, well-run electoral processes should produce evidence that they are correct, not an absence of evidence either way.

We note that other submissions have recommended updating NSW STV counting law to compute weighted transfer values of all votes rather than transferring a random fraction of votes. From the perspective of computation, this would be vastly better: such a scheme would be easier to implement, compute, verify formally and double-check.

In summary, NSW has a seat in the Legislative Council for which the election outcome is close (compared to both the overall election size and the magnitude of problems in the conduct of the election), disputed (at least informally⁴), and surprising (to at least one informed commentator⁵). There is not even enough publicly-available information to allow double-checking for accidental tabulation or counting errors, and no publicly available evidence that the seat has been assigned to the true winning candidate.

CONCLUSION

Running the Legislative Council count without opening either the source code or the vote data to public scrutiny leaves a lack of evidence of the correct outcome. It leaves the count vulnerable to software errors, procedural mistakes, biases in the generation of randomness, and deliberate fraud.

Greater trust in the electronic count can be earned by making the source code and ballot data available to the public, and following the recommendations above for publicly generated random seeding. Eventually the aim should be to replace this code with code that has been formally verified for correctness.

Internet voting is another matter, and this year's iVote run supports our earlier advice that secure, private and verifiable Internet voting is an unsolved problem. Security analysis

4 The No Land Tax party repeatedly told the media they were going to challenge: <http://www.abc.net.au/news/2015-04-02/nsw-election-no-land-tax-party-flags-legal-challenge/6367210>

5 Antony Green had predicted a No Land Tax win: <http://blogs.abc.net.au/antonygreen/2015/04/legislative-council-count-updates.html>

found that an attacker could subvert the iVote voting session, expose the vote that voter intended to cast, substitute a different vote, and sidestep the verification mechanism. Implementing the attack required some skill but no special knowledge that was not publicly available at the time.

The incumbent conservative government has been comfortably re-elected. However, the final seat in the Legislative Council was decided by a difference of 3177 votes, much smaller than the 66,000 votes cast over iVote during the time it was demonstrably vulnerable to manipulation.

It is entirely plausible that, if iVote is allowed to continue at a similar size in future, serious security problems could be exploited without detection to alter a close state election outcome.

Electronic voting could be retained *in polling places* in a transparent and verifiable fashion by incorporating a voter-verifiable permanent paper record.

We reiterate our longstanding recommendation to the NSWEC and others to discontinue Internet voting.

APPENDIX: EXAMINATION OF THE DISTRIBUTION OF PREFERENCES

We argue that *some* choice of random ballots would have produced a different outcome, given at least *some* possible values of the hidden preference data. This justifies the argument for making the full preference data available.

The number we care about is the fraction of votes that will pass to the No Land Tax (NLT) party when the last Liberal candidate is eliminated. This is just an illustrative example: the same sort of reasoning could be applied to Green preferences passing to the Animal Justice Party (AJP).

Note that NSW law requires the preservation of fractions of the immediately following preferences, but not of preferences further down the ballot's list. For example, if when a candidate is eliminated half the immediately following preferences go to the NLT party, then they must be distributed in proportion, but if there's another candidate who gets all the immediately following preferences, followed by NLT as third (or later) preference, then there is no requirement to preserve the NLT fraction.

In the official Distribution of Preferences, 1406 preferences, representing 1.7% of Holly Hughes' votes, were transferred to NLT at the second-last elimination round. Many of those votes would have passed through 9 iterations of random distribution from seated liberal candidates from a starting total of 1.8 million, though by the second-last round there were also thousands of preferences that had been passed from eliminated candidates.

It is easy to see that a higher fraction of NLT preferences at the final Liberal party elimination would have produced an NLT win if all the other numbers remained the same. In particular, if 6% of Holly Hughes' votes had gone to the NLT, this would have provided an extra 3504 votes, enough to exceed the AJP's total. We are not arguing this is a likely hypothesis. The point is to show that random selections do make a difference to the outcome in this particular election.

One possibility is that the initial fractions in the votes are (close to) the same as the official output, with 1.7% of votes that transfer to the Liberal party also having a subsequent NLT preference. Then we can ask how the fraction might have changed through the elimination process to 6%. It's certainly true that this would happen with some nonzero probability. The selection process might just happen to choose more of the votes with a No Land Tax preference further down.

There is also the opposite possibility: that the initial fraction was 6%, which declined to 1.7% through the redistribution and elimination process.

Either way, it is possible that such a change could occur, and could be made more probable by unidentified biases or errors in the software. Any discussion of probability requires knowledge of the real initial preferences, which should be made available, along with the source code, so that the count can be independently repeated.