

Defining Privacy for Weighted Votes, Single and Multi-voter Coercion

Jannik Dreier, Pascal Lafourcade, and Yassine Lakhnech

Université Grenoble 1, CNRS, Verimag, France
{Jannik.Dreier,Pascal.Lafourcade,Yassine.Lakhnech}@imag.fr

Abstract. Most existing formal privacy definitions for voting protocols are based on observational equivalence between two situations where two voters swap their votes. These definitions are unsuitable for cases where votes are weighted. In such a case swapping two votes can result in a different outcome and both situations become trivially distinguishable. We present a definition for privacy in voting protocols in the Applied π -Calculus that addresses this problem. Using our model, we are also able to define multi-voter coercion, i.e. situations where several voters are attacked at the same time. Then we prove that under certain realistic assumptions a protocol secure against coercion of a single voter is also secure against coercion of multiple voters. This applies for Receipt-Freeness as well as Coercion-Resistance.

1 Introduction

Privacy is a key requirement in elections as voters can otherwise be blackmailed, coerced or may be susceptible to vote-buying. Typically privacy is split into three different properties:

- *Vote-Privacy*: The votes are kept private.
- *Receipt-Freeness*: A voter cannot construct a receipt which allows him to prove to a third party that he voted for a certain candidate. This is to prevent vote-buying.
- *Coercion-Resistance*: Even when a voter interacts with a coercer during the entire voting process, the coercer cannot be sure whether the voter followed his instructions or actually voted for another candidate.

The design of complex protocols such as voting protocols is known to be error-prone. This is why formal verification is an ideal tool to ensure the correctness and security of voting protocols. It has already been used to analyze properties such as Verifiability, Privacy, Receipt-Freeness and Coercion-Resistance [1–12].

However, most existing symbolic definitions of Privacy are based on the idea of swapping votes. If the votes are private, a case where Alice votes “yes” and Bob votes “no” should be indistinguishable from a case where Alice votes “no” and Bob votes “yes”. Yet this definition is unsuitable for some situations, for example in companies where votes are weighted according to the proportion of

shares held by each shareholder. Consider the following example: Alice owns 50% of the stocks, and Bob and Carol each hold 25%. The cases where Alice and Bob swap votes are now easily distinguishable if Carol votes “yes” all the time, as the result of the vote is different: 75% vs. 50% vote for “yes”. Note that there are still situations where privacy is ensured in the sense that different situations give the same result. The last outcome (50% yes, 50% no) could - for example - also be announced if Alice votes “yes” and Bob and Carol vote “no”. Protocols supporting vote weights have been proposed, for example Eliasson and Zúquete [13] developed a voting system supporting vote weights based on REVS [14], which itself is based on the protocol by Fujioka et al. [15].

Our Contributions: To address this issue, we define a symbolic privacy notion in the Applied π -Calculus [16] that takes weighted votes into account. Instead of requiring two executions where voters swap votes to be bisimilar, we require two executions to be bisimilar if they publish the same result, independent of the mapping between voters and votes. We analyze the relationship of our notion to the existing swap-based ones and give precise conditions for formally proving the equivalence between them. Then, we generalize our notion to Receipt-Freeness and Coercion-Resistance for weighted votes. We use a variant of the protocol by Eliasson and Zúquete [13] as a case study for our definition, and provide a partially automated proof using ProVerif [17].

In the cases of coercion most existing definitions only consider one attacked voter. Our model also allows to define a case with multiple coerced voters, and we analyze the relationship between this and the single-voter case. In particular, we give a formal proof that single- and multi-voter coercion are equivalent for a given protocol if it satisfies some modularity and de-composability properties. Using two existing protocols, we show that these properties are realistic.

Related Work: Previous research on formal verification of voting protocols concerned privacy properties (privacy, receipt-freeness and coercion-resistance) [1–8], election verifiability [9, 10], or both [11, 12].

In the symbolic model, privacy is usually defined as observational equivalence of two cases where a pair of voters swap their votes [1–5, 8, 12]. The definitions mainly differ in the way they model voting processes and deal with specifics of protocols. Some of them can be verified automatically using standard tools (e.g. ProVerif [17] and ProSwapper [18]). This swap-based approach was not designed for weighted votes and, as explained above, may lead to unexpected results in this case.

The other main approach roots in the computational model. In this case the real-world protocol is compared to an ideal situation and the attacker’s advantage is analyzed [7, 11]. Our symbolic definition is somewhat related to this computational approach, as we also consider some information – the result – to be leaked even in an ideal situation, and only forbid further leakage. Another possibility is to consider the overall advantage of the attacker without comparing it to an ideal situation [19]. This advantage is always non-negligible as in certain situations the votes are always revealed, e.g. in the case of an unanimous vote.

A third approach was proposed by Langer et al. [20]. The authors developed verifiability and privacy notions based on (un-)linkability between a voter and his vote. Their definitions have to be instantiated with a concrete formal process and attacker model. To define unlinkability, they rely on indistinguishability of runs where votes are swapped, with the same issues as described above.

Küsters and Truderung [6] were the first to explicitly consider multi-voter coercion. In their abstract model, Single-Voter Coercion and Multi-Voter Coercion turned out to be different in general. Subsequently they proposed a modified definition of Coercion-Resistance that implies both Single- and Multi-Voter Coercion-Resistance. In our model Single- and Multi-Voter Coercion are equivalent under certain assumptions on the protocol, hence we do not need to change the initial definition. Additionally, the conditions allow us to precisely characterize the difference between both notions.

Outline: In the next section, we present the Applied π -Calculus and recall the privacy definitions given by Delaune et al. [2]. In Section 3, we introduce our privacy definition and show under which condition it is equivalent to the existing ones. Then, in Section 4, we define Single- and Multi-Voter Receipt-Freeness, analyze their relationship and prove their equivalence under certain assumptions. In Section 5 we define Single- and Multi-Voter Coercion-Resistance and again prove their equivalence under the same hypotheses, before concluding in Section 6.

2 Preliminaries

In this section we recall the Applied π -Calculus, introduce our model of voting protocols and present existing privacy definitions.

