

A Model for System-Based Analysis of Voting Systems

Thomas Tjøstheim, Thea Peacock and Peter Y. A. Ryan

July 12, 2007

Abstract

There has recently been keen interest in the threat analysis of voting systems. While it is important to verify the system itself, it has been found that certain vulnerabilities only become apparent when taking a “system-based” view, i.e. considering interactions between the various components of a scheme. Threat analysis has so far been of three main forms: system-based, protocol-level and taxonomy check-lists. We discuss these approaches before presenting a model for system-based analysis of voting systems that is more systematic than previous work. The model is described in detail, and demonstrated with an example from a case study of the Randell-Ryan “Scratch Card” voting system.

1 Introduction

There has been a recent trend towards automated voting systems in an attempt to improve the speed and accuracy of elections, and to encourage voter turn-out. However, many of these new schemes have proven to be flawed, with cases of election fraud, e.g. in the US [12, 23]. “Black box” systems are of particular concern, e.g. those making use of Direct Recording Devices that give no proof that a vote has been correctly recorded [15]. This has generated much interest in research on verifiable voting systems, which have minimal reliance on the players, i.e., voters, election officials, etc., and technical components, such as the hardware and software behaving as intended. Notable examples are Prêt à Voter [8], Punchscan [3], and VoteHere [4], all of which aim to provide a high degree of transparency in the system. While cryptography is often used to enable verifiability without compromising voter privacy, Rivest has shown with the ThreeBallot voting system that this is not, in fact, an absolute necessity [22].

Despite the progress in developing high assurance voting systems, there is nevertheless the need for careful analysis to ensure that requirements such as eligibility, coercion-resistance and accuracy are met. In [14] Karlof et al. carried out a system-based analysis of Chaum’s visual crypto scheme [7] and Neff’s original scheme [19, 18], i.e., taking into account interactions between the various components in each scheme. In doing so, they identified potential threats such as

subliminal channels and “social-engineering”-style attacks. In a similar analysis, Ryan et al. [25] showed that Prêt à Voter [8], is robust against many of the threats mentioned in [14], but identified further possible vulnerabilities such as chain-voting and authority knowledge. See [25] for details.

Although highly useful, this type of analysis is rather ad-hoc and hence, may not uncover all the possible threats in a scheme. At a lower level of abstraction, a protocol-level analysis [20, 16] may be more systematic but is limited to the technical core of the protocol. Another approach is to develop a “catalogue of threats” [1], but perhaps as a reflection of the immensity of this task, aside from [2] there is little work to date in this direction.

In this paper, we propose a model for an analysis of threats in voting systems that is systems-based, but considerably more systematic than previous similar work [14, 25]. While [2] has a largely technical focus and concentrates on DRE systems, our model operates at a higher level of abstraction and is not scheme-specific.

In this model, the main components of a scheme such as the ballot form, voting booth, etc. are identified, and the possible threats to each component, at each phase of the protocol are considered in turn. In this way, it provides a guideline for evaluation of the system with the detail of a protocol-level analysis, but at the same time taking interactions between the various components directly into consideration. An advantage of this model is that apart from offering a more systematic approach to analysis, the components can be selected as appropriate and thus, tailored to the scheme being analysed. In addition, by working through the threat categories in the model, and at the same time applying appropriate reasoning to the scheme, the analyst is arguably better able to identify new threats than if using a catalogue of threats.

We have striven to keep the model as general as possible, hence, it can be used for a range of different systems: from manual, paper-based voting, such as the current UK system, to more sophisticated systems that incorporate, e.g. voting devices and verifiable receipts.

The structure of the paper is as follows. In Section 2 we describe the model in detail, and in Section 3, explain how it might be used to analyse a voting system. In Section 4 we discuss the results and possibilities for future work.

2 A Model for Analysis of Voting Systems

We introduce the model in a step-wise manner, beginning with a simple manual voting system, such as the one currently used in the U.K. We then extend this model to include the capability for automated vote recording and tallying, a paper audit trail and verifiability via receipts, which the voter can check against a Web bulletin board after a vote is cast. As will be seen shortly, this is done by adding the necessary components. Hence, the model offers a high degree of modularity, as the components can be selected as appropriate to the scheme being analysed.

