

Gaining assurance in a voter-verifiable voting system

Eugenio Alberdi, Lorenzo Strigini
Centre for Software Reliability
City University London
London, UK
{eugenio, strigini}@csr.city.ac.uk

Kieran Leach, Peter Ryan
Computing Science
University of Newcastle upon Tyne,
Newcastle upon Tyne, UK
{kieran.leach, Peter.Ryan}@newcastle.ac.uk

Philippe Palanque, Marco Winckler
Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier
Toulouse, France
{palanque, winckler}@irit.fr

Abstract— The literature on e-voting systems has many examples of discussion of the correctness of the computer and communication algorithms of such systems, as well as discussions of their vulnerabilities. However, a gap in the literature concerns the practical need (before adoption of a specific e-voting system) for a complete case demonstrating that the system as a whole has sufficiently high probability of exhibiting the desired properties when in use in an actual election. This paper discusses the problem of producing such a case, with reference to a specific system: a version of the Prêt à Voter scheme for voter-verifiable e-voting. We show a possible organisation of a case in terms of four main requirements – accuracy, privacy, termination and ‘trustedness’– and show some of the detailed organisation that such a case should have, the diverse kinds of evidence that needs to be gathered and some of the interesting difficulties that arise.

Keywords: *assurance case, socio-technical systems, e-voting, cryptotography, trust, security*

I. INTRODUCTION

This paper presents recent progress in understanding the concrete needs of a case for an electronic voting (e-voting) system. “Case” here is understood in the sense derived from safety engineering [1], as in “safety case”, “dependability case” or “assurance case”. “E-voting system” means the combination of vote-casting, ballot transmission and counting, auditing and monitoring systems: a large scale, largely open, socio-technical system (i.e. formed by machine, people, procedures), posing important challenges for the development of a case.

E-voting has been presented as a suitable solution for many of the limitations of conventional paper-based voting technology, such as human errors in vote counting, limited usability of the voting methods, flaws in the application of voting procedures, fraud, etc. After some successful experiences in non-governmental elections, many countries are conducting

experiments in political areas as well [2]. However, e-voting systems are still subject of debate and controversies, which are slowing down the adoption of such systems on a large scale [3]. Several instances of fraud and misconduct and public exposure of shoddy practices in development and certification, for example, have decreased the trust in e-voting systems [4]. In order to increase acceptance of e-voting systems, recent research efforts have focused on making such systems auditable (or verifiable) so that all actions taken during the elections can be inspected and verified by everyone [5, 6, 7]. Most of current developments involve some kind of cryptographic protocol to ensure that basic requirements, such as privacy (i.e. no link can be made between a vote and a voter) and accuracy (i.e. no vote can be altered, deleted or invalidated) are met, without requiring the certification of large quantities of voting equipment, that is, by verifying the election rather than the computer system [8, 9]. In software engineering terms, by run-time checking rather than just by software verification, this is achieved through a high degree of transparency, using a combination of cryptography, monitoring and auditing (e.g. via the Web) to ensure privacy while detecting vote tampering. An example of this development is Prêt à Voter [10], which is used in the case study reported in this paper.

Technical arguments for using these systems have relied on the demonstrable strength of the algorithms: violations of accuracy or privacy are detected with very high probability if the implementations satisfy certain properties and voters and officials follow certain procedures [11]. Thus, if the election completes with very few errors detected, then it will have delivered the correct vote count with very high probability. To deploy such a system in an election, though, decision makers would need a “case” showing that the whole system (not only these

auditing/monitoring algorithms) is sufficiently robust in presence of realistic human behaviour, including errors and re-working of procedures for both normal use and exception handling. Such a case must be complete from four viewpoints: referring to the specific system (software implementation, machines, polling station workers, etc.) rather than design only (algorithm, procedures) in a specific environment or range of environments; covering all required properties of an election (e.g. it is easy to focus on e.g. the requirement for accuracy without adequate consideration for the somewhat conflicting requirement of privacy), and the whole process resulting in satisfaction of these properties (e.g., many papers focus on the efficacy of the safeguards for detection of integrity violations, rather than the probability of a correct final count).

In the remainder of this paper we discuss the advances and difficulties in producing such a case, with reference to the Prêt à Voter e-voting scheme.