2.1 Applied π -Calculus

The Applied π -Calculus [16] is a formal language to describe concurrent processes. The calculus consists of *names* (which typically correspond to data or channels), *variables*, and a *signature* Σ of *function symbols* which can be used to build *terms*. Functions typically include encryption and decryption – for example $\text{enc}(\text{message}, \text{key})$, $\text{dec}(\text{message}, \text{key})$ – hashing, signing etc. Terms are correct (i.e. respecting arity and sorts) combinations of names and functions. We distinguish the type “channel” from other *base* types. To model equalities we use an equational theory E which defines a relation $=_E$. A classical example which describes the correctness of symmetric encryption is $\text{dec}(\text{enc}(\text{message}, \text{key}), \text{key}) =_E \text{message}$. Processes are constructed using the grammars detailed in Figure 1.

The substitution $\{M/x\}$ replaces the variable x with term M . We denote by $fv(A)$, $bv(A)$, $fn(A)$, $bn(A)$ the free variables, bound variables, free names or bound names respectively. A process is *closed* if all variables are bound or defined by an active substitution. The *frame* $\Phi(A)$ of an active process A is obtained by replacing all plain processes in A by 0. This frame can be seen as a representation

$P, Q, R :=$	plain processes	$A, B, C :=$	active processes
0	null process	P	plain process
$P Q$	parallel composition	$A B$	parallel composition
$!P$	replication	$\nu n.A$	name restriction
$\nu n.P$	name restriction (“new”)	$\nu x.A$	variable restriction
if $M = N$ then P	conditional	$\{M/x\}$	active substitution
else Q			
in $(u, x).P$	message input		
out $(u, x).P$	message output		

(a) Plain process

(b) Extended process

Fig. 1. Grammars for *plain* and *extended* or *active* processes

of what is statically known to the exterior about a process. The domain $\text{dom}(\Phi)$ of a frame Φ is the set of variables for which Φ defines a substitution. An evaluation context $C[_]$ denotes an active process with a hole for an active process that is not under replication, a conditional, an input or an output. In the rest of the paper we use the following usual notions of equivalence and bisimilarity based on the original semantics [16] given in our technical report [21].

Definition 1 (Equivalence in a Frame [16]). *Two terms M and N are equal in the frame ϕ , written $(M = N)\phi$, if and only if $\phi \equiv \nu \tilde{n}.\sigma$, $M\sigma = N\sigma$, and $\{\tilde{n}\} \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$ for some names \tilde{n} and some substitution σ .*

Definition 2 (Static Equivalence (\approx_s) [16]). *Two closed frames ϕ and ψ are statically equivalent, written $\phi \approx_s \psi$, when $\text{dom}(\phi) = \text{dom}(\psi)$ and when for all terms M and N we have $(M = N)\phi$ if and only if $(M = N)\psi$. Two extended processes A and B are statically equivalent ($A \approx_s B$) if their frames are statically equivalent.*

The intuition behind this definition is that two processes are statically equivalent if the messages exchanged with the environment cannot be distinguished by an attacker (i.e. all operations on both sides give the same results). This idea can be extended to *labeled bisimilarity*.

Definition 3 (Labeled Bisimilarity (\approx_l) [16]). *Labeled bisimilarity is the largest symmetric relation \mathcal{R} on closed active processes, such that $A \mathcal{R} B$ implies:*

1. $A \approx_s B$,
2. if $A \rightarrow A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ,
3. if $A \xrightarrow{\alpha} A'$ and $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$, then $B \rightarrow^* \xrightarrow{\alpha} \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' .

In this case each interaction on one side can be simulated by the other side, and the processes are statically equivalent at each step during the execution, thus an attacker cannot distinguish both sides.

2.2 Modeling Voting Protocols

We model voting protocols in the Applied π -Calculus as follows.

Definition 4 (Voting Protocol). *A voting protocol is a tuple $(V, A_1, \dots, A_m, \tilde{n})$ where V is the process that is executed by the voter, the A_j 's are the processes executed by the election authorities, and \tilde{n} is a set of private channels. We also assume the existence of a particular public channel res that is only used to publish the result of the election.*

Note that we have only one process for the voters. This means that different voters will execute the same process, but with different variable values (e.g. the keys, the vote etc.). To reason about privacy, we talk about instances of a voting protocol, which we call *voting processes*.

Definition 5 (Voting Process). *A voting process of a voting protocol $(V, A_1, \dots, A_m, \tilde{n})$ is a closed process*

$$\nu \tilde{n}'. (V \sigma_{id_1} \sigma_{v_1} | \dots | V \sigma_{id_n} \sigma_{v_n} | A_1 | \dots | A_l),$$

where $l \leq m$, \tilde{n}' includes the secret channel names \tilde{n} , $V \sigma_{id_i} \sigma_{v_i}$ are the processes executed by the voters, σ_{id_i} is a substitution assigning the identity to a process (this determines for example the secret keys), σ_{v_i} specifies the vote and A_j 's are the election authorities which are required to be honest.

The restricted channel names model private channels. Note that we only model the honest authorities as unspecified parties are subsumed by the attacker.

2.3 Existing Privacy Definitions

Before we can formally define privacy, we need the following two transformations. The first one turns a process P into another process P^{ch} that reveals all its inputs and secret data on the channel ch .

Definition 6 (Process P^{ch} [2]). *Let P be a plain process and ch be a channel name. P^{ch} is defined as follows:*

- $0^{ch} \hat{=} 0$,
- $(P|Q)^{ch} \hat{=} P^{ch}|Q^{ch}$,
- $(\nu n.P)^{ch} \hat{=} \nu n.\text{out}(ch, n).P^{ch}$ if n is a name of base type, $(\nu n.P)^{ch} \hat{=} \nu n.P^{ch}$ otherwise,
- $(\text{in}(u, x).P)^{ch} \hat{=} \text{in}(u, x).\text{out}(ch, x).P^{ch}$ if x is a variable of base type, $(\text{in}(u, x).P)^{ch} \hat{=} \text{in}(u, x).P^{ch}$ otherwise,
- $(\text{out}(u, M).P)^{ch} \hat{=} \text{out}(u, M).P^{ch}$,
- $(!P)^{ch} \hat{=} !P^{ch}$,
- $(\text{if } M = N \text{ then } P \text{ else } Q)^{ch} \hat{=} \text{if } M = N \text{ then } P^{ch} \text{ else } Q^{ch}$.