To derive the model, we first examined the main phases in a typical voting

protocol: voting and tallying. As is the case in most current voting systems, we assume that the pre-election set-up has taken place. Typically, the electoral roll would have been established, and could include setting up of cryptographic keys, printing of paper ballot forms, ensuring that ballot boxes are empty, etc. We also assume that there is a registration process in which voters are authenticated and checked for eligibility. Note however, that a more complete analysis of a voting scheme should also include these processes. We consider them as future extensions to the model.

Taking a high-level view of the protocol, we then isolated the main components involved with each phase. The components in the model will be described in detail shortly. Working through the steps in the voting protocol, we identified possible threats that could occur directly in relation to each component. As we only consider the immediate threats, we avoid the tendency for repetition that can occur when compiling a catalogue of threats.

For uniformity, the possible threats were organised into threat categories, such as “ballot stuffing”, “absence of verifiability”, etc. Although certain threat categories do appear in several components, we only consider the threats that are directly applicable in each case. In an analysis, it is important that the details of the particular scheme be considered with care when deciding whether or not a particular threat category applies, and if so, the way in which the threat may be manifested. It is possible that not all the threat categories in the model will apply in each case, as this clearly depends on the scheme being analysed. However, identifying robustness against a particular threat is useful in highlighting the strengths of a scheme.

Note that we do not directly identify the players in a voting scheme, such as the voters and election officials. However, as will be seen shortly, many of the threats in our model can arise from interactions between certain players and the above components. Note also that we define the components in terms of the generic case, which could be adapted according to the particular system under analysis.

For example, evaluating the threats arising from storage of votes during the voting phase will depend on whether ballot forms are cast into a ballot box or whether votes are recorded on a memory card. The general threat categories are covered by the ballot storage component.

The present model excludes remote voting systems, as this adds considerable complexity to ensuring the coercion-resistance of a scheme. Instead, we start with a model for analysis of booth-based systems, and consider remote voting as a future extension.

2.1 The Base Voting Model

In the base model, we identify two phases in the protocol: voting and tallying. A description of the system, along with the main components involved at each phase, is as follows:

During the voting phase, the voter marks her choice on the ballot form in a booth, then exits and casts the marked form in a ballot box. An official ensures

that she casts one ballot form only, but should not be able to learn her vote choice or to link the voter to her cast ballot form. Here, the main components are the *voting booth*, *ballot form* and *ballot box*.

Tallying commences after the close of voting. Ballot boxes are collected and transferred to a designated tallying place. Officials open the ballot boxes and count the votes, watched by a team of observers. Local results are collated, and the final tally is published. The main components involved are the *voting booth*, *ballot form*, *ballot box* and the *election results*. As we aim for generality with the model, the components are chosen by taking a high-level view of a system.

We next describe the components in terms of their main functions and security requirements. This is necessary for determining relevant threats.

Ballot form - Record the voter’s candidate choice(s). Once the voter has marked her choice, it should not be possible to modify it. There should be no way to link a ballot form to a voter after it has been cast.

Ballot storage - Securely store the cast ballot forms. No voter should be able to cast more than one ballot form into the box, and it should not be possible to insert fake votes.

Voting booth - Provide privacy while the voter marks the ballot form. A voter should be able to make her selection without outside interference, and there should be no opportunity to monitor or record the process.

Election results - The final count of all (legitimate) votes.

Having established the basis for an analysis of threats to a voting system, we now present the first elements of our model. The possible threats associated with each component are given in Figures 1 - 4. There are several points to clarify. Firstly, note that in Figure 1 we list possible threats that could arise from both a paper ballot form and one that is generated by a device. We discuss the differences in Section 3. Secondly, in Figure 3, we show possible threats to a ballot storage component to allow a later extension for automated vote recording. Thirdly, for generality, we have included threats that would apply to schemes that are more sophisticated than the paper-based manual system, such as those that make use of encryption.

Finally, for all components in the model, the property violated is listed alongside each threat. Here, we consider the main properties required of secure systems, i.e., confidentiality, integrity and availability, rather than the traditional requirements of voting systems such as ballot secrecy, accuracy, verifiability, etc [16, 10]. The latter could be regarded as specialisations of the former, and we find that they are rather too fine-grained for a generic model. This is particularly true of voting system requirements which tend to “overlap”, such as coercion-resistance and receipt-freeness. A system may for instance satisfy receipt-freeness, but not coercion-resistance.