II. PRÊT A VOTER: VOTER-VERIFIABLE E-VOTING

Prêt à Voter [10] is an end-to-end cryptographic scheme for booth-based e-voting. The key innovation of the Prêt à Voter approach is that the vote is encoded using a randomised candidate list.

Suppose that our voter is called Anne. At the polling station, Anne chooses at random a ballot form, an example of which is shown in Figure 1. In the booth, Anne makes her selection in the usual way by placing a cross in the right hand column against the candidate of choice. She separates the left and right hand strips along a thoughtfully provided perforation and discards the left hand strip. She is left with the right hand strip which now constitutes her *privacy protected receipt*, as shown in Figure 1.




Figure 1. Ballot form and detached receipt

Anne now exits the booth clutching her receipt, registers with an official and casts her receipt. The receipt is placed over a recording device (an optical reader or similar) which records the random value at the bottom of the strip and records which cell she marked with her X. The original paper receipt is digitally signed and returned to her to keep. After the receipt is recorded, the system verifies the validity of

the ballot and transmits it to a Web Bulletin Board (WBB), an append-only secure web server.

The randomisation of the candidate list on each ballot form ensures that the receipt does not reveal the way she voted, ensuring secrecy. The value printed on the bottom of the receipt (the “onion”) is the key to extraction of the vote. Buried cryptographically in the “onion” is the information needed to reconstruct the candidate order and so extract the vote encoded on the receipt. This information is encrypted with secret keys shared across a number of tellers. Thus, only a quorum of tellers acting together is able to interpret the vote encoded on the receipt.

During and after the election, voters can visit the WBB and confirm that their receipts appear correctly. Once vote-casting is over, the tellers take over and perform anonymising mixes and decryption of the receipts. All the intermediate stages of this process are posted to the WBB and are audited later.

III. VOTING REQUIREMENTS

Various lists of voting requirements have been proposed in the past when discussing e-voting systems (see [6, 12, 11, 13] amongst many others). Examples of proposed requirements include: privacy of the individual ballot, individual verifiability, ballot box security, count integrity, public verifiability, transparency, robustness, fairness, legitimacy, uniqueness. Although useful, these lists tend to be formulated in rather informal ways and mix together the actual top-level requirements with the means assumed to be used to satisfy the requirements (possibly coloured by assumptions about specific implementations).

We have opted for a smaller set of well-defined high level requirements which arguably encompass (and, to some extent, expand on) the requirements proposed by other authors and are consistent with many of the legal and regulatory requirements [14]. Our requirements are: accuracy, privacy, “trustedness” and successful termination. We define these in some detail below; and then introduce what we propose as the top claim for the case for Prêt à Voter.

Accuracy requirement. Accuracy requires that if (and when) the election system declares the election successfully completed, the final election result that it has produced will match (within reasonable margins of error) the voting intentions of all legitimate voters as they enter the ballot booth. Margins of error would be defined in terms of a statistical distribution of the error. This requirement includes some usability requirements (voters being able to record successfully whatever vote they have decided to enter), but it does not deal with

the issue whether voters vote as they “really want” (i.e. without bribing, intimidation etc.) which is dealt with in the other requirements (e.g. the privacy requirement).

Privacy requirement. For this requirement to be met, under no circumstance (not even with the connivance of the voter), shall any person gain from the election system evidence of for whom or for what the voter voted, apart from the vote tallies that the election system is required to publish.

‘Trustedness’ requirement. We use the made up neologism “trustedness” as a term that combines aspects of “trust” and “acceptability”. We say that this requirement has been satisfied if most citizens trust the election process “enough” to take part, using it as required (i.e., they will act on the assumption that the other three requirements are met), and to accept its results. This requirement needs to be stated in a way that it takes into account all the reasons why people may not want to vote because of the system.

Successful Termination requirement. We say that this requirement has been satisfied if the election system declares the election successfully completed, by a deadline specified in its requirements. This implies that the system is robust and resilient enough in the face of accidental and malicious threats that could make it abort (for instance, the election might have to abort if an attack is detected after it has changed many votes in a way that cannot be undone; or if a fault has destroyed a lot of votes already cast). That is, the election system has sufficient ability to resist, or recover from the effects of, the attacks and accidental faults that will occur in the election. So, it has a very high probability to succeed with all the above requirements being met.