In the remainder we assume that $ch \notin \text{fn}(P) \cup \text{bn}(P)$ before applying the transformation. The second transformation does not only reveal the secret data, but also takes orders from an outsider before sending a message or branching.

Definition 7 (Process P^{c_1, c_2} [2]). *Let P be a plain process and c_1, c_2 be channel names. P^{c_1, c_2} is defined as follows:*

- $0^{c_1, c_2} \hat{=} 0$,
- $(P|Q)^{c_1, c_2} \hat{=} P^{c_1, c_2} | Q^{c_1, c_2}$,
- $(\nu n.P)^{c_1, c_2} \hat{=} \nu n.\text{out}(c_1, n).P^{c_1, c_2}$ if n is a name of base type, $(\nu n.P)^{c_1, c_2} \hat{=} \nu n.P^{c_1, c_2}$ otherwise,
- $(\text{in}(u, x).P)^{c_1, c_2} \hat{=} \text{in}(u, x).\text{out}(c_1, x).P^{c_1, c_2}$ if x is a variable of base type, $(\text{in}(u, x).P)^{c_1, c_2} \hat{=} \text{in}(u, x).P^{c_1, c_2}$ otherwise,
- $(\text{out}(u, M).P)^{c_1, c_2} \hat{=} \text{in}(c_2, x).\text{out}(u, x).P^{c_1, c_2}$ where x is a fresh variable,
- $(!P)^{c_1, c_2} \hat{=} !P^{c_1, c_2}$,
- $(\text{if } M = N \text{ then } P \text{ else } Q)^{c_1, c_2} \hat{=} \text{in}(c_2, x).\text{if } x = \text{true then } P^{c_1, c_2} \text{ else } Q^{c_1, c_2}$ where x is a fresh variable and true is a constant.

To hide the output of a process, we use the following definition.

Definition 8 (Process $A^{\text{out}(ch, \cdot)}$ [2]). Let A be an extended process. We define the process $A^{\text{out}(ch, \cdot)}$ as $\nu ch.(A | \text{in}(ch, x))$.

We now recall the privacy definitions given by Delaune et al. [2], which are the bases for many other definitions [1, 3–5, 8, 12]. Their main idea for defining privacy is simple: A protocol respects privacy if any two instances where two voters swap votes are bisimilar.

Definition 9 (Swap-Privacy (SwP) [2]). A protocol satisfies Swap-Privacy (SwP) if for any context S corresponding to a voting process with a hole for two voters and for all votes σ_{v_A} and σ_{v_B} we have

$$S[V\sigma_{id_A}\sigma_{v_A} | V\sigma_{id_B}\sigma_{v_B}] \approx_l S[V\sigma_{id_A}\sigma_{v_B} | V\sigma_{id_B}\sigma_{v_A}].$$

In the literature S may sometimes contain corrupted or coerced voters (e.g. in [8]), here we will suppose that it contains only honest voters and authorities to be able to clearly distinguish single- and multi-voter coercion.

Defining Receipt-Freeness is a bit more complicated as the voter will execute some counter strategy, i.e. a different process, to fake the receipt, but it can still be expressed as a bisimilarity between two situations. In the first situation, the targeted voter votes a and reveals his secret data. In the second situation, he executes another process – the counter-strategy – which allows him to vote b and fake the secret data in a way that both instances are bisimilar. A protocol is receipt-free if such a process – a counter-strategy – exists.

Definition 10 (Swap-Receipt-Freeness (SwRF) [2]). A protocol satisfies Swap-Receipt-Freeness (SwRF) if for any context S corresponding to a voting process with a hole for two voters and for all votes σ_{v_A} and σ_{v_B} there exists a process V' such that $V'^{\text{out}(ch, \cdot)} \approx_l V\sigma_{id_A}\sigma_{v_B}$ and

$$S[(V\sigma_{id_A}\sigma_{v_A})^{chc} | V\sigma_{id_B}\sigma_{v_B}] \approx_l S[V' | V\sigma_{id_B}\sigma_{v_A}].$$

One could define Coercion-Resistance in the same way, but in that case the attacker could force the targeted voter to vote d (and not a) in the situation where he complies with the instructions. This would make both situations trivially distinguishable by just looking at the result. To prevent this, Delaune et al. [2] use a context C that is required to force the voter to vote a , but can otherwise interact in any way with the voter.

Definition 11 (Swap-Coercion-Resistance (SwCR) [2]). *A protocol satisfies Swap-Coercion-Resistance (SwCR) if for any context S corresponding to a voting process with a hole for two voters and for all votes σ_{v_A} and σ_{v_B} there exists a process V' such that for any context C with $C = \nu_{c_1}.\nu_{c_2}.(-|P)$ and $\tilde{n} \cap fn(C) = \emptyset$, $S[C[(V\sigma_{id_A}\sigma_{v_A})^{c_1,c_2}|V\sigma_{id_B}\sigma_{v_B}]] \approx_l VP'_A[(V\sigma_{id_A}\sigma_{v_A})^{chc}|V\sigma_{id_B}\sigma_{v_B}]$ we have $C[V'] \setminus^{out(chc,\cdot)} \approx_l V\sigma_{id_A}\sigma_{v_B}$ and*

$$S[C[(V\sigma_{id_A}\sigma_{v_A})^{c_1,c_2}|V\sigma_{id_B}\sigma_{v_B}]] \approx_l S[C[V']|V\sigma_{id_B}\sigma_{v_A}].$$

Delaune et al. [2] showed that any protocol ensuring (SwCR) ensures (SwRF), and any protocol ensuring (SwRF) ensures (SwP).