It is worth noting that some threats potentially violate more than one property, e.g. chain-voting and early publishing of the election results can undermine both integrity and confidentiality. However, we list these threats under integrity, which could be considered as the top-level requirement of an election: that the final count should accurately reflect the true intention of the voters. In a chain-voting attack, the coercer obtains a fresh ballot form and marks his choice. He then threatens or bribes a voter to cast it at the polling station, in return for an unused form. Hence, one or more voters may be coerced into voting in a certain way, against their free will. Partial results published ahead of time may influence voters who have not yet cast their votes. At the same time, an attacker may be able to make inferences about the identities of voters who have voted or have yet to vote. Clearly, confidentiality is also at risk in both cases.

Threat	Property violated
Identifiable information added by voter/official	Confidentiality
Voter identifiable from ballot form	
Authority knowledge	
Voter's choice incorrectly represented	Integrity
Ballot form spoiled	
Ballot form faked	

Figure 1: **Ballot form**

Threat	Property violated
Voter's activity monitored	Confidentiality
Voter records own choice	
Voter's choice influenced	Integrity
Voter smuggles out unmarked ballot form	

Figure 2: **Voting booth**

Threat	Property violated
Ballot stuffing	Integrity
Ballot spoiling	

Figure 3: **Ballot storage**

We next extend the base model by adding a voting device for automated vote recording.

2.2 Extension 1: Adding a Voting Device

A description of a generic scheme using a voting device is given below.

Threat	Property violated
Early publishing	Integrity
Absence of verifiability	
False/erroneous count	

Figure 4: **Election results**

The device authorises that the voter has the correct credentials to use the device, and then presents the vote choices to the voter. She makes her selection, e.g. on a touchscreen, which the device writes to a storage medium, such as a memory card.

We make several assumptions in this extension to the base model. Firstly, that during registration the voter optionally receives a device, e.g. smart card or one-time password which she presents to the voting machine during authorisation. Secondly, that the device is located in a booth, and the voter is checked against the electoral list during the authorisation process. Further, that the storage media are collected at the end of the voting phase.

After the voting phase has ended, officials collect the storage media from each of the voting machines. The media are transferred to a device which extracts and counts the votes. Note that although we have not identified the counting device as a component, possible threats introduced to a scheme are covered by the election results “component”. As before, the results are collated, and the final tally is published. Note that there could be a network of local counters, or a central counter. With the former, there is possibly greater opportunity for data corruption in transit or early publishing of election results.

Note that the model is still useful for schemes such as Prêt à Voter in which the device only scans the voter’s receipt, as the possible threats can be evaluated as appropriate. This is illustrated shortly in an analysis of the “Scratch Card” voting system which is based on Prêt à Voter.

The functions and requirements of the voting device and storage medium are given below.

Voting device - Authorise the voters’ credentials and present vote choices. Record the voter’s choice and write to storage media. It should not be possible to add any identifying information to a vote choice, alter, duplicate or delete it. The device should not be able to generate fake votes.

Although the function of the storage medium is analogous to the ballot storage component described previously, we re-state them in terms of the physical differences to facilitate visualisation of potential threats.

Storage medium - Store the voter’s choice. Once written to the medium, it should not be possible to alter or delete any data.

The device should be protected against any tampering. Likewise data transfer at the end of the voting period.

Potential threats introduced by adding a device are given in Figure 5.

Threat	Property Violated
Identifiable information added	Confidentiality
Voter’s activity monitored	
Faulty authorisation	Integrity
Voter’s choice incorrectly/not recorded	
Denial of service	Availability

Figure 5: **Voting device**

2.3 Extension 2: Adding a Paper Audit Trail

A second extension to the base model is a paper audit trail. With a voter verifiable paper audit trail (VVPAT) [17] mechanism, a paper copy is made of the voter’s selection and verified by the voter. The copies are securely stored as a back-up in case a manual re-count is necessary, e.g. if automated tallying fails or if the final tally appears suspicious in any way.

We assume that the voting device, such as the one in Section 2.2, produces a printed receipt for each vote cast. As in the “Mercuri Method” [17], the device displays the receipt under a clear screen. The voter verifies the receipt, which is then placed mechanically in a sealed box so that the voter cannot leave the polling station with it. The idea is that if the receipt is incorrect, an official could void the entry and provide the voter with another chance to vote. This is clearly a risk to voter privacy as indicated in the model.

Similar to a VVPAT, a verifiable encrypted paper audit trail (VEPAT) [24] acts as a paper back-up in case a manual recount of votes is necessary, but is intended for schemes in which the voter’s choice is encrypted. Since copies are made of an encrypted vote, the risk to voter privacy is reduced. Since the threats specific to a VEPAT will be covered by the ballot form component, both types can be analysed using the same component in the model.

