Contents lists available at ScienceDirect

Mathematical and Computer Modelling

journal homepage: www.elsevier.com/locate/mcm

An information-theoretic model of voting systems

Ben Hosp *, Poorvi L. Vora

Department of Computer Science, George Washington University, Washington DC 20052, United States

ARTICLE INFO

Article history:
Received 19 May 2008
Accepted 22 May 2008

Keywords:
Voting
Information theory
Privacy
Verifiability

ABSTRACT

This paper presents an information-theoretic model of a vote counting system, and well-defined criteria for evaluating such a system with respect to integrity, privacy and verifiability. The impossibility of achieving perfect integrity, perfect verifiability and perfect privacy follows easily from the information-theoretic approach. The model is applied to the measurement of privacy loss in the ThreeBallot and Farnel voting systems, and finds both systems to have similar privacy loss.

1. Introduction

Elections in several democracies (Brazil, India, the United States, the United Kingdom) have begun to rely upon computerized or electronic voting technology. Yet, the literature does not provide a standard model to compare the electronic vote counting systems with the electromechanical and paper-based systems they have replaced, or to compare them among themselves. This paper presents a voting model that is based on information flow through a vote counting system, and can be used to compare both: simple systems, as well as newer systems that combine the benefits of paper ballots with those of electronic voting. In the remainder of this paper, for brevity, a vote counting system will be referred to as a *voting system*.

It has been our goal to make the model as inclusive as possible, while retaining the essence of a voting system. Votes enter the system, they are processed by the system, which then provides the tally. Without loss of generality, the system is assumed to provide a specification of the algorithm it uses to determine the tally, we denote this algorithm as *SpecifiedVoteCount*. Additionally, the system provides information that leads the voters, and election observers, to trust that *SpecifiedVoteCount* was indeed followed. This information is referred to as a "tally-correctness proof" in the literature on cryptographic voting systems, though the information provided typically does not conclusively prove *SpecifiedVoteCount* was followed, but simply decreases the uncertainty that it was followed. Additionally, other extraneous information may also be provided by the system; for example, logs on the order in which votes were received, who voted, etc. The system may be seen as employing another algorithm to provide its output, this algorithm is denoted *ElectionOutput*.

This paper provides formal definitions of three desirable properties of voting systems, and information-theoretic measures of how well the system achieves these properties. Informally, a system is said to provide (a) integrity when *SpecifiedVoteCount* provides the correct vote tally, (b) privacy when *SpecifiedVoteCount* and *ElectionOutput* together reveal no more information about individual votes than is revealed by the vote tally, and (c) verifiability when *ElectionOutput* provides enough information to determine that the system did indeed follow *SpecifiedVoteCount*. The integrity is measured by the reduction in uncertainty of the tally if *SpecifiedVoteCount* is followed, the privacy loss is measured by the reduction in uncertainty of the individual votes in addition to the reduction due to knowledge of the tally, and the verifiability is measured by the reduction in uncertainty that *SpecifiedVoteCount* was indeed followed.

* Corresponding author.
E-mail addresses: bhosp@gwu.edu (B. Hosp), poorvi@gwu.edu (P.L. Vora).

0895-7177/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.mcm.2008.05.040

It is worth noting that there is a natural tension among integrity, verifiability and privacy; if the system were to reveal all the vote values, each voter were to publicly verify his or her vote, and all votes were tallied, we would have perfect integrity, perfect verifiability and zero privacy. On the other hand, a system could take in the votes, provide no proof, and provide a vote tally. Direct-recording electronic (DRE) voting machines commonly used in elections today are much like this; such a system would provide no verifiability and perfect privacy. Audits of voting systems before voting day, or on voting day but using dummy votes, serve to decrease uncertainty on whether *SpecifiedVoteCount* was indeed followed, and do not reveal information on the actual votes. However such audits can never completely reduce the uncertainty, as there is always some uncertainty regarding what was done with the actual votes.

Cryptographic voting systems such as Prêt à Voter [1,2] and Punchscan [3] attempt to address the tension between privacy and verifiability by processing encrypted values of the votes and providing information on the intermediate steps in the processing. These systems typically depend on the adversary—who is trying to change the tally or find out information on individual votes—being *computationally bounded*; such an adversary does not have the resources to break the encryption in a short enough time period. This is not an unreasonable assumption—the security of electronic commerce, for example, depends on the same assumption. However, because it is also interesting to understand the essential limits of the voting problem and other similar problems, one may also ask what a *computationally unbounded* adversary—who is able to efficiently break modern encryption schemes—is able to achieve.

The contributions of this paper are threefold:

- Some of the more important desirable properties of voting systems—integrity, privacy and verifiability—are carefully defined, and information-theoretic metrics for the measurement of deviation from perfect are presented.
- A simple model of the tally-correctness proof—which applies to several, if not all, current voting systems—is presented. The paper proves that, when voting systems follow this model, perfect integrity, perfect ballot secrecy and perfect tally verifiability cannot be simultaneously achieved when the adversary is computationally unbounded.
- The model is applied to measure the privacy loss in the ThreeBallot [4] and Farnel [5] voting systems of Rivest et al. and Custódio et al. respectively, and finds the systems to be similar with respect to measured privacy loss.

The advantage of this model is (a) it provides a single framework in which to define and measure integrity, privacy and verifiability, and (b) the tradeoffs among these criteria are explicit in this model. It is anticipated that the model and the corresponding measures will be useful in, among others, the following specific instances:

- To study and compare the properties and limitations, not of specific voting system designs, but of entire categories of voting systems, as illustrated in Section 6. A category of voting system may be viewed as a constraint on a certain aspect of the model, such as the proof.
- In the interests of computational efficiency, most proposed cryptographic voting systems (see, for example, Punchscan [3] and Prêt à Voter [1,2]) reveal some information about the individual votes even to computationally bounded adversaries [6,7]. This information may be quantified using the measures proposed in this paper. For example, the randomized partial audits of Jakobsson–Juels–Rivest [8] reveal partial information about votes.
- As illustrated in this paper, the model may be used to evaluate voting systems (such as ThreeBallot and Farnel) that do not use cryptography and expose individual vote information.
- The model can also be used to evaluate election procedures, such as the manner in which voter lists are made public, and the kind of information about voters that is released. For example, the privacy measure of [9] has been used to measure the privacy loss of Californian voters due to precinct-level tallies [10].

This paper is organized as follows. Section 2 describes some modern voting systems to which the model may be applied. Section 3 describes related work, and Section 4 contains an informal list of goals of an election system. This list motivates the more formal definitions provided in Section 5, which presents the mathematical model. Section 6 contains a simple proof that a system in the model cannot provide perfect integrity, perfect privacy and perfect verifiability when the adversary is computationally unbounded. Section 7 contains applications of the model. First, this section provides comparisons of the privacy loss of the ThreeBallot and Farnel voting system. Second, it studies the privacy loss of the Farnel voting system as a function of its parameters. Conclusions and future directions are presented in Section 8.

2. Some example voting systems

This section briefly describes some recently-proposed voting systems which can be examined using the model. One of these systems (Prêt à Voter [1]) uses cryptography, and two (ThreeBallot [4] and Farnel [5]) do not.

2.1. Prêt à Voter

The Prêt à Voter [1,2] ballot consists of two halves printed on a single sheet of paper, side-by-side, separated by a perforation. The left half contains the names of the candidates, in a pseudo-random order. The right half contains spaces for marks against each candidate, as well as an “onion”, whose function will be described later (see Fig. 1). The voter marks the space next to the candidate of her choice, tears the ballot along the perforation, and destroys the left half. The right half is scanned into the polling machine, and the information in it displayed in the virtual ballot box (the exact image is not

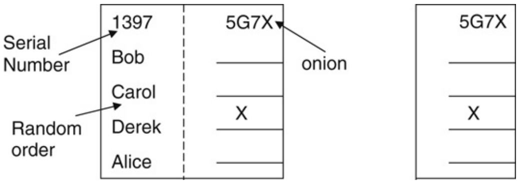

Fig. 1. Left: A Prêt à Voter Ballot; Right: A Prêt à Voter Receipt.

shown in the virtual ballot box, to prevent the leakage of information through messages written on the ballot by the voter, for example). The right half is also the receipt. As the candidate order is unknown, the receipt reveals no information about the vote, and may be considered an encrypted vote.

The information required to obtain the vote from the encrypted receipt is contained, in encrypted form, in the onion. That is, the onion contains information on the candidate ordering contained in the (destroyed) left half of the ballot. The entire virtual ballot box, containing all the receipts, is passed serially through a set of trusted parties. Each trusted party decrypts part of the onion on each vote, uses the information to perform an operation on the corresponding vote, leaves the rest of the onion with the processed vote, shuffles all the votes, and passes them on. The composition of all the operations performed by the trusted parties results in the decryption of the ballots. The decrypted ballots may be tallied by anyone. The input and output to each trusted party is made available on a public bulletin board.

The system also provides proof—to a group formed of the election authority, the candidates, and election observers—that it followed the encryption and decryption processes described above. It decrypts the onions on half of the printed ballots, chosen at random before the election, and after the printing, by the group. It thus demonstrates that the onions hold correct information about the candidate ordering, and, if correctly used by the trusted parties, will result in the correct decryption of votes. These ballots are treated as spoiled and are not used for the election. The decryption process performed by the trusted parties is audited by the group by having each trusted party demonstrate that it correctly performed operations on half of the ballots, chosen randomly after all the processing is performed. The ballots are chosen so as to reveal minimal information about the link between an encrypted ballot in the virtual box and its final decrypted version. The audit information revealed by each trusted party is made public so anyone can check it.

If at least some voters check that their receipts are in the virtual ballot box, a cheating trusted party or election system is caught with high probability. If the receipts are unforgeable, it is not possible to disrupt an election by falsely claiming that the election system is cheating.