3 Defining Privacy

Our privacy definition is based on the observation that - as the result of the vote is always published - some knowledge about the voter's choices can always be inferred from the outcome. The classical example is the case of a unanimous vote where the contents of all votes are revealed just by the result. Yet - as already discussed in the introduction - there can also be other cases where some of the votes can be inferred from the result, in particular in the case of weighted votes. If for example Alice holds 66% of the shares and Bob 34%, both votes are always revealed when announcing the result: If one option gets 66% and the other 34%, it is clear which one was chosen by Alice or Bob. However, if we have a different distribution of the shares (e.g. 50%, 25% and 25%), some privacy is still possible as there several situations with the same result. Thus our main idea: If two instances of a protocol give the same result, an attacker should not be able to distinguish them. Note that this includes the classic definition where votes are swapped, if this give the same result.

3.1 Formal Definition

To express this formally, we need to define the result of an election. As defined above, we suppose that the result is always published on a special channel *res*. The following definition allows us to hide all channels except for a specified channel c , which we can use for example to reason about the result on channel *res*.

Definition 12 ($P|_c$). *Let $P|_c = \nu\tilde{c}h.P$ where $\tilde{c}h$ are all channels except for c , i.e. we hide all channels except for c .*

Now we can formally define our privacy notion: If two instances of a protocol give the same result, they should be bisimilar.

Definition 13 (Vote-Privacy (VP)). *A voting protocol ensures Vote-Privacy (VP) if for any two instances $VP_A = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^A} | \dots | V\sigma_{id_n}\sigma_{v_n^A} | A_1 | \dots | A_l)$ and $VP_B = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^B} | \dots | V\sigma_{id_n}\sigma_{v_n^B} | A_1 | \dots | A_l)$ we have*

$$VP_A|_{res} \approx_l VP_B|_{res} \Rightarrow VP_A \approx_l VP_B.$$

A simple interpretation of this definition is that everything apart from the result on channel *res* has to remain private. This obviously relies heavily on the notion of “result” and the modeling of the protocol. Typically the result will only contain only the sum of all votes, which corresponds to a simple and intuitive understanding of privacy.

Some protocols may leak some additional information, for example the number of ballots on the bulletin board. For instance in the protocol by Juels et al. [11] voters can post fake ballots. In this case, the above definition of the result may lead to a too restrictive privacy notion, since two situations with the same votes but a different number of fakes are required to be bisimilar. To address this issue, we can include the number of ballots in the result if we want to accept the additional leakage. This gives very fine-grained control about the level of privacy we want to model.

Note that if the link between a voter and his vote is also published as part of the result on channel *res*, our definition of privacy may be true although this probably does not correspond to the intuitive understanding of privacy. This is however coherent within the model since everything apart from the result is private; simply the result itself leaks too much information.

3.2 Link to Existing Definitions

To establish the relationship of our definition and the existing ones, we need to formally characterize their difference. Intuitively the swap-based definition assumes that swapping two votes will not change the result. This can be formalized as follows: If two instances of the protocol with the same voters give the same result, then the votes are a permutation of each other, and vice versa. This precludes weighted votes, thus the name “Equality of Votes”.

Definition 14 (Equality of Votes (EQ)). *A voting protocol respects Equality of Votes (EQ) if for any $VP_A = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^A}|\dots|V\sigma_{id_n}\sigma_{v_n^A}|A_1|\dots|A_l)$ and $VP_B = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^B}|\dots|V\sigma_{id_n}\sigma_{v_n^B}|A_1|\dots|A_l)$ we have*

$$VP_A|_{res} \approx_l VP_B|_{res} \Leftrightarrow \exists\pi : \forall i : \sigma_{v_i^B} = \sigma_{v_{\pi(i)}^A},$$

where π is a permutation.

This allows us to formally prove that our definition is equivalent to the existing ones if (EQ) holds.

Theorem 1 (Equivalence of Privacy Definitions). *If a protocol respects (EQ), then (VP) and (SwP) are equivalent.*

The full formal proof can be found in our technical report [21]. Intuitively, because of (EQ), two instances of a protocol can only have the same result if the votes are a permutation of each other. As any permutation can be written as a sequence of simple permutations (swaps), (SwP) is enough to generate any possible permutation, which gives (VP). Conversely, the definition of (SwP) becomes just a particular case of (VP).

It is easy to see that this condition (EQ) is necessary: If a protocol uses weighted votes (e.g. Alice 66%, Bob 34%), it may satisfy (VP), but not (SwP).

Similarly, consider the following example: In the official result announced on channel *res*, a pre-selected candidate always wins - this could be the case if the authorities are dishonest and want to manipulate the election outcome. If however at the same time the ballots on the bulletin board allow to calculate the result, such a protocol may ensure (SwP) – if the ballots cannot be linked to the voters –, but not (VP) because two instances with a different outcome based on the ballots will have the same “result” on *res*. Note that such a protocol would contradict (EQ) because we have instances where the votes are not a permutation of each other, but still give the same result.

3.3 Example: A Variant of FOO

Eliasson and Zúquete [13] propose an implementation of a voting system supporting vote weights based on REVS [14], which itself is based on the protocol by Fujioka et al. [15], often referred to as “FOO”.

Informal Description: The protocol by Fujioka et al. [15]. is split into three phases. In the first phase, the voter obtains the administrator’s signature on a commitment to his vote:

- Voter V_i chooses his vote v_i and computes a commitment $x_i = \xi(v_i, k_i)$ for a random key k_i .
- He blinds the commitment using a blinding function χ , a random value r_i and obtains $e_i = \chi(x_i, r_i)$.
- He signs e_i and sends the signature $s_i = \sigma_{V_i}(e_i)$ together with e_i and his identity to the administrator A .
- The administrator checks if V_i has the right to vote, has not yet voted, and if the signature s_i is correct. If all tests succeed, he signs $d_i = \sigma_A(e_i)$ and sends it back to V_i .
- V_i unblinds the signature and obtains $y_i = \delta(d_i, r_i)$. He checks the signature.

In the second phase, the actual voting takes place:

- Voter V_i sends (x_i, y_i) to the collector C through an anonymous channel.
- C checks the administrator’s signature and enters (x_i, y_i) into a list.