Possible threats to a scheme arising from the paper audit trail are shown in Figure 6.

Threat	Property violated
Voter identifiable from receipt	Confidentiality
Voter’s choice noted by official	
Mismatch between voter’s choice and paper copy	Integrity

Figure 6: **Paper audit trail**

2.4 Extension 3: Adding a Web Bulletin Board (WBB) and Verifiable Receipts

A final extension to the base model is a WBB and verifiable receipts. This is to enable the analysis of schemes which allow verifiability without compromising voter privacy, such as Prêt à Voter. In such schemes the voter receives a receipt, which may bear an encrypted value, e.g. the voter's selection. Ideally, there should be mechanisms that allow the voter to check that her vote has been encrypted correctly. She later checks her receipt against a WBB to ensure that it has been correctly recorded by the system. However, the ThreeBallots scheme enables verifiability without the use of encryption. To achieve this, ballot forms are constructed in such a way that the portion retained by the voter as a receipt cannot be used as proof of a vote. Details can be found in [22].

In schemes which utilise encryption, the encrypted votes are typically passed through anonymising tabulation servers before final tallying. The final count is posted to a WBB, so can be verified by anyone. Further details can be found in e.g. [8, 7, 5, 9]. Note that the model can be used for schemes which do not use cryptography, as the possible threats can be evaluated as appropriate. This is illustrated in a forthcoming paper in which we use our model to analyse potential threats in Prêt à Voter and the ThreeBallots scheme: the former uses encryption, whereas the latter does not. The unifying requirement is that it should not be possible to link the voter's receipt to her (unencrypted) vote. However, with the possibility of verifying a receipt, coercion becomes a serious potential threat. This is identified in our model, and discussed in Section 3.

The WBB and verifiable receipt are defined below. Possible threats arising from these components are given in Figures 7 and 8.

Verifiable receipt - Enables the voter to check that her vote has been correctly recorded by the system, without compromising voter privacy. There should be proof of authenticity, such as a verifiable digital signature, so that neither the system nor the voter can falsely claim that the receipt is invalid. It should not be possible for the voter to prove her vote using the receipt.

WBB - This should be a publicly-accessible, write-only medium. The voter should be allowed access to verify that her receipt has been correctly recorded by the system. In addition, anyone should be able to verify that the intermediate decryptions of encrypted votes and/or the final tally is correct from postings to the WBB.

This completes the model, and in the next section, we discuss ways in which it may be used. Note that the model does not include certain threats such as forced abstention due to, e.g. shortage of election equipment, complicated voter registration, etc., as these are generally due to forces outside the system, and need to be addressed by means other than improvements in the protocol.

Threat	Potential threat
Voter identifiable from receipt	Confidentiality
Authority knowledge	
Receipt discarded/ surrendered	Integrity
Invalid signature	
Faked receipt	

Figure 7: Verifiable receipt

Threat	Potential threats
Monitoring access to the WBB	Confidentiality
Voter presented with fake WBB	Integrity
WBB modified	
Denial of service	Availability

Figure 8: WBB

3 Applying the Model

In this section, we describe the way in which our model could be used to identify potential threats in one of the more robust versions of the “Scratch Card” voting system [21]. This is a version of Prêt à Voter [8], which aims to promote voter understandability. The scheme offers receipt-freeness and limited voter verifiability without the use of encryption. It provides a good exemplar for an analysis as all the various components in the model can be demonstrated to their full extent.

3.1 Threat Analysis of the Randell-Ryan “Scratch Card” Voting System

An overview of the scheme is as follows. The voter randomly chooses a ballot form, an example of which is shown in Figure 9. A randomised candidate list is printed in the left hand column (LHC). Below this is a code identification number (CIN): the key to the candidate ordering. The same CIN appears at the foot of the right hand column (RHC), but is concealed with a scratch strip. Overprinted on the scratch strip is the receipt identification number (RIN).

In the privacy of the booth, the voter marks an “X” against her chosen candidate in the RHC. The LHC is detached and dropped into a clearly-marked LHC ballot box. Outside the booth, and in the presence of an official, a photocopy is made of the RHC, while the original goes into a clearly-marked RHC ballot box. The voter retains the photocopy as a receipt (see Figure 10), and can use it to check that her “encrypted” vote has been correctly recorded by the system. For example, the RIN and position of the “X” could be shown on a publicly-accessible Web bulletin board (WBB).