2.2. The Farnel voting system

In the basic Farnel system [5], a "basket" containing equal numbers of each possible ballot is set up. After each voter places his voted ballot in the basket, the basket signs it, shuffles its contents, and gives her back a random signed ballot. The voter places this ballot in the ballot box and takes a copy of it as her receipt. At the end of the polling period, the basket is opened and its contents are placed in the virtual ballot box. Voters check that the receipts they possess, all indexed by serial number, are in the virtual ballot box. The votes in the (physical and virtual) ballot box are then counted, yielding the vote tally plus the "dummy" ballots which were in the basket originally.

It is trusted that the basket performs an (unverified) perfect shuffle. If the basket were not assumed to be honest, it could, for example, delete votes and replace them with others, whose receipts would be provided to the voter, who cannot check that the receipt given her does indeed represent the hitherto-unreleased vote of a previous voter.

In another version, the box contains all votes cast at any given time, and provides only a copy of a randomly-chosen previously-cast vote to the voter.

2.3. The threeballot voting system

The ThreeBallot ballot consists of the list of candidates, arranged one below the other in a fixed, pre-defined order, and three columns next to the candidates. To choose a candidate, the voter marks two of the three columns corresponding to the candidate. For all other candidates, the voter marks exactly one column. The three columns are separated out and cast separately, each as a ballot. Each ballot has an associated serial number, though the serial numbers are independent. The voter scans in all three and takes exactly one home with her; this is her receipt. A trusted checker checks that the scanned ballots are not overvotes or undervotes, that is, that the voter has not marked two columns for more than one candidate, and that she has not marked more than two columns, or fewer than one column, for any candidate.

All ballots are posted online with the corresponding serial numbers. Each voter checks if her receipt is on the bulletin board. She does not know if the other two ballots are there too (and unchanged), but because the voting system cannot guess which of the three she took home with her, it will be caught with high probability if it changes even a few ballots. Anyone can tally the votes—the winner will be the candidate with the most marks. The number of votes obtained by each candidate is the number of marks less the total number of voters.

The assumptions of this scheme are that the voting system cannot anticipate which ballot will be kept by the voter, and that the scanner is trusted not to allow overvotes or undervotes, and does not need to be audited. The system is resistant to direct coercion, since a voter can produce any desired mark pattern on his receipt, regardless of his actual vote. (Note that in general we could have B (rather than 3) ballot-columns, in which case chosen candidate-rows would have $B-1$ marks and unchosen candidate-rows would have $B-2$ marks.)

3. Related work

Perhaps the earliest list of voting system requirements is due to Shamos, [11]. Many papers in the WOTE 2001 [12] and WEST 2002 [13] workshops also include overviews of voting system requirements [14–16]. None provide a means of measuring performance with respect to the requirements. Papers on evaluating voting technologies include [17], and several other papers from the NIST Workshop on Threats to Voting Systems [18], in particular [19,20], provide an evaluation with respect to threats to count integrity. A mathematical definition of voting system privacy, and a related entropy-based privacy measure, which our work draws heavily from, is due to Coney et al. [9]. Information flow and entropy are commonly used in the security literature and data mining literature—see, for example, entropy-based measures of anonymity of Diaz et al. [21] and Serjantov and Danezis [22], and entropy-based privacy measures of Agrawal and Aggarwal, [23], and Vora [24] that are very similar to our measure of average privacy loss. Our worst-case privacy loss measure is based on the measure of Evfimievski et al. [25]. Finally, this paper builds on an initial presentation of the model at WOTE '06 [26].

4. Election goals

This section provides a brief list of informally-stated desirable properties of voting systems; the goals have been drawn from prior work such as [11,14–16,27]. This list provides the motivation for the more formally stated desirable properties—our model focuses only on integrity, privacy and verifiability.

1. **Usability:** Ballots should be "cast as intended", meaning that an otherwise valid voter who intends to cast a vote for choice X should not be thwarted by election procedures or technology.
2. **Integrity:** Ballots should be "counted as cast", meaning that the voting system should declare that a particular choice received m votes if and only if exactly m ballots marked for that choice were cast. In proportional representation elections the fraction of votes obtained by each choice is important; in first-past-the-post elections, the actual votes are not as important as long as the margin is sufficient.
3. **Privacy:** The secret ballot principle should apply to the election; a particular voter should not have the contents of her ballot associated with her in any way by anyone—even with the collusion of many parties, including election officials and other voters—beyond the association implied by knowledge of the tally. Notably, this privacy should be *involuntary*—that is, even a set of colluding parties that includes the voter herself should not be able to prove the contents of her ballot, beyond that implied by the tally, once she has left the polling place. Note that knowledge of the tally does allow a set of $n-1$ colluding voters, out of a total of n voters, to determine the vote of the non-colluding voter.
4. **Verifiability:** Both the general public—including non-voting observers—as well as the individual voter should be able to rest assured that the above goals have been met. Such assurance should not require real-time observation of election procedures or secret information.
   Dispute-freeness is a special kind of verifiability proposed by Kiayias and Yung [27], where disputes raised by various parties as to the validity of the election are decidable based on information that is publicly-available. In other words, the dispute resolution procedure is universally verifiable.
5. **Robustness:** Errors and failures should be detectable and fixable without impact upon the other goals.

Our model provides a set of definitions and a system of metrics which measure how well a given voting system accomplishes goals 2, 3 and 4. An election is modeled as a communication channel, and the metrics are designed to measure the information flow through such a channel. The metrics are based on the concept of information-theoretic entropy, but are anticipated to work equally well given computational assumptions. In this model, the concepts of Integrity, Privacy, and Verifiability may be thought of as follows:

- **Integrity** refers to the capability of the voting system design to communicate information about the actual or accurate vote totals.
- **Privacy** refers to the resistance of the voting system to providing information about individual votes.
- **Verifiability** refers to the extent to which the voting system is able to demonstrate that it used its full capability (which was measured by Integrity).

5. The model

In this section we describe our model and translate some of the goals of the previous section into mathematical conditions in the model.

We will consider that there are n voters casting ballots in an election. Let V_i be a discrete random variable representing the ith voter's ballot (or rather the votes it contains); $V_i \in \mathcal{V}$, the set of all possible ballots in the election. We will use V^* as shorthand for the ordered list of ballots, $[V_1, V_2, \ldots V_{n-1}, V_n]$; $V^* \in \mathcal{V}^n$, the set of all n-tuples with elements from $\mathcal{V}$. Let V^Σ be the vote count, in the form: "In the race for Governor, 600 votes for Alice, 400 for Bob; on Proposition 242, 580 votes for Yes, 420 votes for No, ..." V^Σ is therefore also a discrete random variable, but it is a deterministic function of V^*.

Let $\widehat{V^\Sigma}$ represent the vote count output by the voting system used to conduct the election, and let E, an ordered tuple of random variables, be its entire "output". E contains, in addition to $\widehat{V^\Sigma}$, also other information: for example, system logs, proofs to convince the observer that the voting system behaved as claimed, the number of voters who voted, the order in which they arrived, etc. Assume that the voting system declares two algorithms, i.e. two sets of well-defined steps, *SpecifiedVoteCount* and *ElectionOutput*, that, when applied to the votes, V^*, produce $\widehat{V^\Sigma}$ and E respectively. For example, *SpecifiedVoteCount* might consist of (a) the encryption of votes, followed by (b) partial decryption and shuffling by each of a set of entities, known as *mixes* [28] in the literature on cryptography, and finally (c) a tally of the final decrypted votes. In this case, *ElectionOutput* would be the algorithm that provides the input and output to each mix.

We do not assume that either algorithm *SpecifiedVoteCount* or *ElectionOutput* is deterministic, and we denote the random variables influencing the output of these algorithms by $X_1, X_2, \ldots, X_m$, where the numbering is not related to voter order. These random variables may represent random numbers used to seed pseudo-random number generators, for example, or voting system inputs provided by voters that are not related to their votes. X^* is the vector of all random variables, $X^* = [X_1, X_2, \ldots, X_{m-1}, X_m]$ and is independent of V^*. The random variables are also themselves independent of one another, i.e. $\mathcal{I}(X_i; X_j) = 0$, $i \neq j$, and $\mathcal{I}(V^*; X^*) = 0$. Thus *SpecifiedVoteCount* (V^*, X^*) is denoted $\widehat{V^\Sigma}$ and *ElectionOutput* (V^*, X^*) is denoted E.

5.1. Preliminaries

We use the notion of entropy to define the mathematical goals of a perfect voting system and to measure deviation from perfect. As defined by Shannon [29], entropy is a mathematical measure of the uncertainty in a random variable. We concern ourselves only with discrete random variables, and measure entropy in bits. The entropy of discrete random variable X that takes on value x with probability $p_X(x)$ is [29]

$$\mathcal{H}(X) = -\sum_x p_X(x) log_2 p_X(x)$$

with the understanding that $p_X(x) log_2 p_X(x) = 0$ when $p_X(x) = 0$. Roughly speaking, the entropy of a random variable is understood to be the average number of bits required to represent it.

The computational entropy of a random variable, roughly speaking, is the average number of bits required to represent it under the constraint that the algorithm generating the bits from the random variable is feasible in the computational model [30]. In certain instances, when secrecy is provided by computational assumptions, it is more appropriate to use computational entropy over "Shannon" (or absolute) entropy. The paper addresses, for ease of exposition, only information-theoretic entropy, and points out instances when computational entropy is relevant; the use of computational entropy appears to be a straightforward extension of this work. The fact that we do not address computational entropy explicitly should not be taken to imply that we require or recommend the use of only Shannon entropy in all cases.

When two random variables X and Y are not independent, knowing the value of one reduces the uncertainty of the other. Let $p_{X|Y}(x; y)$ be the probability that random variable X takes on value x when random variable Y has fixed value y. If $\mathcal{H}(X|Y)$ is the uncertainty in X if Y is known,

$$\mathcal{H}(X|Y) = \sum_y p_Y(y) \left[-\sum_x p_{X|Y}(x; y) log_2 (p_{X|Y}(x; y)) \right] = \sum_y p_Y(y) Z(y).$$

It is the average value of the random variable $Z(Y) = -\sum_x p_{X|Y}(x; Y) log_2 p_{X|Y}(x; Y)$. $Z(Y)$ is the uncertainty in the random variable X given a particular value of random variable Y, and is hence commonly denoted, through some abuse of notation, by $H(X|Y = y)$.