When all ballots are cast or when the deadline is over, the counting phase begins:

- The collector publishes the list of correct ballots.
- V_i verifies that his commitment appears on the list and sends r_i together with the commitment’s index l on the list to C using an anonymous channel.
- The collector C opens the l -th ballot using r_i and publishes the vote.

Adding Vote Weights: In [13] Eliasson and Zúquete discuss several possibilities on how to implement weights in this protocol:

- including the weight in the vote (which requires trusting the voter for correctness or zero-knowledge proofs to verify the weight)
- using different keys when the vote is signed by the administrator, where each key corresponds to a different weight
- using multiple ballots per voter, i.e. if for example voter A holds 70% and voter B 30% of the shares, voter A sends 7 and voter B 3 ballots.

We implemented the latter variant in the Applied π -Calculus. Using a manual proof (see [21] for details) we can show that

$$VP_A|_{res} \approx_l VP_B|_{res} \Rightarrow \sum_{i=1}^n v_i^A * w_i = \sum_{i=1}^n v_i^B * w_i. \tag{1}$$

Using a python script available on our website [22] that generates all cases to check based on the number of voters and the discrete weight distribution, we can use Proverif to then establish (2) which gives that this variant ensures (VP).

$$\sum_{i=1}^n v_i^A * w_i = \sum_{i=1}^n v_i^B * w_i \Rightarrow VP_A \approx_l VP_B \tag{2}$$

4 Receipt-Freeness

In this section we define receipt-freeness for weighted votes. We first consider the case where only one voter is attacked, then we define multi-voter attacks.

4.1 Single-Voter Receipt-Freeness (SRF)

We combine the idea by Delaune et al. (Def. 10) with our definition of Privacy: If two instances of a voting protocol give the same result, they should be bisimilar even if one voter reveals his secret data in one case or fakes it in the other.

Definition 15 (Single-Voter Receipt Freeness (SRF)). *A voting protocol ensures Single-Voter Receipt Freeness (SRF) if for any voting processes $VP_A = \nu \tilde{n}.(V\sigma_{id_1}\sigma_{v_1^A} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^A} \mid A_1 \mid \dots \mid A_l)$, $VP_B = \nu \tilde{n}.(V\sigma_{id_1}\sigma_{v_1^B} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^B} \mid A_1 \mid \dots \mid A_l)$ and any number $i \in \{1, \dots, n\}$ there exists a process V'_i such that we have $V'_i \setminus^{out}(chc_i, \cdot) \approx_l V\sigma_{id_i}\sigma_{v_i^B}$ and*

$$VP_A|_{res} \approx_l VP_B|_{res} \Rightarrow VP'_A \left[(V\sigma_{id_i}\sigma_{v_i^A})^{chc_i} \right] \approx_l VP'_B [V'_i],$$

where VP'_A and VP'_B are like VP_A and VP_B , but with holes for the voter $V\sigma_{id_i}$.

As for (VP), our definition is equivalent to the existing one based on swapping if the protocol ensures (EQ), which is the case if it does not use weighted votes. Similarly to swap-based definitions, (SRF) is stronger than (VP). The proof is analogous to the proof in the swap-based model (see [21] for details).

4.2 Multi-voter Receipt-Freeness (MRF)

We now generalize the idea of Receipt-Freeness to the case where multiple voters are attacked. Instead of only considering one attacked voter i , we consider a set I of attacked voters. To be receipt-free, it should be possible for all attacked voters to fake the receipt. Note that we assume that there is always at least one honest voter, except for the case with only one voter.

Definition 16 (Multi-Voter Receipt Freeness (MRF)). *A voting protocol ensures Multi-Voter Receipt Freeness (MRF) if for any voting processes $VP_A = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^A} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^A} \mid A_1 \mid \dots \mid A_l)$, $VP_B = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^B} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^B} \mid A_1 \mid \dots \mid A_l)$ and any subset $I \subset \{1, \dots, n\}$, $I \neq \{1, \dots, n\}$ if $n > 1$, then there exists processes V'_i such that we have $\forall i \in I : V_i^{\wedge out(chc, \cdot)} \approx_l V\sigma_{id_i}\sigma_{v_i^B}$ and*

$$VP_A|_{res} \approx_l VP_B|_{res} \Rightarrow VP'_A \left[\begin{array}{c} | \\ \hline \mid \\ \hline \end{array} (V\sigma_{id_i}\sigma_{v_i^A})^{chc_i} \right] \approx_l VP'_B \left[\begin{array}{c} | \\ \hline \mid \\ \hline \end{array} V'_i \right],$$

where VP'_A and VP'_B are like VP_A and VP_B , but with holes for all voters $V\sigma_{id_i}$, $i \in I$.

By choosing $I = \{i\}$ we obtain that (MRF) implies (SRF). Under certain conditions the converse is also true. To prove this, we define a “generalized voting process” which is like a voting process, but some voters might be under attack.

Definition 17 (Generalized Voting Process). *A Generalized Voting Process is a voting process VP with variables for the voter’s processes that can either be a “normal” voter or a voter communicating with the intruder, i.e. $VP = \nu\tilde{n}.(V_1 \mid \dots \mid V_n \mid A_1 \mid \dots \mid A_l)$ where $V_i \approx_l V\sigma_{id_i}\sigma_{v_i}$ or $V_i^{\wedge out(chc_i, \cdot)} \approx_l V\sigma_{id_i}\sigma_{v_i}$.*

The next definition captures the key properties required for our proof. It expresses two modularity conditions of a voting protocol.

Definition 18 (Modularity (Mod)). *A voting protocol is modular (Mod) if it is composable and decomposable. A voting protocol is composable if for any generalized voting processes VP_A and VP_B there exists a generalized voting process VP such that $VP \approx_l VP_A|VP_B$. A voting protocol is decomposable if any generalized voting process $VP = \nu\tilde{n}.(V_1 \mid \dots \mid V_n \mid A_1 \mid \dots \mid A_l)$ can be decomposed into processes $VP_i = \nu\tilde{n}_i.(V_i \mid A_1^i \mid \dots \mid A_l^i)$ where*

$$VP \approx_l VP_1 \mid \dots \mid VP_n. \quad (3)$$

Imagine a protocol where in order to escape coercion the voters can claim that a certain ballot on the bulletin board is their ballot, but it was actually prepared by some honest authority to allow the voters to create a fake receipt. If we suppose that this ballot exists only once no matter how many voters are attacked, it would be enough for a single voter to fake his receipt. However we cannot compose two instances with one attacked voter each, as they would use the same fake ballot which would be noticeable for the attacker. Hence the above definition also

captures the fact that faking the receipt to escape coercion can be done by each voter independently.