King	
Queen	
Knight	
Rook	
513170 (CIN)	023169 (RIN)

Figure 9: **Scratch Card ballot form**

X
023169 (RIN)

Figure 10: **Photocopied receipt**

At the close of voting, the scratch strips are removed from each RHC, revealing the CIN as shown in Figure 11. The votes can then be recovered by matching the LHCs to the corresponding RHCs. Note that the only purpose of the CIN, is to link the LHC to the RHC during the tallying phase.

Ballot auditing is carried out under the supervision of officials pre-, post- and during the election period. Voters and independent auditors take random ballot forms, scratch off the RINs and check that the CINs match the corresponding candidate order. Although the scheme boasts simplicity, the unwieldy tabulation process is a disadvantage. In addition, voters must rely on the correctness of the tabulation as the scheme does not provide verifiability of the final tally. See [21] for a discussion.

We now carry out a threat analysis of the scheme, first identifying the main components from the model: the *ballot form*, *voting booth*, *ballot storage*, *voting device*, *verifiable receipt*, *WBB* and *election results*.

3.1.1 Ballot form

Threats to confidentiality:

X
513170 (CIN)

Figure 11: **Countable vote**

- Identifiable information added - Only the RIN and the voter's mark is recorded, so unless the correspondence between CIN and candidate order is leaked by the authority or the CIN-RIN noted by an official (see below), the RHC cannot later be identified at the WBB.
- Voter identifiable from ballot form - A potential threat, if an authority notes down the CIN-RIN correspondence from the RHC and the CIN-candidate order correspondence from the LHC, a voter would be able to prove her vote to that election official. A suggested mitigation is to have independent authorities for the LHCs and RHCs, the above attack would then require the cooperation of two dishonest election officials.
- Authority knowledge is a potential threat, as information about CIN-RIN and CIN-candidate list pairings could be leaked during creation, storage and distribution of ballot forms. A possible countermeasure is to have the CINs in the LHCs covered by scratch strips, which would only be removed during the tabulation process. Note that this would protect against authority knowledge during storage and distribution of ballots, but not during creation of ballots. However, it is possible to distribute the creation of ballots, by first covering the LHC CINs with scratch strips, and getting a different group of ballot clerks to print the candidate order on the LHC.

It is interesting to note that in identifying the voter from information on a legally marked ballot form, the attacker makes use of a subliminal channel. In contrast, if a voter is identified from e.g. marks added to the ballot form by a dishonest official, the information flow is via an illegal channel.

Threats to integrity:

- Voter's choice incorrectly represented - The voter's choice could be incorrectly represented if there are multiple ballots with identical CINs. A RHC could then be incorrectly linked to a LHC with a different candidate order.
- Ballot form spoiled - A possible threat if the LHC CIN does not match the RHC CIN. However, this should be caught during both random pre-auditing and auditing during the election.
- Ballot form faked - This could be done with knowledge of how CINs are formed, but the chance of a faked ballot being caught during auditing should act as a deterrent. Anti-counterfeiting devices would be another possible mitigation against this attack. Note that [21] does not describe formation of the CINs.

3.1.2 Voting booth

Threats to confidentiality

- Voter monitored - A possible threat, e.g. with a hidden camera in the booth.
- Voter records own choice - A voter could e.g. use a camera phone, to prove the correspondence between candidate list and RIN, and later prove how she voted by showing to her scanned receipt (RHC) at the WBB.

Note that the above would be threats in almost any scheme, but should nevertheless be evaluated in an analysis.

Threats to integrity:

- Voter choice influenced - A possible threat, e.g. by a subliminal message in the booth.
- Voter smuggles out unmarked ballot form - Chain-voting is a potential threat. The coercer marks the ballot and can later check the RIN of that ballot against the WBB, to ensure that the voter has complied with his instructions.

3.1.3 Ballot box

Threats to integrity:

- Ballot stuffing - This could be carried out e.g. by corrupt officials
- Spoiling - A possible threat, e.g. ballot forms could be lost or substituted by a dishonest election official. Having a VVPAT mechanism [17] in place is a possible mitigation. However, note that the use of the WBB only ensures that the ballots enter the counting phase.

Both attacks would require a certain amount of coordination as the CINs and RINs on the faked/substituted LHCs and/or RHCs would have to be correctly matched. The suggested mitigation in [21] is for LHCs and RHCs to be handled by independent authorities.

