

# Model Checkers Are Cool: How to Model Check Voting Protocols in UPPAAL

Wojciech Jamroga<sup>1,2</sup>, Yan Kim<sup>1</sup>, Damian Kurpiewski<sup>2</sup>, and Peter Y. A. Ryan<sup>1</sup>

<sup>1</sup> SnT, University of Luxembourg

<sup>2</sup> Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

**Abstract** The design and implementation of an e-voting system is a challenging task. Formal analysis can be of great help here. In particular, it can lead to a better understanding of how the voting system works, and what requirements on the system are relevant. In this paper, we propose that the state-of-art model checker UPPAAL provides a good environment for modelling and preliminary verification of voting protocols. To illustrate this, we present an UPPAAL model of Prêt à Voter, together with some natural extensions. We also show how to verify a variant of receipt-freeness, despite the severe limitations of the property specification language in the model checker.

## 1 Introduction

The design and implementation of a good e-voting system is highly challenging. Real-life systems are notoriously complex and difficult to analyze. Moreover, elections are *social* processes: they are run by humans, with humans, and for humans, which makes them unpredictable and hard to model. Last but not least, it is not always clear what *good* means for a voting system. A multitude of properties have been proposed by the communities of social choice theory (such as Pareto optimality and nonmanipulability), as well as secure voting (ballot secrecy, coercion-resistance, voter-verifiability, etc.). The former are typically set for a very abstract view of the voting procedure, and consequently miss many real-life concerns. For the latter ones, it is often difficult to translate the informal intuition to a formal definition that is commonly accepted.

In a word, we deal with processes that are hard to understand and predict, and evaluate them against criteria that are not exactly clear. Formal analysis can be of great help here: perhaps not in the sense of providing the ultimate answers, but rather to strengthen our understanding of both how the voting system works and how it should work. The main goal of this paper is to propose that model checkers from distributed and multi-agent systems can be invaluable tools for such an analysis.

**Model checkers and UPPAAL.** Much research on model checking focuses on the design of logical systems for a particular class of properties, establishing their theoretical characteristics, and development of verification algorithms. This obscures the fact that a model checking framework is valuable as long as it is

actually *used* to analyze something. The analysis does not necessarily have to result in a “correctness certificate”. A readable model of the system, and an understandable formula capturing the requirement are already of substantial value.

In this context, two features of a model checker are essential. On the one hand, it should provide a *flexible model specification language* that allows for modular and succinct specification of processes. On the other hand, it must offer a *good graphical user interface*. Paradoxically, tools satisfying both criteria are rather scarce. Here, we suggest that the state of the art model checker UPPAAL can provide a nice environment for modelling and preliminary verification of voting protocols and their social context. To this end, we show how to use UPPAAL to model a voting protocol of choice (in our case, Prêt à Voter), and to verify some requirements written in the temporal logic **CTL**.

**Contribution.** The main contribution of this paper is methodological: we demonstrate that specification frameworks and tools from distributed and multi-agent systems can be useful in analysis and validation of voting procedures. An additional, technical contribution consists in a reduction from model checking of temporal-epistemic specifications to purely temporal ones, in order to verify a variant of receipt-freeness despite the limitations of UPPAAL.

**Structure of the paper.** We begin by introducing the main idea behind modelling and model checking of multi-agent systems, including a brief introduction to UPPAAL (Section 2). In Section 3, we provide an overview of Prêt à Voter, the voting protocol that we will use for our study. Section 4 presents a multi-agent model of the protocol; some interesting extensions of the model are proposed in Section 6. We show how to specify simple requirements on the voting system, and discuss the output of model checking in Section 5. The section also presents our main technical contribution, namely the model checking reduction that recasts knowledge-related statements as temporal properties. We discuss related work in Section 7, and conclude in Section 8.

## 2 Towards Model Checking of Voting Protocols

Model checking is the decision problem that takes a model of the system and a formula specifying correctness, and determines whether the model satisfies the formula. This allows for a natural separation of concerns: the model specifies how the system is, while the formula specifies how it should be. Moreover, most model checking approaches encourage systematic specification of requirements, especially for the requirements written in modal and temporal logic. In that case, the behavior of the system is represented by a transition network, possibly with additional modal relations to capture e.g. the uncertainty of agents. The structure of the network is typically given by a higher-level representation, e.g., a set of agent templates together with a synchronization mechanism.