Top Claim of the Case. We posited that specifying absolute numerical (probabilistic) requirements of these four categories would be difficult for the stake holders in the adoption of an election system, and thus the requirement would be stated in relative terms, i.e., the case should demonstrate that, for each requirement, the Prêt à Voter system implementation (machines, procedures, voter training, etc.) under consideration is at least as good as the standard paper based system (henceforth, POPS, short for “Plain Old Paper System”) that it replaces.

IV. STRUCTURING THE CASE

The purpose of a “case” is to show that the whole body of evidence really proves the claim made. In designing a complex argument (a case), as for any other design, it is important to follow a structure that will make the completed case mentally manageable for the writer and reader alike. Therefore we start by

illustrating the general structure that the case for Prêt à Voter may have and then proceed to some interesting examples of details of claims, sub-claims, arguments and supporting evidence that could be deployed.

The work described here is based on an interdisciplinary collaboration, including software engineers, human factor specialists and cryptography and security experts. Various activities were conducted to develop an integrated system description, which would then partly inform the construction of the case. These activities included: task modelling and HAZOP-style analyses [15], probabilistic modelling (already introduced partly in [3]), work on design rationale management [16] and a small-scale trial of Prêt à Voter, which looked at usability and acceptability aspects of the system [17].

A. Case Skeleton

Our work in the development, organisation and structuring of the case was guided by the following considerations:

- at an abstract level the main claims we want to prove about Prêt à Voter are the four requirements defined earlier,
- to obtain a comprehensible structure for a case, we organise the lower level claims of the case around identifiable functions in Prêt à Voter, which may or may not map one-to-one to physical (hardware, software, or human) components,
- to prove the claims for the 4 top requirements, with respect to POPS, one can then identify homologous functions in both systems and demonstrate an improvement with respect to each function rather than for Prêt à Voter as a whole; clearly, the specific POPS considered will determine the evidence available and the types of comparisons performed..

Additionally, for functions of the voting system that are susceptible to be attacked by adversaries or corrupted by technical faults and human error, we observed that the claim will naturally need to be structured into four “stages” or “layers”, as also typical of safety reasoning. Each layer will be associated with a sub-claim about a probability parameter. For instance, if we think of attacks aiming to affect accuracy rather than to violate privacy:

- the *attack*, “foul play”, or fault, which occurs with a certain probability,
- the *data corruption* caused by that attack (or fault); this corruption also occurs with a certain probability, conditional on the attack (or fault),
- the *detection* of the corruption; the detection also has an associated conditional probability (failure of detection may cause accuracy violations),