3.1.4 Voting device

Threats to confidentiality:

- Identifiable information added by device - As the device only scans the receipt this is not a particular problem.
- Voter choice incorrectly/not recorded - This is a possible threat, but would be discovered if voters are diligent in checking their receipts on the WBB. Another countermeasure is to have a VVPAT mechanism in place.
- Voter's activity monitored - This could be carried out e.g. via wireless connection, but as long as the CIN-RIN pairings are not revealed until the time of counting, the voter's choice cannot be learned from the RHC scanned by the device.

Threats to integrity:

- Faulty authorisation - Since the device does not authorise anything, this is not a threat to the scheme.

Threats to availability:

- DoS - A possible threat, e.g. due to device failure. However, the voter does not face the possibility of losing a vote if unable to scan her receipt, as may be the case with some touchscreen voting machines.

3.1.5 Verifiable receipt

Threats to confidentiality:

- Voter identifiable from receipt - This is not a threat (assuming correct operation of the scheme) as the LHCs and RHCs are both cast at the time of voting, and the voter cannot prove correspondence between RIN and candidate order. However, this is a potential threat if a corrupt official notes the CIN-RIN and CIN-candidate list correspondences on the voter's ballot form.

Randomisation attacks are also possible. With this, an attacker could require e.g. that the first candidate is marked, regardless of which candidate ordering is used. The level of threat is determined by the extent a voter can pick a ballot of her own choosing and the number of candidates in an election. In the case of few candidates, it might be easy for the voter to pick a ballot where she can vote as she wishes while satisfying the coercer. However, as Ben Adida points out in [5], a more complex randomisation attack is possible by forcing a voter to vote for a candidate on the ballot form that is determined by the ballot identifier (RIN). A randomisation attack may benefit the low key candidates as the votes would be spread evenly across the candidates.

- Authority Knowledge - Kleptographic attacks [11] are a possible threat, where e.g. a cryptographic operation on the RIN or CIN would give away information about the corresponding candidate list. Such an attack would obviously require a lot of searching, and would be dependent on how the RIN and CIN numbers are generated.
- Discarded receipts/surrendered receipts may indicate receipts that will not be checked and hence could be altered without detection. A possible countermeasure is to have a VVPAT-style mechanism in place.

Threats to integrity:

- Invalid signature - A possible threat if the mechanism for digitally signing receipts is malicious or fails. Likewise the mechanism for checking the signature on a receipt. The voter is then unable to prove an incorrectly recorded receipt.

- Fake receipt - A voter could falsely claim to be disenfranchised with a fake receipt. A suggested mitigation is to frank the receipts [21].

For both the above, a possible countermeasure is to digitally sign the receipts and then have immediate checks on the signatures.

3.1.6 WBB

Threats to confidentiality:

- Monitoring access to the WBB - This is not a particular threat as without knowledge of the RIN-candidate list correspondence, the value of the voter's vote cannot be learned from postings to the WBB. However, see threat under WBB modified.

Threats to integrity:

- The voter could be presented with a fake WBB, e.g. in a spoofing attack, and be misled into believing her vote has been recorded correctly when in reality, it has been changed.
- WBB modified - There is a risk that the WBB could be modified after the voter has checked her receipt. The WBB is supposed to be a write-only medium, but this seems hard to achieve in practice. Apart from the challenges of implementing a write-only WBB, a practical issue is how to handle detected errors. Voters can e.g. complain if they cannot find their receipt at the WBB or if the position of the voter's choice has been shifted. A write-only WBB would soon get quite disorganised if the old errors are kept, and additional columns with the corrected postings are added. Note however, that an attacker set on altering the election results in the "Scratch Card" voting system would not actually need to modify the WBB, since the scheme does not, in any case, provide verifiability of the final tally.

For both the attacks, a VVPAT mechanism is a possible countermeasure.

Threats to availability:

- DoS - A possible threat, e.g. due to network overload, power failure, etc.

3.1.7 Election results

Threats to integrity:

- Early publishing - A potential threat. To mitigate this, vote counting at local stations, the final tally and publishing of results should be carefully synchronised.

- Absence of verifiability - As the voter is only able to check that her receipt has been correctly recorded on the WBB, this should be regarded as a potential threat.
- False/erroneous count - There is a danger that this could go undetected, as the scheme offers limited verifiability. Again, a VVPAT mechanism is a possible countermeasure.