Another property we need for our proof is *Correctness*, i.e. the fact that if in two instances the voters' choices are the same, they give the same result¹.

Definition 19 (Correctness (Cor)). *A voting protocol is correct if for any generalized voting processes $VP_A = \nu\tilde{n}_A.(V_{1,A}|\dots|V_{n,A}|A_1|\dots|A_l)$ and $VP_B = \nu\tilde{n}_B.(V_{1,B}|\dots|V_{n,B}|A_1|\dots|A_l)$ with for any i and $X \in \{A, B\}$: $V_{i,X} \setminus^{out(chc_i, \cdot)} \approx_l V\sigma_{id_i}\sigma_{v_i}$, we have*

$$VP_A|_{res} \approx_l VP_B|_{res} \quad (4)$$

It is easy to see that Correctness is implied by Equality of Votes as the identity is a permutation, hence any protocol ensuring (EQ) ensures (Cor) [21]. Putting everything together, we are able to prove the equivalence of (SRF) and (MRF).

Theorem 2. *If a protocol is modular, correct and ensures Single-Voter Receipt Freeness, it also ensures Multi-Voter Receipt Freeness.*

The full proof is given in our technical report [21]. The main idea is that we can decompose an instance with multiple attacked voters into instances with at most one attacked voter, where we can apply the single-voter assumption, and recompose the result. Note that the assumptions (Mod), (EQ) are satisfied by many well-known protocols (e.g. [15, 23, 24]), we illustrate this on an example.

Remark. We have to be careful when modeling protocols using a full PKI. If we model the PKI inside the voting process, decomposing a protocol would result in two instances using different keys, which will most probably be visible to an attacker and the bisimilarity (3) will not hold. A possible solution could be to externalize the PKI into a context K such that $K[VP] \approx_l K[VP_1|VP_2]$, which ensures that VP_1 and VP_2 use the same keys. This would allow us to obtain the same result for protocols such as [25, 26].

4.3 Example: Protocol by Okamoto

The protocol by Okamoto [24] uses trapdoor commitments to achieve (SwRF), but it is not (SwCR) [2].

Informal Description: The protocol is split in 3 phases. In the first phase the voter obtains a signature on a commitment to his vote from the administrator:

- Each voter V_i chooses his vote v_i and computes a trapdoor commitment $x_i = \xi(v_i, k_i, td_i)$ for a random key k_i and a trapdoor td_i .
- V_i blinds the commitment using a blinding function χ , a random value r_i and obtains $e_i = \chi(x_i, r_i)$.

¹ This does not entirely cover intuitive correctness as it will be fulfilled by protocols always giving the same result independently from the votes, but it will fail for a protocol announcing a random result.

5 Coercion-Resistance

After discussing Receipt-Freeness, we now define Coercion-Resistance. As before, we start with Single-Voter Coercion-Resistance.

5.1 Single-Voter Coercion (SCR)

In this case, we combine (VP) with (SwCR): If two instances of a voting protocol give the same result, they should be bisimilar even if one voter interacts with the attacker in one case or only pretends to do so in the other case. The coercion is modeled by the context C that interacts with the voter and tries to force him to vote for a certain candidate.

Definition 20 (Single-Voter Coercion-Resistance (SCR)). *A voting protocol ensures Single-Voter Coercion-Resistance (SCR) if for any voting processes $VP_A = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^A} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^A} \mid A_1 \mid \dots \mid A_l)$, $VP_B = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^B} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^B} \mid A_1 \mid \dots \mid A_l)$ and any number $i \in \{1, \dots, n\}$ there exists a process V'_i such that for any context C_i with $C_i = \nu c_1.\nu c_2.(_ \mid P_i)$ and $\tilde{n} \cap fn(C) = \emptyset$, $VP'_A \left[C_i \left[(V\sigma_{id_i}\sigma_{v_i^A})^{c_1, c_2} \right] \right] \approx_l VP'_A \left[(V\sigma_{id_i}\sigma_{v_i^A})^{chc_i} \right]$ we have $C_i [V'_i]^{out(chc, \cdot)} \approx_l V\sigma_{id_i}\sigma_{v_i^B}$ and*

$$VP_A|_{res} \approx_l VP_B|_{res} \Rightarrow VP'_A \left[C_i \left[(V\sigma_{id_i}\sigma_{v_i^A})^{c_1, c_2} \right] \right] \approx_l VP'_B \left[C_i [V'_i] \right],$$

where VP'_A and VP'_B are like VP_A and VP_B , but with a holes for the voter $V\sigma_{id_i}$.

As above, we can easily link this definition to the existing swap-based definition using (EQ): If a protocol respects (EQ), (SCR) and (SwCR) are equivalent. The proof given in [21] is similar to the (SRF) case.

5.2 Multi-voter Coercion (MCR)

We now discuss Multi-Voter Coercion-Resistance. To model the case where multiple voters are attacked, we consider the set I of attacked voters.