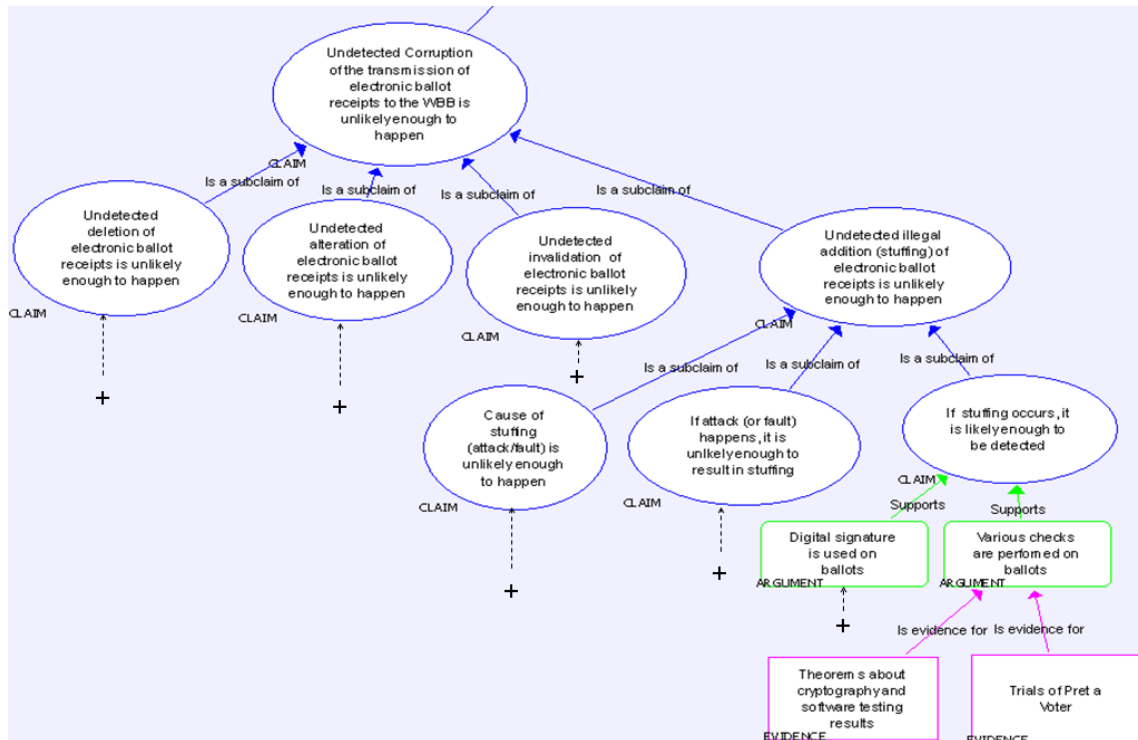


Figure 2. Partial decomposition for the sub-claim about “transmission of digital receipts” (showing part of the claim-argument-evidence tree)

- the *recovery* from the detected corruption and the conditional probability of that recovery (failure of recovery may lead to non-termination).

We will see that there might be countless such attack-corruption-detection-recovery chains, each for one category of corruption causes. Part of the challenge is to partition such scenarios into a small set of categories amenable to sound argumentation.

As noted, the top claim of our case for Prêt à Voter is decomposed into 4 top sub-claims, each referring to the satisfaction of one of the 4 voting requirements introduced above. Each sub-claim will state that, with Prêt à Voter, the corresponding requirement will be met, with high enough probability, at the election for which the case is being written. We use the phrase “with high enough probability” to mean that for an election using Prêt à Voter those requirements are, at least, as likely to be met as they would be in a conventional POPS-based election. For requirements expressed as probability distributions (e.g. of the size of numerical errors), the claim will be that Prêt à Voter is “stochastically better” than POPS.

To illustrate our claim decomposition approach, we focus now on the sub-claim concerning accuracy.

We have subdivided the accuracy claim in three further sub-claims, each relevant to a different accuracy-related component of the e-voting process, namely: 1) *Voter accuracy*, which addresses the voter’s task of filling the form at the voting booth; 2) “*Ballot box*” *integrity*, which, in Prêt à Voter, refers to what happens to the ballots while they are being scanned, digitally transmitted to and stored in the WBB before the tellers’ cryptographic transformations; the counterpart in POPS is what happens to the ballot forms while being deposited in the box and, later, while they remain stored there; 3) *Ballot count integrity*, which refers, in Prêt à Voter, to the cryptographic transformations conducted by tellers on the digital ballot receipts and the automatic count; its counterpart POPS is the manual counting of the ballots after these are taken out of the box.

Each of the above processes can be further subdivided into smaller components, and their corresponding sub-claims, until we reach subclaims for which sound arguments can be conveniently stated using available evidence. We might then still decompose each such sub-claim into sub-claims. For example, as noted earlier, many corruption detection processes can be subdivided into sub-claims corresponding to the layers of defence in the attack-corruption-detection-recovery chain introduced earlier.

B. Example: “Ballot Box” Integrity Claim

We present here an example developed for a sub-claim of the “Ballot Box” Integrity sub-claim of the top “Accuracy” claim.

“Ballot box” integrity involves, at least, two separate functions in Prêt à Voter: 1) the transmission of encrypted electronic ballot receipts from the recording device (scanner in the polling station, which digitised the ballot form filled in by the voter) to the Web Bulletin Board (WBB), which, at this stage, is our electronic “ballot box”; 2) the storage of the electronic ballot receipts in the WBB.

The description that follows focuses on the first sub-component (i.e., the one about ballot transmission). The claim we want to make here is that undetected corruption of the transmission of encrypted electronic ballot receipts to the WBB is at least as unlikely as it is for its counterpart in POPS.

In POPS the corruption of this part of the system would involve illegitimately or mistakenly removing, replacing, adding (“stuffing”), or damaging ballots as they go into the ballot box. Ballot box “stuffing” in POPS is well known. In Prêt à Voter, the corruption of this part of the system can occur in four different ways: 1) a valid electronic receipt is not transmitted to the WBB (the vote has been “deleted”); 2) the value of an electronic receipt is *altered* so that, when decrypted, it will give a different value from that entered by the voter; 3) the value of the electronic receipt is altered, but so as to make the ballot *invalid* (equivalent to making a paper ballot unreadable); 4) extra digital ballot receipts, valid in appearance, but not corresponding to a genuine vote are added (equivalent to “stuffing” the ballot box).