From the analysis above, it is clear that having a VVPAT mechanism in place would counter many of the threats to the integrity of the scheme. We next investigate whether or not this would add any further threats.

3.1.8 Paper audit trail

Threats to confidentiality:

- Voter identifiable from receipt - See the “verifiable receipt” component above for a discussion of this potential threat.
- Voter’s choice noted by official - This is not a threat assuming the scheme operates as intended. However, the above also applies.

Threats to integrity:

- Mismatch between voter’s choice and paper copy - Not a threat, as in the “Scratch Card” scheme, two copies could be made of the RHC: one of which the voter retains as a receipt, the other to act as a paper back-up.

It appears that a VVPAT mechanism would not introduce any additional threats, at least threats that may not have been present before. However, it may magnify any existing threats to confidentiality.

The analysis shows that the main problems with the scheme are firstly, that it only offers partial verifiability as the voter is only able to verify that her receipt has been correctly recorded by the system. Secondly, the voter may be open to coercion if the CIN-RIN correspondence on her ballot form, together with her ID is noted by an official. While a possible countermeasure for the latter has been suggested, the former requires trust in the correctness of the tabulation process.

The analysis also demonstrates that the model offers a systematic way to carry out a threat analysis of voting systems, i.e., by identifying the main phases and components in a scheme, and evaluating potential threats in direct relation to each component during a run of the protocol, taking into account its particular design aspects. We have aimed for generality so that the model is adaptable, and found this to be the case in the analysis of the “Scratch Card” scheme. The appropriate components could be readily selected from the model, and the vulnerabilities evaluated against the threat categories provided.

It should be noted that while every effort has been made to ensure completeness of our model, given the open-endedness of systems it is difficult to guarantee that it captures all possible threats.

In the next section we discuss the results and mention some limitations of the work.

4 Discussion and Future Work

We have presented a model for the systematic analysis of threats to voting systems that can be applied to a wide range of different schemes. This is further demonstrated in a forthcoming paper in which we use the model to analyse potential threats in Prêt à Voter [8] and the ThreeBallot voting scheme [22].

In anticipation of some of these threats, error detection mechanisms have been built into many current schemes, e.g. randomised partial checking [13] of the mix process to ensure correct decryption of votes without compromising voter privacy. While a systematic threat analysis is valuable for identifying the need for error detection mechanisms, it can also be useful for assessing the effectiveness of any that are existing within the scheme, especially when taking interactions between the various players and components into consideration.

A further step from an analysis such as the one performed above could involve assessing the likelihood of certain threats occurring. This goes beyond the model: not only assessing the potential threats but also the probability of their occurrence, and could involve a more complex and informed analysis of the scheme in relation to both the sociological and technical aspects of its environment. Estimating the security of a scheme would then require balancing the probability of the threats occurring against the effectiveness of any error detection mechanisms that may be in place. Bryans et al. discuss this issue in [6], and make a distinction between accidental and malicious error. Our model can be used for analysing potential threats through accident or malice, e.g. a user interface could be deliberately confusing, or confusing due to poor design.

Bryans et al. also mention the need to consider threats to the reputation of a voting system [6]. Interestingly, transparency in a voting scheme could work against it, e.g. a large number of reported errors in recorded votes could dissuade voters from using it, and cause it to be abandoned altogether. As previously mentioned, there has been a recent move towards increased transparency in voting systems as a way to provide verifiability and to reduce dependency on the “correctness” of the system. However, possible threats to the reputation of the system are worth careful consideration. Although we briefly touched on this issue in the previous section, our model does not directly analyse threats to reputation, as this lies outside the current (largely technical) scope. Once again, the analyst would need to merge the sociological and technical aspects of a scheme in assessing the strength of its reputation.

Another important point raised in [6] is the importance of error handling and recovery strategies, alongside error detection mechanisms. This is a currently neglected field in research on voting systems, and error handling and recovery is lacking in many current voting schemes. This is a highly complex issue, involving decisions not only on the way in which recovery should be effected,

but also when the appropriate mechanisms should be invoked. It is likely that decisions would have to be made as to when margins of error are regarded as insignificant, and when they become unacceptable. Patterns of error may have to be studied, e.g. in deciding whether a particular error is accidental or malicious. This may, in turn, affect decisions on how best to deal with the error.