The reduction in entropy in one variable, due to the other being known, is termed mutual information. It is the average value of the random variable $\mathcal{H}(X) - H(X|Y = y)$, and is defined as follows:

$$\mathcal{I}(X; Y) = \mathcal{H}(X) - \mathcal{H}(X|Y) = \sum_y p_Y(y)[\mathcal{H}(X) - H(X|Y = y)]. \tag{1}$$

It can be shown that:

$$\mathcal{I}(X; Y) = \mathcal{I}(Y; X).$$

5.2. Integrity

Election integrity requires that algorithm *SpecifiedVoteCount*, if followed, produce the *correct* vote count, V^Σ. In this model, integrity does not address the issue of whether *SpecifiedVoteCount* is indeed followed by the voting system; this is covered by the property of verifiability (see Section 5.4). This distinction between election integrity and verifiability allows us, for example, to observe that hand counting of votes provides lower integrity than does good electronic counting. Even if the hand count algorithm is followed honestly, it will typically not provide an exact count (see Example 1).

Election integrity may be defined more precisely as follows:

Definition 1. A voting system provides **perfect integrity** if $\widehat{V^{\Sigma}} = V^{\Sigma}\ \forall\, p_{V^*},\ X^*$.

Even if the system does not provide perfect integrity, the uncertainty in V^{Σ} is generally reduced on knowledge of *SpecifiedVoteCount*(V^*). The reduction in uncertainty, $\mathcal{I}(V^{\Sigma}; \widehat{V^{\Sigma}})$ could range anywhere from zero (indicating election results independent of the cast ballots, and hence an election with zero integrity) up to a maximum of $\mathcal{H}(V^{\Sigma})$ (indicating perfect integrity). One could use a normalized value of this reduction in uncertainty to measure election integrity.

Definition 2. The **integrity measure** of an voting system is

$$\mathfrak{I} = \frac{\mathcal{I}(V^{\Sigma}; \widehat{V^{\Sigma}})}{\mathcal{H}(V^{\Sigma})}.$$

We assume that the system treats all voters similarly (if not, then there is more than one voting system). Our uniform approach while determining measures has been that, when a system treats a particular user or groups of users differently from others, we are actually dealing with multiple systems, and each has its own measure. Thus, while the entropy-based measure averages over outputs $\widehat{V^{\Sigma}}$, it does not average over populations that are treated distinctly by the system. This is so as to draw attention to inequities among the voting systems. For example, in the US, one county might use hand counts, another optical scan systems, and yet another an electronic voting system. The tallies from several counties would be used to determine the statewide winner of a particular race, yet it would not be appropriate to provide a single measure that attempts to evaluate the "combined" voting system used. It would be appropriate to point out that the system of one county provides, say, less privacy than does that of another county.

Notice that $\mathfrak{I}$ is a function of the probability distribution of V^*, $p_{V^*}(v^*)$. This and subsequent measures could be made independent of any particular distribution by taking the worst case and maximizing (or minimizing) over all possible distributions, as with the privacy measure of [9]. We have chosen not to do this for two reasons. First, it is our opinion that voting systems should be evaluated for the situations where they are expected to be used. Different jurisdictions may have dramatically different electorates, legal provisions, etc., which make a particular system more or less appropriate for use, and the yardstick used to measure voting systems should reflect this. Second, the worst case may be a boundary case (such as a unanimous election) unlikely to occur in the real world. Basing a model on such cases is likely to yield uninteresting, trivial results. For example, the worst-case behavior of the Farnel voting system, examined in Section 7.3 is seen to arise when the entropy in the vote distribution is very small.

Example 1. Consider a voting system that produces a vote count through hand counting, where $\widehat{V^{\Sigma}}$ is the average of N hand counts. The uncertainty in $\widehat{V^{\Sigma}}$ is the uncertainty due to hand counting. The integrity is not perfect, and the integrity measure increases with N. Note that the integrity is imperfect *even when the hand count is performed honestly and as specified by SpecifiedVoteCount*. Whether the system actually does count the votes correctly and honestly is addressed through the property of verifiability, see Section 5.4.

We assume, wlog, that an integrity value of one implies perfect integrity – that is, that $\widehat{V^{\Sigma}} \neq f(V^{\Sigma})$ for f a deterministic invertible function that is not the identity. If $\widehat{V^{\Sigma}} = f(V^{\Sigma})$, $\widehat{V^{\Sigma}} \neq V^{\Sigma}$, but the integrity would be one. However, the value of $\widehat{V^{\Sigma}}$ would contain all the information necessary to obtain V^{Σ}, which can be obtained by applying the inverse of f to $\widehat{V^{\Sigma}}$. Note that this is true whether the measure used is "Shannon entropy" or computational entropy, because a computational integrity of one means that V^{Σ} can be determined from $\widehat{V^{\Sigma}}$ in the computational model.

5.3. Privacy

Election privacy requires that the voting system not reveal additional information about the values of individual or specific votes beyond that revealed by the tally. In this section, we define perfect privacy, and a metric to measure deviation from perfect. We also provide *maximum information loss* and *specific-case privacy loss* measures.

5.3.1. The definition of privacy

Perfect privacy is defined by Coney et al. [9] as being provided when no information at all is revealed about individual votes. This is trivially not possible when an accurate vote tally is revealed. Hence we define a system as providing perfect privacy when it provides no more information about V^* other than V^{Σ}. In the definition, we use the fact that $E = $ *ElectionOutput*(V^*, X^*); that is, the privacy definition assumes that the (possibly probabilistic) relationship between the output of the voting system and the individual votes is known. Note that *ElectionOutput* is not restricted to *SpecifiedVoteCount* (V^*, X^*), it includes any other information the system may reveal; such as encrypted receipts and proofs of correctness of mixnets.

Definition 3. A voting system is said to provide **perfect privacy** if V^* is conditionally independent of E after conditioning on V^Σ, i.e.,

$$p_{V^*|V^\Sigma}(v^*; v^\Sigma) = p_{V^*|V^\Sigma,E}(v^*; v^\Sigma; e)$$

for all v^*, v^Σ, e.

Deviation from perfect privacy can be measured using a measure similar to a normalized form of that in [9], which in turn derives from the conditional privacy loss measure of [23].

Definition 4. The **privacy loss**, $\mathfrak{L}$, of a voting system and process is

$$\mathfrak{L} = \frac{\mathcal{I}(V^*|V^\Sigma; E)}{\mathcal{H}(V^*|V^\Sigma)}.$$

Notice that, like the integrity measure $\mathfrak{I}$, $\mathfrak{L}$ is a function of the probability distribution on the vote, p_{V^*}. It ranges between zero (indicating that the voting system reveals nothing at all other than the vote tally) and one (indicating that the voting system reveals all ballots exactly) for any non-trivial p_{V^*}. Privacy loss may be distinguished as voluntary and involuntary; in the former, the adversary obtains no information from the voter, and in the latter, the voter is assumed to collude with the adversary [9].

Example 2. Consider a precinct with a single polling machine that provides Voter-Verified Paper Audit Trail (VVPAT) records on a paper reel which maintains the order of the vote. Suppose further that election officials maintain a record of who voted, in the order of arrival. Trivially, E consists of the ordered list of votes. Hence $\mathcal{H}(V^*|E) = 0$, and the privacy loss of this system is one.

The measure $\mathfrak{L}$ is used to compare the privacy properties of the Farnel and ThreeBallot voting systems in Section 7.

5.3.2. A privacy measure based on maximum information loss

In the field of data mining, it has been noted that a measure such as $\mathfrak{L}$, based on mutual information (a measure of average entropy loss, over the conditioning variable, for a specific distribution), has certain limitations. Most importantly, it does not identify cases when the entropy of V^* decreases significantly for a particular value of $\mathit{ElectionOutput}(V^*, X^*) = e$, the conditioning variable, which occurs with small probability [25]. As the average and maximum values of the entropy reduction both play a role in characterizing it, and because both have played a significant role as privacy measures in the field of data mining, we define the privacy measure based on maximum information loss. Note that the maximum is not taken over several probability distributions p_{V^*}, but over several values of e for a fixed distribution p_{V^*}. This is for two reasons, as mentioned earlier. First, we want the measures to be a function of the distribution p_{V^*}, which represents the specifics of the election for which the voting system will be used. Second, maximizing over several probability distributions is likely to result in a marginal case, such as a unanimous election, unlikely to occur in the real world. On the other hand, maximizing over several values of e, could result in a case that is rare, but can still occur for a specific election, represented by p_{V^*}.

Definition 5. The **maximum information loss privacy measure**, $\mathfrak{L}_m$, of a voting system and process is

$$\mathfrak{L}_m = \frac{\max_e[\mathcal{H}(V^*|V^\Sigma) - H(V^*|V^\Sigma; E = e)]}{\mathcal{H}(V^*|V^\Sigma)}.$$

Note that, like $\mathfrak{L}$, $\mathfrak{L}_m$ is a function of the probability distribution $p_{V^*}(v^*)$.

5.3.3. Specific-case privacy loss

In the previous sections, we defined privacy loss taking into consideration the information contained in E. In this section, we focus on the situation where the adversary uses a subset of the information in E to obtain information on a subset of the votes. This approach allows us to bound the adversary in simple ways: perhaps the adversary is not sophisticated enough to obtain all the output of the voting system and use it to obtain the relationships among all the votes. For example, an adversary waiting outside the polling booth might use only the voter's receipt in Farnel, to estimate only her vote (this example is examined in more detail in Section 7.3). We believe this will be a somewhat more useful measure than $\mathfrak{L}$ or $\mathfrak{L}_m$ for the examination of real voting systems, while $\mathfrak{L}$ or $\mathfrak{L}_m$ will be more useful in stating properties of voting systems and impossibility results.

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040

Definition 6. The **amount of specific-case privacy loss** on vote subset V^S due to information A, $\mathfrak{L}_{V^S,A}$, of a voting system and process is

$$\mathfrak{L}_{V^S,A} = \frac{\mathcal{I}(V^S|V^\Sigma; A)}{\mathcal{H}(V^S|V^\Sigma)}. \tag{2}$$

An example of the use of this measure is presented in Section 7.3.