Definition 21 (Multi-Voter Coercion-Resistance (MCR)). *A voting protocol ensures Multi-Voter Coercion-Resistance (MCR) if for any voting processes $VP_A = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^A} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^A} \mid A_1 \mid \dots \mid A_l)$, $VP_B = \nu\tilde{n}.(V\sigma_{id_1}\sigma_{v_1^B} \mid \dots \mid V\sigma_{id_n}\sigma_{v_n^B} \mid A_1 \mid \dots \mid A_l)$ and any subset $I \subset \{1, \dots, n\}$, $I \neq \{1, \dots, n\}$ if $n > 1$, there exists processes V'_i such that for any contexts $C_i, i \in I$ with $C_i = \nu c_1.\nu c_2.(_ \mid P_i)$ and $\tilde{n} \cap fn(C) = \emptyset$, $VP'_A \left[\prod_{i \in I} C_i \left[(V\sigma_{id_i}\sigma_{v_i^A})^{c_1, c_2} \right] \right] \approx_l VP'_A \left[\prod_{i \in I} (V\sigma_{id_i}\sigma_{v_i^A})^{chc_i} \right]$ we have $\forall i \in I : C_i [V'_i]^{out(chc, \cdot)} \approx_l V\sigma_{id_i}\sigma_{v_i^B}$ and*

$$VP_A|_{res} \approx_l VP_B|_{res} \Rightarrow VP'_A \left[\prod_{i \in I} C_i \left[(V\sigma_{id_i}\sigma_{v_i^A})^{c_1, c_2} \right] \right] \approx_l VP'_B \left[\prod_{i \in I} C_i [V'_i] \right],$$

where VP'_A and VP'_B are like VP_A and VP_B , but with holes for all voters $V_{\sigma_{id_i}}, i \in I$.

As for (MRF), (MCR) implies (SCR), and (MCR) resp. (SCR) is stronger than (MRF) resp. (SRF) (for the proofs, see the technical report [21]). We also have equivalence between (SCR) and (MCR) under the same assumptions as in the case of Receipt-Freeness using a similar proof.

5.3 Example: Bingo Voting

Bingo Voting was developed by Bohli et al. [23] to achieve coercion-resistance as well as individual and universal verifiability by using a trusted random number generator (RNG) and a voting booth.

Informal Description: We consider an election with k voters and l candidates. The protocol is split into three phases: The pre-voting phase, the voting phase and the post-voting phase. In the pre-voting phase, the voting machine generates k random values $n_{i,j}$ for every candidate p_j . It commits to the $k \cdot l$ pairs $(n_{i,j}, p_j)$ and publishes the shuffled commitments.

In the voting phase, the voter enters the voting booth and selects the candidate he wants to vote for on the voting machine. The RNG generates a random number r which is transmitted to the voting machine and displayed to the voter. The voting machine chooses for each candidate a dummy vote except for the voter's choice. For this candidate the random value from the RNG is used and the receipt (a list of all candidates and the corresponding dummy or real votes) is created. Finally the voter checks that the number displayed on the RNG corresponds to the entry of his candidate on the receipt.

In the post-voting phase the voting machine announces the result, publishes all receipts and opens the commitments of all unused dummy votes. The machine also generates non-interactive zero-knowledge proofs that each unopened commitment was actually used as a dummy vote in one of the receipts.

Analysis: The protocol satisfies (SwRF) as the receipt contains only random numbers, and it is impossible for the attacker to know which entry corresponds to the random value generated by the RNG [4]. It also ensures (SwCR) as voting takes places inside a secured voting booth. This was formally proven in the DKR-model [4], and we use the same model to show that it satisfies (MCR). As before, it is easy to see that the protocol ensures Equality of Votes and hence Correctness as votes are not weighted. By analyzing the structure of the voting process, we can see that the protocol also is *modular*. In the case of n voters, we have the following voting process

$$\nu\text{privCh}M_1 \dots \nu\text{privCh}M_n. \nu\text{privCh}R_{M_1} \dots \nu\text{privCh}R_{M_n}. \\ \nu\text{privCh}R_1 \dots \nu\text{privCh}R_n. (V_1 | \dots | V_n | M_{1,\dots,n;l} | R_1 | \dots | R_n)$$

where R_i are the trusted random number generators, $M_{1,\dots,n;l}$ is the voting machine process for n voters from 1 to n and l candidates, and $\text{privCh}M_i$,




Fig. 2. Relations among the notions. $A \xrightarrow{C} B$ means that under the assumption C a protocol ensuring A also ensures B.

privChRM_i and privChR_i are the private channels between the voter and the voting machine, the RNG and the voting machine, and the RNG and the voter respectively. For $k \in \{1, \dots, n - 1\}$, this can be rewritten as

$$\begin{aligned} & \nu\text{privChM}_1 \dots \nu\text{privChM}_k. \nu\text{privChRM}_1 \dots \nu\text{privChRM}_k. \\ & \nu\text{privChR}_1 \dots \nu\text{privChR}_k. (V_1 | \dots | V_k | M_{1, \dots, k; l} | R_1 | \dots | R_k) | \\ & \nu\text{privChM}_{k+1} \dots \nu\text{privChM}_n. \nu\text{privChRM}_{k+1} \dots \nu\text{privChRM}_n. \\ & \nu\text{privChR}_{k+1} \dots \nu\text{privChR}_n. (V_{k+1} | \dots | V_n | M_{k+1, \dots, n; l} | R_{k+1} | \dots | R_n) \end{aligned}$$

as $M_{1, \dots, n; l} \approx_l M_{1, \dots, k; l} | M_{k+1, \dots, n; l}$. This can be easily seen from the applied π -code [21]. It is easy to see that this also works for composing processes. Hence we have all necessary conditions and obtain that Bingo Voting ensures (MCR).

6 Conclusion

We presented an intuitive definition of privacy for voting protocols that generalizes to situations with weighted votes. We extended the definition to include Receipt-Freeness and Coercion-Resistance as well. We considered situations where only one voter is under attack, and others where multiple voters are attacked. We were able to show that - under the assumptions that votes are not weighted and correctly counted - the single voter case is equivalent to (SwP), (SwRF), (SwCR) as defined by Delaune et al. [2]. Moreover, we proved that the multi-voter case is equivalent to the single-voter case if the protocol is correct (Cor) and respects a modularity condition (Mod). This condition allows us to compose and decompose protocols, which expresses the fact the different parts of the protocol are independent. Figure 2 summarizes our results. Finally, we illustrated our work by analyzing two existing protocols. As future work, we would like to translate these symbolic definitions to the computational setting.

Acknowledgments. This work was partly supported by the ANR project ProSe (decision ANR 2010-VERS-004).