It would be highly useful to have a systematic model not only for the threat analysis, but also for dealing with any errors or security breaches that may occur as a result of these threats. This could take the form of a model for the analysis of potential threats based on the components in a scheme, in conjunction with a series of “decision trees” offering possible ways to handle such threats should they occur. We envisage this as a possible extension of our model, and a subject of future work.

Acknowledgments

The authors would like to thank Jeff Yan for many helpful comments.

References

- [1] Workshop on developing an analysis of threats to voting systems, 2005. <http://vote.nist.gov/threats/index.html>.
- [2] The machinery of democracy: Protecting elections in an electronic world (full report), 2006. Brennan Centre for Justice, NYU School of Law, <http://www.brennancenter.org>
- [3] Punchscan, 2006. <http://www.punchscan.org>.
- [4] VoteHere, 2006. <http://www.votehere.net/default.php>.
- [5] B. Adida. *Advances in Cryptographic Voting Systems*. PhD thesis, Massachusetts Institute of Technology, 2006.
- [6] J. Bryans, B. Littlewood, P. Y. A. Ryan, and L. Strigini. E-voting: Design for dependability. In *Availability, Reliability and Security (ARES)*. IEEE, 2006.
- [7] D. Chaum. Secret-ballot receipts: True voter-verifiable elections. *IEEE Security and Privacy*, 2(1):38–47, January-February 2004.
- [8] D. Chaum, P. Y. A. Ryan, and S. Schneider. A practical, voter-verifiable election scheme. In *European Symposium on Research in Computer Security*, number 3679 in Lecture Notes in Computer Science. Springer-Verlag, 2005.
- [9] G. Danezis. *Better Anonymous Communications*. PhD thesis, University of Cambridge, 2004.

- [10] A. Fujioka, T. Okamoto, and K. Ohta. A practical secret voting scheme for large scale elections. In *Workshop on the Theory and Application of Cryptographic Techniques: Advances in Cryptology*, pages 244–251. ACM, 1992.
- [11] M. Gogolewski, M. Klonowski, P. Kubiak, M. Kutylowski, A. Lauks, and F. Zagorski. Kleptographic attacks on e-voting schemes. In *Workshop on Electronic Voting and E-Government in the UK*, 2006.
- [12] A. Gumbel. *Steal This Vote*. Thunder’s Mouth Press, U.S.A, 2005.
- [13] M. Jakobsson, A. Juels, and R. Rivest. Making mix nets robust for electronic voting by randomized partial checking. In *USENIX Security Symposium*, pages 339–353, 2002.
- [14] C. Karlof, N. Sastry, and D. Wagner. Cryptographic voting protocols: A systems perspective. In *USENIX Security Symposium*, 2005.
- [15] T. Kohno, A. Stubblefield, A. D. Rubin, and D. S. Wallach. Analysis of an electronic voting system. In *Symposium on Security and Privacy*. IEEE, 2004.
- [16] S. Kremer and M. Ryan. Analysis of an electronic voting protocol in the applied pi-calculus. In *European Symposium on Programming*, number 3444 in Lecture Notes in Computer Science, pages 186–200. Springer-Verlag, 2005.
- [17] R. Mercuri. A better ballot box? *IEEE Spectrum Online*, October 2002.
- [18] A. Neff. A verifiable secret shuffle and its application to e-voting. In *Conference on Computer and Communications Security*, pages 116–125. ACM, 2001.
- [19] A. Neff. Practical high certainty intent verification for encrypted votes, 2004. <http://www.votehere.net/documentation/vhti>.
- [20] T. Peacock. *Guess My Vote: a Study of Opacity and Information Flow in Voting Systems*. PhD thesis, School of Computing Science, Newcastle University, 2006.
- [21] B. Randell and P. Y. A. Ryan. Voting technologies and trust. *IEEE Security & Privacy*, October 2006.
- [22] R. L. Rivest. The ThreeBallot voting system. Unpublished draft, <http://theory.lcs.mit.edu/~rivest/Rivest-TheThreeBallotVotingSystem.pdf>, 2006.
- [23] A. Rubin. *Brave New Ballot: The Battle to Safeguard Democracy in the Age of Electronic Voting*. Morgan Road, 2006.

- [24] P. Y. A. Ryan. Verified encrypted paper audit trails. Technical Report CS-TR-966, University of Newcastle upon Tyne, 2006.
- [25] P. Y. A. Ryan and T. Peacock. Prêt à voter: a systems perspective. Technical Report CS-TR-929, University of Newcastle upon Tyne, 2005.