Figure 2 shows a possible decomposition for the sub-claim about electronic receipt transmission, taking into account the 4 possible deviations of the process. The graphical formalism is that of ASCE (Assurance and Safety Case Environment) [18], a powerful “safety case editor” that we use in this project. Sub-claims are represented by ellipsoid shapes, arguments by “rounded rectangles” and supporting evidence by plain rectangles. For each deviation, a claim is made that the undetected occurrence of the deviation is unlikely enough (i.e., at least as unlikely as the corresponding deviation in POPS). Each of these 4 sub-claims can be further decomposed into at least 3 lower sub-claims, namely, the “attack”, the “corruption” and the “detection” parts of the “attack-corruption-detection-recovery chain” discussed earlier. We do not consider the “recovery” part since detection should be enough to avoid accuracy violations.

The figure makes this decomposition explicit for only one of the possible deviations, namely, ballot “stuffing” but similar decompositions are implicit for the other 3 deviations. Similarly, the lower levels of the argumentation are again made explicit only for the “detection” sub-claim under the “stuffing” sub-claim. Although not expanded in the figure, both the “attack” and the “corruption” sub-claims are decomposed in similar ways.

The “stuffing” detection sub-claim is shown to be supported by two arguments:

1. One concerns the use of a digital signature for every ballot “onion”, a “hash” with the authority’s private key, which allows a series of corruption detection mechanisms
2. Corruption detection mechanisms in Prêt à Voter include a variety of checks (e.g., the polling station PC compares the receipt against the list of audited, cast or invalidated receipts in the WBB).

The bottom part of the tree in the figure corresponds to the evidence supporting one of the arguments for the sub-claim. For POPS, evidence would need to come from experiments about the effectiveness of detection procedures (historical records would not help, as undetected incidents would not normally be recorded). For Prêt à Voter, the evidence would come from a combination of proof of the efficacy of the cryptographic algorithms and statistical testing of the software implementation of the algorithms. For other functions that include human actions, trials would be needed to estimate the frequency of human error (including the effects of potential reduced vigilance due to reliance on computer support, to repeated false alarms, etc).

The partial case sketched so far already highlights some interesting challenges. A line of argumentation that is difficult to represent in the case is, for example, the deterrent effect that the easy detection of ballot “stuffing” corruption can have in the potential attackers of the system: because attackers know that their attack is very likely to be detected, they are less likely to attempt it. In this way the claim about detection becomes a piece of evidence supporting the claim about the likelihood of the cause of the attack occurring. However, as often happens in cases, we may want not to use this extra “comforting” argument since assessing its strength would be too difficult.

The easy detection of the kind of corruption we consider here has also implications for the top claim “successful termination”. In the discussion so far, we have not considered the “recovery” part of the “attack-corruption-detection-recovery chain” because it fits better in claims for the termination requirement

than in those for the accuracy requirement. It is important to note that easy detection (important to meet the accuracy requirement) can be in conflict with the termination requirement, e.g., the existence of good detection mechanism might prompt adversaries to attempt to corrupt data, even though they know in advance that they will be detected, just for the purpose of creating so many incidents to prevent successful termination. Detection that is effective but occurs too late for recovery actions (like cancelling an invalid ballot and asking the voter to cast a new vote) might jeopardise not only the termination but the ‘trustedness’ requirement as well.

V. FINAL REMARKS

We have presented preliminary results of a multidisciplinary collaboration towards a case to support the use of an e-voting system implementing the Prêt à Voter scheme. We identified important components of such a case and proposed a possible organisation of the case in terms of four main requirements: accuracy, privacy, successful termination and ‘trustedness’. We have also shown examples of critical sub-claims (limited, in this paper, to the accuracy requirement) and of the kinds of argumentation and evidence needed to support those sub-claims.

ACKNOWLEDGMENT

This work was partly supported by the European Union Framework Program 6 via the ReSIST Network of Excellence, IST-4-026764-NOE. Alberdi and Strigini were also in part supported by the U.K. Engineering and Physical Sciences Research Council via project INDEED, “Interdisciplinary Design and Evaluation of Dependability,” EP/E000517/1.