5.4. Verifiability

This section addresses verification: how one may determine that the voting system is actually following algorithm *SpecifiedVoteCount*.

Consider the tuple E as consisting of:

- *Tally*, a purported vote count output by the system.
- P: the proof—consisting of claims used to prove that algorithm *SpecifiedVoteCount* was followed, and truth values of verified claims.
- all other information released by the system, perhaps some of it unintentional.

5.4.1. The proof

In this section we provide more detail on the model of a proof.

Definition 7. A **voting system proof** is a collection of *voter-verifiable claims*, *voting system claims*, *inferences*, and a set of variables denoting the truth of claims and inferences that have been verified. A single voter-verifiable claim is verifiable by agreement between the voting system and a single voter, the other claims and inferences are verifiable by the public.

We first define voter-verifiable claims. These are claims about the combination of a vote with some random variables from X^*, or claims about some random variables by themselves. In either case, it is required that random variables used to make claims about V_i not be reused to make a claim about V_j, for $i \neq j$. This is so as to prevent voter i to determine information about V_j through knowledge of the claim about V_j and the reused random variable.

Definition 8. If f and f' denote any deterministic functions, and $\mathbf{X}_i$ and $\mathbf{X}'_i$ denote subsets of X^*, an ith-voter-verifiable claim is either:

- a *vote-dependent voter-verifiable claim*, of the form $f(V_i, \mathbf{X}_i) = y$, or
- a *vote-independent voter-verifiable claim*, of the form $f'(\mathbf{X}'_i) = y'$.

For some y, y', such that $(\mathbf{X}_i \cup \mathbf{X}'_i) \cap (\mathbf{X}_j \cup \mathbf{X}'_j) = \emptyset$ for $i \neq j$. Each ith-voter-verifiable claim is verifiable by agreement between a single, specific voter who voted after voter i, and the voting system. The number of ith-voter-verifiable claims is c, a constant independent of i. Voter i cannot verify more than c claims.

Example 3. Consider the Prêt à Voter system. Let $\sigma(k, :)$ be the permutation printed on ballot k. Let $X_{(k,2)}$, be the string on the top right hand corner. A vote-dependent voter-verifiable claim is of the form $\sigma(k, V) = y_{(k,1)}$ where $y_{(k,1)}$ is a specific value. For the example receipt of Fig. 1 in Section 2, $\sigma(k, :)$ is the permutation that takes the alphabetically-ordered candidate list *Alice, Bob, Carol, Derek* to *Bob, Carol, Derek, Alice* or the permutation taking (1, 2, 3, 4) to (4, 1, 2, 3). V is the vote for *Derek*. $k = 1397$, and $y_{(k,1)} = 3$, the position of the mark on the receipt, and the value of $\sigma(1397, \mathit{Derek})$. A vote-independent voter-verifiable claim is of the form $X_{(k,2)} = y_{(k,2)}$, which states that the string on the top right-hand corner takes on the value $y_{(k,2)}$. By verifying her receipt, the voter is verifying that the permutation printed on the ballot encodes her vote to position 3, and that the string 5$G$7$X$ is the string on the top right-hand corner.

We make the following clarification regarding the voter-verifiable claims:

- If the voter and voting system do not agree on the truth of the claim, an appropriate dispute resolution process is assumed to follow, that is assumed to result in a correct resolution of the truth value of the claim. In cryptographic systems, the dispute resolution process typically uses digital signatures, in paper-based systems it could be an authentication seal or stamp.
- As maybe noted by Example 3, in voting systems that provide receipts, as long as a voter has a receipt and agrees that it represents her experience in the polling booth, and as long as the receipt is authenticated by the voting system, it represents an agreement between the voter and the system, and hence represents a voter-verifiable claim. A third party can now trivially verify that this receipt is in the virtual ballot box, that is, a third party can now trivially communicate that this claim has been verified by agreement between the voter and the voting system. However, it is only the voter and the voting system that can agree that the receipt is a correct representation of the interaction in the polling booth.

- The condition that each claim be a function of at most one vote implies that the model does not allow, for example, two voters to agree on the sum of their votes and to verify the encryption of that sum. The reason is as follows. The problem of determining the verifiability of a claim about two or more votes is similar to that of determining the verifiability of the voting system. The truth of this claim will either follow from the truth of other claims about single votes, and will hence be an inference, or will require an agreement between the voters whose votes are involved, thus revealing information on the other votes to a single voter.
- We do not state explicitly that a claim about vote V_i is only verifiable by voter i. If a Farnel basket exists, voter i could verify a claim about some V such that V was cast before V_i. We implicitly assume that all voters do not cast their votes simultaneously, because the existence of a Farnel basket and simultaneous voting is simply the existence of a trusted shuffle, and a trusted shuffle is equivalent, when the tally count can be performed without error, to a trusted voting system.

The truth value of each voter-verifiable claim has an associated entropy, based on the joint entropy of variables $V_i, X_1, X_2, \ldots, X_s$. In the information-theoretic setting, the entropy values will be independent of whether f is one-way or not, as the adversary will be assumed powerful enough to obtain any information leaked by the function f. If a voter-verifiable claim is verified to be true by the voter and the voting system, the entropy of the truth value of the claim is now reduced to zero. If it is verified to be false by the voter and the voting system, this too reduces the entropy of the truth value to zero. If a voter is not involved in the verification of a voter-verifiable claim, its entropy is unchanged. The truth value of all verified claims is part of the proof.

Voting system claims are claims about the random variables used in the computation of voter-verifiable claims and the processing of votes. The random variables may be functions of other random variables, however claims cannot be made by combining random variables used to process one vote with those used to process another. This is to prevent a voter from determining relationships between her vote and another vote through relationships between the random variables.

Definition 9. Let f denote any deterministic function, and $\mathbf{X}$ a subset of X^* such that, $(\mathbf{X}_i \cup \mathbf{X}_i') \cap \mathbf{X} \neq \emptyset \Rightarrow (\mathbf{X}_j \cup \mathbf{X}_j') \cap \mathbf{X} = \emptyset$, for $i \neq j$. A **voting system claim** is of the form $f(\mathbf{X}) = y$. It is verifiable by anyone if the voting system provides enough information on $\mathbf{X}$.

Example 4. For Prêt à Voter, an example voting system claim is $P_{k_3}(g_3, P_{k_2}(g_2, P_{k_1}(g_1, P_{k_0}(g_0, D_0)))) - X_{(k,2)} = 0$ (see Example 3) for random numbers g_0, g_1, g_2, g_3 and D_0, and public keys of the mixes $P_{k_0}, P_{k_1}, P_{k_2}, P_{k_3}$. That is, the voting system claims that the string on the upper right-hand corner, $X_{(k,2)}$, is the onion, formed by nested encryption of some random numbers using the public keys of the mixes the votes will go through. Further, another voting system claim is that $\sigma(k, :)$ is the cyclic shift by the value $\sum_{i=0}^{3} hash(g_i) \ mod \ C$ where C is the number of candidates. Both claims are publicly verifiable by examining ballot k, and spoiling it in the process. The first claim is verifiable by requiring the system to provide the random variables g_0, g_1, g_2, g_3 and D_0, and checking that the encryption does indeed provide the string $y_{(k,2)}$, verified by the voter as being the string in the upper right hand corner, the position of the onion. The second claim is verifiable by checking that the value of $\sum_{i=0}^{3} hash(g_i) \ mod \ C$ does indeed correspond to the cyclic shift of the permutation printed on the ballot.

A *cryptographic commitment* is a value that the committer commits to, and cannot change at a later time. It is typically kept secret for a time interval. The commitment is said to be *opened* when the value is revealed. For example, the onions are commitments made by the Prêt à Voter voting system to the random numbers used. The random numbers cannot be changed once committed to, but are kept secret till opened, during the spoiling of the ballots. Thus cryptographic commitments are a special type of voting system claim, and the verification of the claim corresponds to the opening of the commitment.

When a commitment is randomly chosen from among many, opened and verified to be consistent with other claims about it, the likelihood of this being true for the other commitments increases. Similarly, the process used to generate the voting system claims typically ensures that their truth values are not independent. For example, in Prêt à Voter, (Examples 3 and 4), the truth claims $P_{k_3}(g_3, P_{k_2}(g_2, P_{k_1}(g_1, P_{k_0}(g_0, D_0)))) - X_{(k,2)} = 0$ are verified on randomly chosen unfilled ballots—i.e. for $k \in \mathcal{K} \subset \mathcal{B}$, where $\mathcal{B}$ is the set of all printed ballots—which are spoilt through the verification. This verification decreases the uncertainty of the truth of similar claims about unspoilt ballots, i.e. for $k \in \bar{\mathcal{K}}$, the complement of $\mathcal{K}$. Typically, if voting system claims are correctly verified, the uncertainty in the truth of (some) related claims is reduced.

An *inference* is a logical conclusion following from the truth of some claims. For example, if voters were to verify the homomorphic encryptions of their votes, the appropriate combination of the encrypted values would be the encrypted value of the tally. In general, the truth of voter-verified claims and publicly-verified claims reduces the uncertainty in the truth of the inferences. The concluding inference is that the true vote tally is a particular value, $\widehat{V^{\Sigma}}$. The verifiability of a voting system is the amount of certainty in the truth of this inference.

Before concluding this section, we present two more examples.

Example 5. Consider the Farnel voting system. The voter-verifiable vote-dependent claims are that posted contents correspond exactly to the receipts and the initialization ballots, i.e. $V_j = y_k$ where y_k is the receipt, k its serial number, and j the unknown voter or the basket initializer (which is unknown even after verification of this claim). Voting system claims are claims about the ballots used to initialize the basket, and votes found in the basket at the end of voting. They are verifiable by examining the basket at the beginning and end of the election, and watching it through the election. There are no other claims. An inference is that the tally may be obtained by tallying the y_k, and subtracting out the initialization ballots.