References

1. Backes, M., Hritcu, C., Maffei, M.: Automated verification of remote electronic voting protocols in the applied pi-calculus. In: IEEE Computer Security Foundations Symposium, pp. 195–209 (2008)
2. Delaune, S., Kremer, S., Ryan, M.: Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security* 17, 435–487 (2009)
3. Delaune, S., Kremer, S., Ryan, M.D.: Verifying Privacy-Type Properties of Electronic Voting Protocols: A Taster. In: Chaum, D., Jakobsson, M., Rivest, R.L., Ryan, P.Y.A., Benaloh, J., Kutylowski, M., Adida, B. (eds.) *Towards Trustworthy Elections*. LNCS, vol. 6000, pp. 289–309. Springer, Heidelberg (2010)
4. Dreier, J., Lafourcade, P., Lakhnech, Y.: Vote-Independence: A Powerful Privacy Notion for Voting Protocols. In: Garcia-Alfaro, J., Lafourcade, P. (eds.) *FPS 2011*. LNCS, vol. 6888, pp. 164–180. Springer, Heidelberg (2012)
5. Dreier, J., Lafourcade, P., Lakhnech, Y.: A formal taxonomy of privacy in voting protocols. In: *First IEEE International Workshop on Security and Forensics in Communication Systems (ICC 2012 WS - SFCS)* (2012)
6. Küsters, R., Truderung, T.: An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In: *2009 IEEE Symposium on Security and Privacy (S&P 2009)*, pp. 251–266. IEEE Computer Society (2009)
7. Moran, T., Naor, M.: Receipt-Free Universally-Verifiable Voting with Everlasting Privacy. In: Dwork, C. (ed.) *CRYPTO 2006*. LNCS, vol. 4117, pp. 373–392. Springer, Heidelberg (2006)
8. Smyth, B., Cortier, V.: Attacking and fixing helios: An analysis of ballot secrecy. In: *Proceedings of the 24th IEEE Computer Security Foundations Symposium (CSF 2011)*, pp. 297–311. IEEE (2011)
9. Kremer, S., Ryan, M., Smyth, B.: Election Verifiability in Electronic Voting Protocols. In: Gritzalis, D., Preneel, B., Theoharidou, M. (eds.) *ESORICS 2010*. LNCS, vol. 6345, pp. 389–404. Springer, Heidelberg (2010)
10. Smyth, B., Ryan, M.D., Kremer, S., Kourjeh, M.: Towards Automatic Analysis of Election Verifiability Properties. In: Armando, A., Lowe, G. (eds.) *ARSPA-WITS 2010*. LNCS, vol. 6186, pp. 146–163. Springer, Heidelberg (2010)
11. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*. WPES 2005, pp. 61–70. ACM (2005)
12. Kremer, S., Ryan, M.: Analysis of an Electronic Voting Protocol in the Applied Pi Calculus. In: Sagiv, M. (ed.) *ESOP 2005*. LNCS, vol. 3444, pp. 186–200. Springer, Heidelberg (2005)
13. Eliasson, C., Zúquete, A.: An electronic voting system supporting vote weights. *Internet Research* 16(5), 507–518 (2006)
14. Joaquim, R., Zúquete, A., Ferreira, P.: Revs - a robust electronic voting system. In: *IADIS International Conference e-Society 2003*, Lisboa, Portugal, June 3-6 (2003)
15. Fujioka, A., Okamoto, T., Ohta, K.: A Practical Secret Voting Scheme for Large Scale Elections. In: Zheng, Y., Seberry, J. (eds.) *AUSCRYPT 1992*. LNCS, vol. 718, pp. 244–251. Springer, Heidelberg (1993)
16. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2001*, pp. 104–115. ACM, New York (2001)
17. Blanchet, B., Abadi, M., Fournet, C.: Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming* 75(1), 3–51 (2008)

18. Klus, P., Smyth, B., Ryan, M.D.: Proswapper: Improved equivalence verifier for proverif (2010), <http://www.bensmyth.com/proswapper.php>
19. Küsters, R., Truderung, T., Vogt, A.: A game-based definition of coercion-resistance and its applications. In: Proceedings of the 2010 23rd IEEE Computer Security Foundations Symposium. CSF 2010, pp. 122–136. IEEE Computer Society, Washington, DC (2010)
20. Langer, L., Jonker, H., Pieters, W.: Anonymity and Verifiability in Voting: Understanding (Un)Linkability. In: Soriano, M., Qing, S., López, J. (eds.) ICICS 2010. LNCS, vol. 6476, pp. 296–310. Springer, Heidelberg (2010)
21. Dreier, J., Lafourcade, P., Lakhnech, Y.: On defining privacy in the presence of weighted votes and the equivalence of single and multi-voter coercion. Technical Report TR-2012-2, Verimag Research Report (March 2012), <http://www-verimag.imag.fr/TR/TR-2012-2.pdf>
22. Dreier, J.: The code and scripts used to automatically verify the examples (2011), <http://www-verimag.imag.fr/~dreier/papers/foo-weighted-code.zip>
23. Bohli, J.M., Müller-Quade, J., Röhrich, S.: Bingo Voting: Secure and Coercion-Free Voting Using a Trusted Random Number Generator. In: Alkassar, A., Volkamer, M. (eds.) VOTE-ID 2007. LNCS, vol. 4896, pp. 111–124. Springer, Heidelberg (2007)
24. Okamoto, T.: An electronic voting scheme. In: Proceedings of the IFIP World Conference on IT Tools, pp. 21–30 (1996)
25. Lee, B., Boyd, C., Dawson, E., Kim, K., Yang, J., Yoo, S.: Providing Receipt-Freeness in Mixnet-Based Voting Protocols. In: Lim, J.-I., Lee, D.-H. (eds.) ICISC 2003. LNCS, vol. 2971, pp. 245–258. Springer, Heidelberg (2004)
26. Wen, R., Buckland, R.: Masked Ballot Voting for Receipt-Free Online Elections. In: Ryan, P.Y.A., Schoenmakers, B. (eds.) VOTE-ID 2009. LNCS, vol. 5767, pp. 18–36. Springer, Heidelberg (2009)