REFERENCES

- [1] R.E. Bloomfield, S. Guerra, A. Miller, M. Masera and C.B. Weinstock, "International Working Group on Assurance Cases (for Security)," *IEEE Security & Privacy* 4 (2006), no. 3, pp. 66- 68.
- [2] H. Geser, "E-voting projects in Switzerland. Internet Voting. Present states and future perspectives," IPSA Research Committee 05, Marburg, 2002.
- [3] J.W. Bryans, P.Y.A. Ryan, B. Littlewood and L. Strigini, "E-voting: dependability requirements and design for dependability," In *Workshop on Dependability and Security in e-Gov (DeSeGov 2006)*, at *First Int. Conf. on Availability, Reliability and Security, ARES 2006*, IEEE Computer Society Press, 2006, pp. 988-995.
- [4] P. Kocher and B. Schneier, "Insider risks in elections," *Communications of the ACM* 47 (2004), no. 7, pp. 104.
- [5] A. Antoniou, C. Korakas, C. Manolopoulos, A. Panagiotaki, D. Sofotassios, P. Spirakis and Y. Stamatiou, "A Trust-Centered Approach for Building E-Voting Systems," *E-Voting and Identity*, LNCS 4896/2007, 2007, pp. 366-377.
- [6] O. Cetinkaya and D. Cetinkaya, "Verification and Validation Issues in Electronic Voting," *The Electronic Journal of e-Government* 5 (2007), no. 2, pp. 117-126.
- [7] A. Prosser, K. Schiessl and M. Fleischhacker, "E-Voting: Usability and Acceptance of Two-Stage Voting Procedures," In *Proc. of EGOV 2007*, Springer LNCS 4656, 2007, pp. 378-387.
- [8] B. Adida and R.L. Rivest, "Scratch & vote: self-contained paper-based cryptographic voting," In *Proceedings of the 5th ACM Workshop on Privacy in Electronic Society, WPES '06*, ACM, New York, NY, 2006, pp. 29-40.
- [9] D. Chaum, "Secret-Ballot Receipts: True Voter-Verifiable Elections," *IEEE Security & Privacy* 2 (2004), no. 1, pp. 38-47.
- [10] D. Chaum, P.Y.A. Ryan and S. Schneider, "A practical, voter-verifiable election scheme," In *European Symp. on Research in Computer Security, number 3679 in LNCS*, Springer-Verlag, 2005, pp. 118-139.
- [11] P.Y.A. Ryan and T. Peacock, "Prêt à Voter: a Systems Perspective," School of Computing Science Technical Report CS-TR:929, Newcastle University, 2005.
- [12] A. Fujioka, T. Okamoto and K. Ohta, "A Practical Secret Voting Scheme for Large Scale Elections," *Proc. of the Workshop on the Theory and Application of Cryptographic Techniques: Advances in Cryptology*, Springer-Verlag, 1993.
- [13] T. Storer and R. Lock, "Accuracy: The fundamental requirement for voting systems," In *Proceedings of ARES 2009*, 2009, pp.
- [14] L.M.D. Gritzalis and S. Katsikas, "Revisiting legal and regulatory requirements for secure e-voting," In *Proceedings of the 17th International Conference on Information Security (SEC2002)*, Ghonaimy, A., M.T. El-Hadidi and H.K. Aslan (Editors), vol. 214, Kluwer, Cairo, Egypt., 2002, pp. 469-480.
- [15] N. Kaing, M. Winckler and P. Palanque, "Task Models and Hazop Analysis for Prêt-à-voter and Manual voting," IHCS-IRIT Tech. Report, May, 2008.
- [16] P. Palanque, M. Winckler, R. Bernhaupt, E. Alberdi, L. Strigini and P. Ryan, "AROVE-v: Assessing the resilience of open verifiable E-voting systems," In *7th European Dependable Computing Conference (EDCC-7) 2008*, pp. n/a.
- [17] M. Winckler, R. Bernhaupt, P. Palanque, D. Lundin, K. Leach, P. Ryan, E. Alberdi and L. Strigini, "Assessing the usability of open verifiable e-voting systems: a trial with the system Prêt à Voter," In *Proc. of ICE-GOV*, Turisat, 2009, pp. 281-296.
- [18] <http://www.adelard.com/web/hnav/ASCE>.