Example 6. Consider the ThreeBallot voting system. Recall that a voter fills out B multiballots, each a ballot for C candidates. To choose a particular candidate, she puts a mark for that candidate in $B - 1$ multiballots, and she puts marks for all other candidates in $B - 2$ multiballots. The voter-verifiable vote-related claims are of the form $f(V_k, X_k, i) = Y_i$ where X_k denotes the voter's unknown strategy, and i the multiballot serial number. It is assumed that a voter casts exactly B multiballots and that they add up to a valid vote. There are no other claims. The inference is that the tally may be obtained by adding up all the marks for each candidate, and subtracting $(B - 1) \times n$, to obtain the final score for each candidate.

5.4.2. The definition of verifiability

Perfect verifiability occurs when there is no uncertainty whether the voting system used *SpecifiedVoteCount* to produce $\widehat{V^{\Sigma}}$, given the correctness proof provided by the system. We denote by T the random variable representing the truth of $Tally = \widehat{V^{\Sigma}}$, where *Tally* is the tally output by the voting system.

Definition 10. A voting system is **perfectly verifiable** when $\mathcal{H}(T|P) = 0\ \forall p_{V^*} \forall\ Tally$.

Equivalently, $\mathcal{I}(T; P) = \mathcal{H}(T)$. Note that a voting system may be perfectly verifiable even if the proof shows that the vote count was not obtained through its declared algorithm; we simply require that a system provide enough information to check its result.

Definition 11. The **verifiability measure** of a voting system is

$$\mathfrak{V} = \frac{\mathcal{I}(T; P)}{\mathcal{H}(T)}.$$

Notice that we do not define verifiability as a measure of the uncertainty in $V^{\Sigma} = Tally$; that is, we do not define it in terms of how close the purported vote count is to the true one. Such a definition would be a combination of our definition of verifiability (which connects *Tally* to $\widehat{V^{\Sigma}}$) and our definition of integrity (which connects $\widehat{V^{\Sigma}}$ to V^{Σ}). For verifiability, we simply require the system to demonstrate that it is indeed using *SpecifiedVoteCount*, which was quantified by its integrity measure.

Example 7. Consider a voting system that makes the following common "black box" DRE claim:

- During polling, V^* (and nothing else) goes in.
- After polling, V^{Σ} (and nothing else) comes out.

Here, $SpecifiedVoteCount(V^*, X^*) = V^{\Sigma}$, and the integrity is perfect. However, $E = Tally$; that is, $P = \emptyset$. Hence, $\mathcal{H}(T|P) = \mathcal{H}(T)$ and $\mathfrak{V} = 0$; the system has zero verifiability.

5.4.3. A special case: First-past-the-post elections

The above definition of verifiability is appropriate for use in elections where the exact vote tally is of interest, such as a legislative election involving proportional representation. However, we are often concerned only with the identity of the winning candidate, not with the margin of victory. (This is the case in "first-past-the-post" elections, including almost all elections in the US.) In such elections a weaker definition of verifiability can apply. We denote by W the random variable representing the truth of the sentence "The correct winner was declared in this election." (Note that by "correct winner", we mean the winner according to $\widehat{V^{\Sigma}}$.) A system provides perfect single-winner-verifiability if the uncertainty in W is zero, given the correctness proof of the system.

Definition 12. A single-winner voting system is **perfectly single-winner-verifiable** when $\mathcal{H}(W|P) = 0$.

Equivalently, $\mathcal{I}(W; P) = \mathcal{H}(W)$. The measure of single-winner-verifiability may be defined as follows.

Definition 13. The **single-winner-verifiability measure** of a voting system is

$$\mathfrak{W} = \frac{\mathcal{I}(W; P)}{\mathcal{H}(W)}.$$

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040

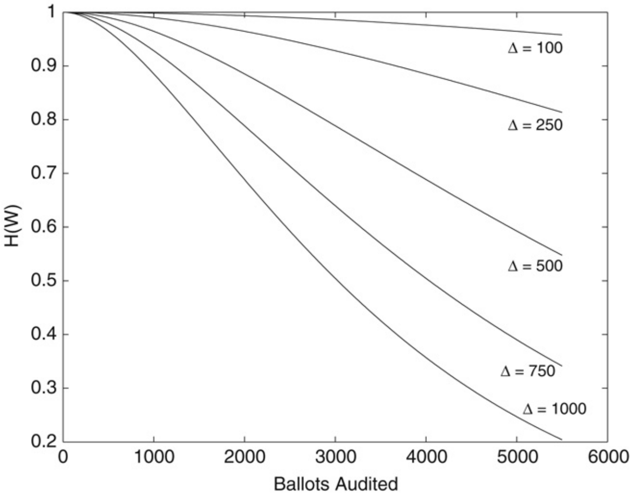

Fig. 2. Winner verifiability vs audit size for various margins of victory.

The probability that W is true is a function of the margin of victory in the election. If Δ is the margin of victory in number of votes, $p_D(f)$ is the probability of determining that f votes were modified, and $p_F(f)$ is the probability of these votes being modified, then the probability that W is true is the sum of the probabilities that the fraud is smaller than Δ and hence does not change the election outcome, added to the probabilities that it is greater and is detected.

$$p_W(W = True) = \sum_{f=0}^{\Delta} p_F(f) + \sum_{f=\Delta}^{n} p_D(f)p_F(f).$$

Example 8. Consider a single-winner plurality election. Each voter casts one vote, and the single choice that received the greatest total number of votes wins. Consider an independent end-to-end verification (E2E) voting system, such as PunchScan [3] or Prêt à Voter [1], where a set of m votes can be audited to test if the votes were correctly recorded and correctly contributed to the tally. The correctness of the audit results form part of the proof P (For example, P would include confirmation from m voters that their encrypted receipts were in the virtual ballot box, it would also include the proofs of correctness of the onions and/or mixnet decryption in Prêt à Voter). The value of $p_D(f)$ depends on the quality of the proof P, that is, on m.

Fig. 2 plots $\mathfrak{W}$ as a function of m (the audit size) for various values of Δ. We use the values of $p_D(f)$ derived in [31]. Not having information to posit the probability of fraud, we make the uniform assumption: the occurrence of fraud is as likely as the non-occurrence. Further, we would wish to examine the smallest fraud that could have changed the outcome of the election. Hence we assume that F is 0 with probability $\frac{1}{2}$ and $\frac{\Delta}{2}$ with probability $\frac{1}{2}$. This is simply an illustrative example; for a real election or a real voting system, the probability models would depend on the specifics of the political environment, and the ease of committing fraud.

6. Tradeoff between verifiability and privacy

In this section, we illustrate the use of the model in studying a general tradeoff among integrity, verifiability and privacy, proving that a voting system cannot achieve perfect integrity, perfect verifiability and perfect privacy.

Theorem. *A voting system cannot have perfect integrity, perfect privacy and perfect verifiability.*

Proof. Suppose the system has perfect integrity. That is, $V^{\Sigma} = \widehat{V^{\Sigma}}$. T—the truth of the statement $Tally = \widehat{v^{\Sigma}}$—is hence the truth of the statement $Tally = v^{\Sigma}$. For perfect verifiability,

$$\mathcal{H}(T|P) = 0 \ \forall \ Tally \ \forall p_{V^*}$$
$$\Rightarrow \mathcal{H}(V^{\Sigma}|P) = 0 \ \forall \ p_{V^*}.$$

That is, all values of V^* and X^* that satisfy the verified claims give the same value of V^{Σ}.

Consider the set of ith-voter-verifiable claims, and all voting system claims about the random variables in $\mathbf{X}_i$ and $\mathbf{X}'_i$. Denote the verified claims amongst these as $\mathbf{Y}_i$. Suppose V_i is not fixed by $\mathbf{Y}_i$, and there are at least two distinct values of V_i, $v_i \neq v'_i$ with corresponding values of $\mathbf{X}_i$ and $\mathbf{X}'_i$—$(v_i, \mathbf{x}_{(i,1)}, \mathbf{x}'_{(i,1)})$ and $(v'_i, \mathbf{x}_{(i,2)}, \mathbf{x}'_{(i,2)})$—that satisfy $\mathbf{Y}_i$. Choose $v_1, v_2, \dots, v_{i-1}, v_i, v_{i+1}, \dots, v_n$ and random variables X^*, with $\mathbf{X}_i = \mathbf{x}_{(i,1)}$ and $\mathbf{X}'_i = \mathbf{x}'_{(i,1)}$, so that all verified claims are satisfied. Change only the values of V_i and $\mathbf{X}_i$ and $\mathbf{X}'_i$ to $(v'_i, \mathbf{x}_{(i,2)}, \mathbf{x}'_{(i,2)})$. Then the new values also satisfy all the claims, because (i) they satisfy the values of $\mathbf{Y}_i$, and (ii) changes in V_i, $\mathbf{X}_i$ and $\mathbf{X}'_i$ do not affect $\bar{\mathbf{Y}}_i$, the set of all other verified claims. But the new values also result in a different value of V^{Σ}, which is a contradiction. Hence $\mathcal{H}(V_i|\mathbf{Y}_i) = 0$.

In the absence of a Farnel basket, voter i would verify the ith-voter-verifiable claims of $\mathbf{Y}_i$. Further, all voting system claims about the random variables in $\mathbf{X}_i$ and $\mathbf{X}'_i$ would also be verified, and voter i's vote would be associated with her, hence privacy is not perfect.

If a Farnel basket were assumed to exist, the ith-voter-verifiable claims of $\mathbf{Y}_i$ would be verified by voters j such that $j \geq i$. If some bogus claims, such as those initializing the Farnel basket, are not used, then voter i verifies the claims relating to her vote. Hence, we assume the existence of some number of known bogus claims. Consider the first voter to cast a vote. She creates a number of first-voter-verifiable claims, and she verifies as many as she creates. Some of the claims she verifies are bogus, though it is not known which ones are bogus. By casting her vote, she has biased the distribution of possible claims she will verify, originally consisting only of bogus claims, towards her claims. Hence the claims she verifies reduce the entropy of her vote, hence privacy is not perfect. ■

7. Privacy losses: Farnel and ThreeBallot voting systems

In this section, we compare the privacy loss of the Farnel and ThreeBallot voting systems, the main E2E voting systems that do not explicitly use cryptography to encrypt votes, yet provide voter-verifiable receipts. We also examine the specific-case privacy loss due to knowledge of single receipts in the Farnel voting system, and study its relationship with various system parameters.

7.1. Privacy loss in the Farnel voting system

In this section we measure the privacy loss in the Farnel voting system. The main source of privacy loss is that voter i's receipt provides information about voter j's vote, for $j \leq i$. For illustrative purposes, we assume the election is expected to be keenly contested, and that voters are independent; that is, that the votes are identically, independently and uniformly distributed.

Consider a Farnel election with C candidates and a basket capacity of B ballots. (Strictly, in order to properly initialize the basket, B must be a multiple of C.) Let $\mathcal{V} = \{1, 2, \dots, C\}$. We denote the receipt obtained by the ith voter by R_i ($R_i \in \mathcal{V}$). Voter i receives their own ballot as their receipt with probability $\frac{1}{B+1}$. They could also receive the ballot of a previous voter that happens to be equal to their ballot. These two possibilities are combined to derive a formula for an upper bound on $\mathcal{L}$ in the appendix (Appendix A.1), assuming a uniform distribution of the basket's contents to begin with, as well as a near-uniform election. Fig. 3 plots the value of this upper bound on $\mathfrak{L}$ as a function of the basket size and the number of candidates when the number of votes is very large. Note that the upper bound increases with the number of candidates on the ballot. This is because, with a fixed basket size, a larger number of candidates implies fewer ballots per candidate in the basket, and hence a greater perturbation of the basket distribution by the voter's vote.

7.2. The privacy loss in the ThreeBallot voting system

In this section we examine the privacy loss of a ThreeBallot voting system with N candidates (rows) and B subballots (columns) in order to compare with the privacy loss of the Farnel system. As with the Farnel system, we assume the votes are uniformly, identically and independently distributed.

Note that the optimal strategy for a voter is to take a multiballot that completely hides her vote, that is, one filled uniformly at random. Such a receipt reveals no information on the vote. However, assuming that the typical voter does not use a strategy in filling or picking the multiballot(s), we examine the case when each mark is equally likely given the voter's choice, and each multiballot is equally likely to be chosen as a receipt. The expression for $\mathfrak{L}$ is derived in Appendix A.2.

Fig. 4 plots the function $\frac{\mathcal{I}(V|R)}{\mathcal{H}(V)}$ which is $(1 - \frac{\mathcal{H}(V^{\Sigma})}{\mathcal{H}(V^*)}) \times \mathfrak{L} \approx \mathfrak{L}$ for large n. While the privacy loss of Farnel and of ThreeBallot appear to be of a similar magnitude, the Farnel privacy loss increases with an increase in the number of candidates.

7.3. A specific-case privacy loss in the Farnel voting system

In this section, we illustrate the use of another privacy measure in our model, that of specific case privacy loss, see Eq. (2). Perhaps the simplest and most obvious privacy attack in the Farnel voting system is to determine a voter's vote by looking at her receipt. This is a simple attack, requiring no sophisticated equipment, algorithms, or mathematics. In this section, we measure the specific-case privacy loss due to this attack. A is the kth voter's receipt R_k, and V^S is the single vote of that voter, V_k.

Consider a Farnel election with C candidates and a basket capacity of B ballots. Let the fraction of votes in the basket that correspond to choice i be α_i. To determine $p_{V_k|R_k}$, which would be the goal of an attacker, we assume that the fraction of votes for i in the community is β_i (that is, $p_{V_k}(V_k = i) = \beta_i$). We further assume that p_{V_k} is non-trivial, i.e $0 < \beta_i < 1$.

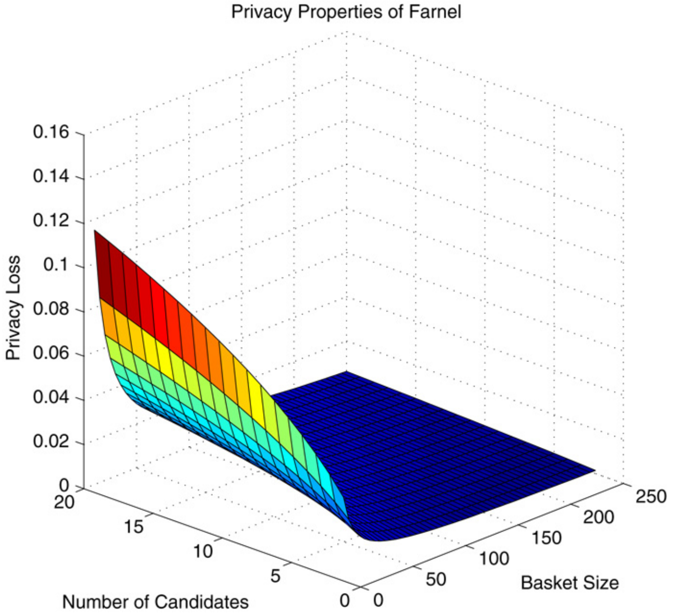

Fig. 3. An upper bound on $\mathcal{L}$: Farnel, as a function of B and C, when the votes are independently, identically and uniformly distributed, the basket is initialized with an equal number of votes for each candidate, and the number of voters is large.

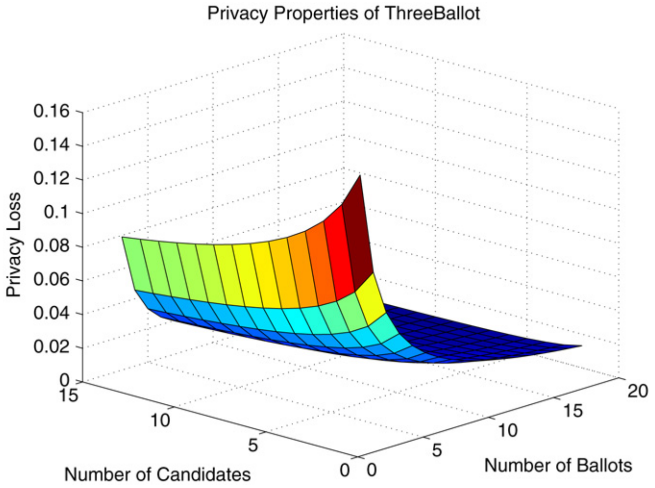

Fig. 4. $\mathcal{L}$: ThreeBallot as a function of B and C, when the votes are independently, identically and uniformly distributed and the number of voters is large.

Appendix A.3 derives the various formulae for probabilities. From the formulae it is clear that, if a voter possesses a receipt for candidate j, her probability of having voted for candidate j is greater than β_j, which is the probability of her having voted for candidate j if her receipt were not known. Also her probability of having voted for a candidate other than j is smaller than it was when her receipt was not known.

Fig. 5 plots $\mathfrak{L}_{V,R}$ as a function of B and C assuming the β_i are all equal, and the α_i are all equal. That is, that the election is close. (We assume the basket distribution follows that of the votes at steady state.) Observe that it behaves much as the privacy loss $\mathfrak{L}$, see Fig. 3.

We also examine the case of $C = 2$ in order to examine the specific case privacy loss as a function of Farnel system parameters. Appendix A.4 provides details of the derivation of the formula. Fig. 6 plots $\mathfrak{L}_{V,R}$ as a function of β and α when $C = 2$ and $B = 5$, using the probability distributions of (7). This provides the privacy loss for an early voter before the basket distribution mimics the vote distribution. Observe that the privacy loss decreases as the voter distribution moves away from uniform, and as the Farnel basket distribution moves towards uniform. We also observed that privacy loss, in general, as expected, decreases as basket size increases. Finally, Fig. 7 plots the privacy loss, for $\alpha = \beta$, assuming the basket distribution is that of the votes at steady state, as a function of B and β. Notice that the privacy loss is least for an even vote distribution.

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040

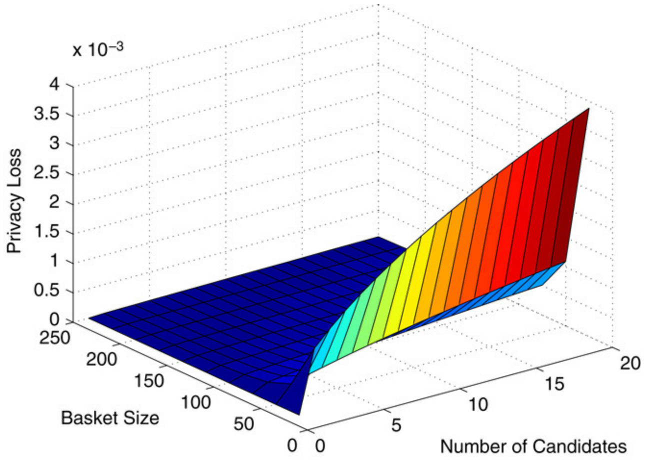

Fig. 5. $\mathcal{L}_{V,R}$: Farnel, as a function of C and B, when the votes are independently, identically and uniformly distributed and the number of voters is large.

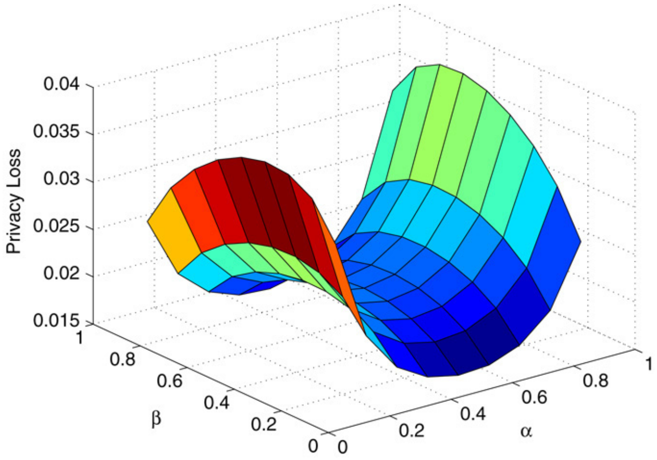

Fig. 6. $\mathcal{L}_{V,R}$: Farnel, as a function of α and β, for $B = 5$ and $C = 2$.

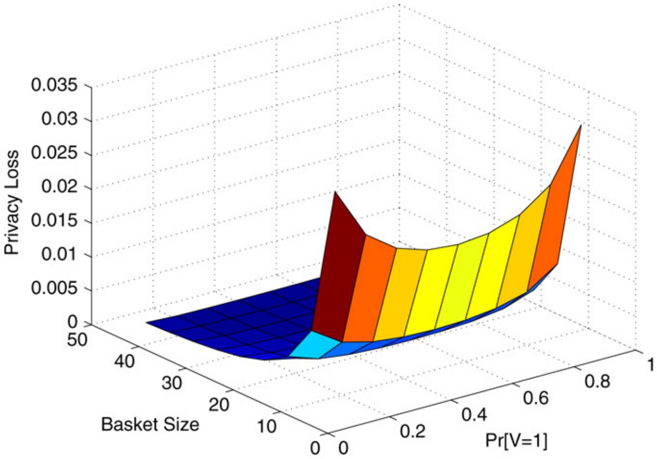

Fig. 7. $\mathcal{L}_{V,R}$: Farnel, as a function of $\beta = \alpha$, $B = 5$, $C = 2$.

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040

8. Conclusions

We have presented the beginnings of an information-theoretic approach to rating voting systems for integrity, privacy and verifiability. We have used this framework to show that tradeoffs exist among integrity, verifiability and privacy, and have applied it to the measurement of the privacy of the ThreeBallot and Farnel voting systems. Future directions include the use of this model to measure other voting systems. In particular, this model could be used to measure the privacy loss of randomized partial audits of mixnets [6]. It could also be used to examine the possibility of obtaining verifiability and privacy in voting systems that do not use cryptography.

Acknowledgment

Both of the authors were supported in part by NSF award SGER 0505510.

Appendix A

A.1. An upper bound on $\mathfrak{L}$ in the Farnel voting system

In this section we derive the privacy loss for the Farnel voting system assuming a close election and a uniform initialization of the basket.

Recall the Farnel voting system with C candidates and basket size B. A voter casts vote V_i in the basket, and obtains a receipt, denoted R_i, $R_i, V_i, \quad \in \mathcal{V}$. The receipt is chosen at random from among the contents of the basket. The basket is initialized with B votes. If we assume a uniform random distribution of these B votes, and a close contest, the basket contains a near uniform distribution. If a voter casts vote V_i, the number of receipts in the basket that represent V_i are $\frac{B}{C} + 1$, and the total number of possibilities in the basket are $B + 1$. Hence,

$$p_{R_i|V_i}(v_i; v_i) = \frac{\frac{B}{C} + 1}{B + 1} = \frac{B + C}{C(B + 1)} = \frac{B + C}{CB + C}.$$

Because future voters may also receive voter i's receipt, their receipts also leak information about voter i's vote. Because there is a $\frac{1}{B+1}$ chance that each voter between voters i and j (including voter i) removes voter i's ballot from the possible receipts if it is there, the probability of the jth receipt given the ith vote is:

$$p_{R_j|V_i}(r_j; v_i) \begin{cases} \left(\left(\frac{B}{B+1}\right)^{j-i} \times \frac{B+C}{CB+C}\right) + \left(\left(1 - \left(\frac{B}{B+1}\right)^{j-i}\right) \times C^{-1}\right) & r_j = v_i \\ \left(\left(\frac{B}{B+1}\right)^{j-i} \times \frac{B}{CB+C}\right) + \left(\left(1 - \left(\frac{B}{B+1}\right)^{j-i}\right) \times C^{-1}\right) & r_j \neq v_i \end{cases} \tag{3}$$

where $j \geq i$. The value of $p_{R_j|V_i}(r_j|v_i)$ depends only on whether $r_j = v_i$ or not; we denote $p_{R_j|V_i}(r_j|v_i)$ by P_1 when $r_j = v_i$ and by P_2 when $r_j \neq v_i$. Assuming V_i is uniformly distributed, note that

$$p_{R_j}(r_j) = \sum_{v_i=1}^{C} p_{R_j|V_i}(r_j; v_i)C^{-1} = (P_1 + (C - 1)P_2)C^{-1} = C^{-1}$$

and the jth receipt is also uniformly distributed. Further, we have, using Bayes' rule:

$$\begin{aligned} p_{V_i|R_j}(v_i; r_j) &= \frac{p_{R_j|V_i}(r_j; v_i)p_{V_i}(v_i)}{p_{R_j}(r_j)} \\ &= \frac{p_{R_j|V_i}(r_j; v_i)C^{-1}}{C^{-1}} \\ &= p_{R_j|V_i}(v_j; r_i). \end{aligned}$$

Finally,

$$H(V_i|R_j = r_j) = -P_1 \log_2 P_1 - (C - 1)P_2 \log_2 P_2$$

is independent of r_j.

Hence we get,

$$\mathcal{H}(V_i|R_j) = -\sum_{R_j} p_{R_j}(r_j)H(V_i|R_j = r_j) = H(V_i|R_j = r_j).$$

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040

If R^* is the (ordered) vector of receipts, then we have:

$$\mathcal{I}(V_i; R^*) \leq \sum_{j=i}^{|V^*|} \mathcal{I}(V_i; R_j)$$

$$\mathcal{I}(V^*|R^*) \leq \sum_{i=1}^{|V^*|} \sum_{j=i}^{|V^*|} \mathcal{I}(V_i; R_j) = n\mathcal{I}(V_i; R^*).$$

Further,

$$\mathcal{H}(V^*|V^\Sigma; R^*) + \mathcal{H}(V^\Sigma) = \mathcal{H}(V^*, V^\Sigma|R^*) = \mathcal{H}(V^*|R^*)$$

as V^Σ is completely determined by V^*. Similarly, $\mathcal{H}(V^*|V^\Sigma) + \mathcal{H}(V^\Sigma) = \mathcal{H}(V^*)$. This allows computation of an upper bound on $\mathfrak{L}$ as follows:

$$\begin{aligned}\mathfrak{L} &= \frac{\mathcal{I}(V^*; V^\Sigma; R^*)}{\mathcal{H}(V^*|V^\Sigma)} \\ &= \frac{\mathcal{H}(V^*|V^\Sigma) - \mathcal{H}(V^*|V^\Sigma; R^*)}{\mathcal{H}(V^*|V^\Sigma)} \\ &= \frac{\mathcal{H}(V^*) - \mathcal{H}(V^\Sigma) - \mathcal{H}(V^*|R^*) + \mathcal{H}(V^\Sigma)}{\mathcal{H}(V^*) - \mathcal{H}(V^\Sigma)} \\ &= \frac{\mathcal{I}(V^*; R^*)}{\mathcal{H}(V^*) - \mathcal{H}(V^\Sigma)} \\ &\leq \frac{n\mathcal{I}(V_i; R^*)}{n\mathcal{H}(V) - \mathcal{H}(V^\Sigma)} \\ &\leq \frac{n \sum_{j=i}^{|V^*|} \mathcal{I}(V_i; R_j)}{n\mathcal{H}(V) - \mathcal{H}(V^\Sigma)}.\end{aligned}$$

Hence:

$$\left(1 - \frac{\mathcal{H}(V^\Sigma)}{n\mathcal{H}(V)}\right) \times \mathfrak{L} \leq \frac{\sum_{j=i}^{|V^*|} \mathcal{I}(V_i; R_j)}{\mathcal{H}(V_i)}.$$

Note that the ratio of the entropy in the tally, $\mathcal{H}(V^\Sigma)$, to that in the votes, $n\mathcal{H}(V_i)$, is a function of the number of votes, and is smaller for a larger number of votes. Hence, for a large enough number of votes,

$$\mathfrak{L} \leq \frac{\sum_{j=i}^{|V^*|} \mathcal{I}(V_i; R_j)}{\mathcal{H}(V_i)}$$

is approximately correct.

A.2. Derivation of $\mathfrak{L}$ for the ThreeBallot voting system

Recall the ThreeBallot voting system with B multiballots and C candidates. The ballot consists of C rows and B columns. To choose a candidate, the voter marks $B-1$ columns for that voter, and $B-2$ for the others.

Suppose (wlog) that a certain voter votes for the first candidate (row 1) on the ballot. Using [32], we derive the probability distribution on the voter's receipt, given the value of the vote. Also suppose that there are B ballot-columns and C candidates (rows). If there is a mark in row 1, then there are $\binom{B-1}{B-2} = B-1$ possible ways to distribute the remaining $B-2$ marks among the remaining $B-1$ ballot-columns. Otherwise there is only $\binom{B-1}{B-1} = 1$ possible way to distribute the remaining $B-1$ marks. Similarly, if there is a mark in row i (where $2 \leq i \leq C$), then there are $\binom{B-1}{B-3}$ ways to distribute the remaining $B-3$ marks; if not, there are $\binom{B-1}{B-2} = B-1$ ways to distribute the remaining $B-2$ marks.

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040

Now, from [32], for each candidate, we have $B\binom{B}{B-1}\binom{B}{B-2}^{C-1}$ possible mark-choice patterns on the receipt (B possible receipts, and $\binom{B}{B-1}\binom{B}{B-2}^{C-1}$ ways to fill a ballot). Since each receipt can appear in any of the B possible columns; we can factor out $\frac{B}{B}$, leaving us with:

$$p_{R_i|V_i}[r_i|v_i=1]=\frac{\binom{B-1}{B-1-m_1}\prod_{j=2}^{C}\binom{B-1}{B-2-m_i}}{\binom{B}{B-1}\binom{B}{B-2}^{C-1}}$$

where m_i is 1 if there is a mark on row i of the receipt, and 0 otherwise. If the vote is actually for candidate x (instead of candidate 1), we can still use this formula, but instead of using m_i, we use $m_{\pi_x(i)}$, where π_x is a permutation of $\{1, 2, \ldots, C\}$ with $\pi_x(1)=x$, $\pi_x(x)=1$, and $\pi_x(i)=i$ otherwise. Therefore,

$$p_{R_i|V_i}[r_i|v_i]=\frac{\binom{B-1}{B-1-m_{\pi_{v_i}(1)}}\prod_{j=2}^{C}\binom{B-1}{B-2-m_{\pi_{v_i}(j)}}}{\binom{B}{B-1}\binom{B}{B-2}^{C-1}}$$

is the probability distribution for the receipt given the vote. By applying Bayes' Rule, we get:

$$\begin{aligned}p_{V_i|R_i}[v_i|r_i]&=\frac{p_{R_i|V_i}(r_i;v_i)p_{V_i}(v_i)}{p_{R_i}(r_i)}\\&=\frac{p_{R_i|V_i}(r_i;v_i)\times C^{-1}}{\sum_{v_i=1}^{C}p_{R_i|V_i}(r_i;v_i)\times C^{-1}}\\&=\frac{\binom{B-1}{B-1-m_{\pi_{v_i}(1)}}\prod_{j=2}^{C}\binom{B-1}{B-2-m_{\pi_{v_i}(j)}}}{\sum_{v_i=1}^{C}\binom{B-1}{B-1-m_{\pi_{v_i}(1)}}\prod_{j=2}^{C}\binom{B-1}{B-2-m_{\pi_{v_i}(j)}}}.\end{aligned}$$

This allows us to compute $\mathfrak{L}$ for a ThreeBallot election with N rows (candidates) and B columns as $\mathcal{I}(V;R)$.

A.3. Derivation of $\mathfrak{L}_{V,R}$ for the Farnel voting system for uniform vote distribution

In this section we derive the specific-case privacy loss for Franel, assuming a uniform vote distribution. That is, we measure the privacy loss for an individual vote V_k when the adversary has access only to the voter's receipt, R_k.

We observe the following:

$$p_{R_k|V_k}(R_k=j;V_k=i)=\begin{cases}\dfrac{\alpha_jB+1}{B+1} & i=j\\[2ex]\dfrac{\alpha_jB}{B+1} & i\neq j.\end{cases}\tag{4}$$

In this example it is very clear how the voting system behaves as a communication channel transmitting information about the vote to the adversary. The channel input is the vote V_k, the output is the receipt R_k. The equation above provides the error probabilities of the channel, i.e. the probabilities with which the receipt does not correspond to the vote. It also provides the probability of correct transmission.

By Bayes rule we obtain:

$$p_{V_k|R_k}(V_k=i;R_k=j)=\begin{cases}\dfrac{\beta_i(\alpha_jB+1)}{\alpha_jB+\beta_j}>\beta_i=p_{V_k}(V_k=i) & i=j\\[2ex]\dfrac{\beta_i\alpha_jB}{\alpha_jB+\beta_j}<\beta_i=p_{V_k}(V_k=i) & i\neq j.\end{cases}\tag{5}$$

To see how this changes with B and C, we assume that $\alpha_j=\beta_j=C^{-1}$.

$$p_{V_k|R_k}(V_k=i;R_k=j)=\begin{cases}\dfrac{\frac{B}{C}+1}{B+1} & i=j\\[2ex]\dfrac{\frac{B}{C}}{B+1} & i\neq j.\end{cases}\tag{6}$$

We can now easily derive $\mathcal{L}_{V,R} = \frac{\mathcal{I}(V_k|V^{\Sigma};R_k)}{\mathcal{H}(V_k|V^{\Sigma})}$, by noting that, for a large number of votes, $\mathcal{H}(V_k|V^{\Sigma}) \approx \mathcal{H}(V_k)$ as the tally reveals no more information about V_k than is already known through knowledge of its probability distribution. Hence $\mathcal{L}_{V,R} \approx \frac{\mathcal{I}(V_k;R_k)}{\mathcal{H}(V_k)}$.

A.4. Derivation of $\mathcal{L}_{V,R}$ for Farnel when $C = 2$

In this section we derive the specific-case privacy loss for Farnel assuming $C = 2$ and letting β_j and α_j vary.

Let $\beta_1 = \beta$, then $\beta_2 = 1 - \beta$, similarly $\alpha_1 = \alpha$ and $\alpha_2 = 1 - \alpha$, then, Eq. (5) is:

$$p_{V_k|R_k}(v_k; r_k) = \begin{cases} \dfrac{\beta(\alpha B + 1)}{\alpha B + \beta} & v_k = r_k = 1 \\ \dfrac{(1-\beta)\alpha B}{\alpha B + \beta} & v_k = 2; r_k = 1 \\ \dfrac{(1-\beta)((1-\alpha)B + 1)}{(1-\alpha)B + (1-\beta)} & v_k = r_k = 2 \\ \dfrac{\beta(1-\alpha)B}{(1-\alpha)B + (1-\beta)} & v_k = 1; r_k = 2 \end{cases} \tag{7}$$

using which we obtain the privacy loss for various values of B, α and β.

References

[1] P.Y.A. Ryan, A Variant of the Chaum Voter-verifiable Scheme, Tech. Rep. CS-TR: 864, School of Computing Science, Newcastle University, 2004.
[2] D. Chaum, P.Y.A. Ryan, S.A. Schneider, A practical, voter-verifiable election scheme, Tech. Rep. CS-TR: 880, School of Computing Science, Newcastle University, 2004.
[3] S. Popoveniuc, B. Hosp, An introduction to Punchscan, in: IAVoSS Workshop On Trustworthy Elections, WOTE, 2006.
[4] R.L. Rivest, W.D. Smith, Three voting protocols: ThreeBallot, VAV, and Twin, in: Proceedings of EVT'07 (Electronic Voting Technology Workshop), 2007.
[5] R. Araujo, R.F. Custodio, J. van de Graaf, A verifiable voting protocol based on Farnel, in: IAVoSS Workshop On Trustworthy Elections, WOTE, 2007.
[6] M. Gomulkiewicz, M. Klonowski, M. Kutylowski, Rapid mixing and security of Chaum's visual electronic voting, in: ESORICS, 2003.
[7] D. Chaum, J. van de Graaf, P.Y.A. Ryan, P.L. Vora, Secret Ballot elections with unconditional integrity, IACR eprint archive, no. 2007/270.
[8] M. Jakobsson, A. Juels, R.L. Rivest, Making mix nets robust for electronic voting by randomized partial checking, in: USENIX Security, 2002, pp. 339–353.
[9] L. Coney, J.L. Hall, P.L. Vora, D. Wagner, Towards a privacy measurement criterion for voting systems, in: National Conference on Digital Government Research, May 2005.
[10] C. Crutchfield, D. Molnar, D. Turner, Approximate measurement of voter privacy loss in an election with precinct reports, 2006. VSRW'06.
[11] M.I. Shamos, Electronic voting — evaluating the threat, Computers, Freedom and Privacy (1993).
[12] WOTE01 Workshop on Trustworthy Elections, 2002. http://www.vote.caltech.edu/wote01/.
[13] WEST 2002 Workshop on Election Standards and Technology, 2002. http://www.vote.caltech.edu/west02/presentations.html.
[14] E. Gerck, Voting system requirements, in: Workshop on Trustworthy Elections, WOTE, 2001.
[15] D. Jefferson, Requirements for electronic and internet voting systems in public elections, in: Workshop on Trustworthy Elections, WOTE, 2001.
[16] D.W. Jones, End-to-end standards for accuracy in paper-based systems, 2002. http://www.cs.uiowa.edu/ jones/voting/west02/.
[17] D.W. Jones, Evaluating voting technology, 2001. http://www.cs.uiowa.edu/ jones/voting/uscrc.html.
[18] Workshop on threats to voting systems. National Institute of Standards and Technology (NIST), October 2005. http://vote.nist.gov/threats/.
[19] D.W. Jones, Threats to voting systems, NIST Workshop on Threats to Voting Systems, October 2005.
[20] Developing an analysis of threats to voting systems: Workshop summary. National Institute of Standards and Technology (NIST), October 2005, http://vote.nist.gov/threats/workshop_summary.pdf.
[21] C. Diaz, S. Seys, J. Claessens, B. Preneel, Towards measuring anonymity, in: Privacy Enhancing Technologies (PET), 2002.
[22] A. Serjantov, G. Danezis, Towards an information theoretic metric for anonymity, in: Privacy Enhancing Technologies (PET), 2002.
[23] D. Agrawal, C.C. Aggarwal, On the design and quantification of privacy preserving data mining algorithms, in: Proceedings of the 20th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS, 2001.
[24] P.L. Vora, An information-theoretic approach to inference attacks on random data perturbation and a related privacy measure, IEEE Transactions on Information Theory 53 (August) (2007) 2971–2977.
[25] A. Evfimievski, J. Gehrke, R. Srikant, Limiting privacy breaches in privacy preserving data mining, in: Proceedings of the 22nd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, PODS, 2003.
[26] B. Hosp, P.L. Vora, An information-theoretic model of voting systems, in: IAVoSS Workshop on Trustworthy Elections, WOTE, 2006.
[27] A. Kiayias, M. Yung, Self-tallying elections and perfect ballot secrecy, in: Public Key Cryptography — 5th International Workshop on Practice and Theory in Public Key Cryptosystems, 2002, pp. 141–158.
[28] D.L. Chaum, Untraceable electronic mail, return address, and digital pseudonym, Communication of ACM (February) (1981).
[29] T.M. Cover, J.A. Thomas, Elements of Information Theory, Wiley-Interscience, 1991.
[30] A.C. Yao, Theory and applications of trapdoor functions, in: Proc. of the 23rd Annual IEEE Symposium on Foundations of Computer Science, 1982.
[31] C.A. Neff, Election confidence: A comparison of methodologies and their relative effectiveness at achieving it, 2003. http://www.votehere.com/papers/ElectionConfidence.pdf.
[32] J. Clark, A. Essex, C. Adams, On the security of ballot receipts in E2E voting systems, in: IAVoSS Workshop On Trustworthy Elections, WOTE, 2007.

Please cite this article in press as: B. Hosp, P.L. Vora, An information-theoretic model of voting systems, Mathematical and Computer Modelling (2008), doi:10.1016/j.mcm.2008.05.040