We begin with a brief overview of UPPAAL, the model checker that we will use in later sections. A more detailed introduction can be found in [3].

## 2.1 Modeling in UPPAAL

The system model in UPPAAL consists of a set of concurrent processes. The processes are defined by templates, each possibly having a set of parameters. The parameterized templates are used for defining a large number of almost identical processes. Every template consists of *locations*, *edges*, and optional local declarations. An example template is shown in Figure 2; we will use it to model the behavior of a voter.

Locations are graphically represented as circles. *Initial* locations are marked by a double circle. *Committed* locations are marked by the circled letter C. If any process is in a committed location, then the next transition must involve an edge from one of the committed locations. Those are used for creating atomic sequences or encoding synchronization between more than two components.

The edges are annotated by selections (in yellow), guards (green), synchronizations (teal), and updates (blue). The syntax of expressions mostly coincides with that of C/C++. *Selections* bind the identifier to a value from the given range in a nondeterministic way. *Guards* enable the transition if and only if the guard condition evaluates to true. *Synchronizations* allow processes to synchronize over a common channel *ch* (labeled *ch?* in the receiver process and *ch!* for the sender). Note that a transition on the side of the sender can be fired only if there exists an enabled transition on the receiving side labeled with the same channel identifier, and vice versa. *Update* expressions are evaluated when the transition is taken. For a synchronizing transition, the update expressions on the sender side are executed before the receiver ones. Straightforward value passing over a channel is not allowed; instead, one has to use shared global variables for the transmission.

For convenience, we will place the selections and guards at the top or left of an edge, and the synchronizations and updates at the bottom/right.

## 2.2 Specification of Requirements

To specify requirements, UPPAAL uses a fragment of the temporal logic **CTL** [10]. **CTL** allows for reasoning about the possible execution paths of the system by means of the *path quantifiers* E (“there is a path”) and A (“for every path”). A path is a maximal<sup>1</sup> sequence of states and transitions. To address the temporal pattern on a path, one can use the *temporal operators*  $\bigcirc$  (“in the next moment”),  $\square$  (“always from now on”),  $\diamond$  (“now or sometime in the future”), and U (“until”). For example, the formula  $A\square(\text{has\_ballot}_i \rightarrow A\diamond(\text{voted}_{i,1} \vee \dots \vee \text{voted}_{i,k}))$  expresses that, on all paths, whenever voter *i* gets her ballot form, she will eventually cast her vote for one of the candidates  $1, \dots, k$ . Another formula,  $A\square\neg\text{punished}_i$  says that voter *i* will never be punished by the coercer.

More advanced properties usually require a combination of temporal modalities with *knowledge operators*  $K_a$ , where  $K_a\phi$  expresses “agent *a* knows that  $\phi$  holds.” For example, formula  $E\diamond(\text{results} \wedge \neg\text{voted}_{i,j} \wedge \neg K_c\neg\text{voted}_{i,j})$  says that the

---

<sup>1</sup> I.e., infinite or ending in a state with no outgoing transitions.

(a)	Obelix	
	Idefix	
	Asterix	
	Panoramix	
		7304944

X
7304944

Figure 1: (a) Prêt à Voter ballot form; (b) Receipt encoding a vote for “Idefix”

coercer  $c$  might not know that voter  $i$  hasn’t voted for candidate  $j$ , even if the results are already published. Moreover,  $\text{A}\Box(\text{results} \rightarrow \neg K_c \neg \text{voted}_{i,j})$  expresses that, when the results are out, the coercer won’t know that the voter refused to vote for  $j$ . Intuitively, both formulas capture different strength of receipt-freeness for a voter who has been instructed to vote for candidate  $j$ .

### 3 Outline of Prêt à Voter

In this paper, we use UPPAAL for modeling and analysis of a voting protocol. The protocol of choice is Prêt à Voter. A short overview of Prêt à Voter is presented here; the full details can be found, for example, in [24].

Most verifiable voting systems follow the pattern: at the time of casting, an encryption or encoding of the vote is created and posted to a public bulletin board. The voter can later check that her encrypted ballot appears correctly. The set of posted, ballots are then processed in some verifiable way to reveal the tally or outcome. Much of this is effectively a secure distributed computation, and as such is well-established and understood cryptography. The really challenging bit, because it involves interactions between the users and the system, is the creation of the encrypted ballots. This has to be in a way that assures the voter that her vote is correctly embedded, while avoiding any coercion or vote buying threats.

The key innovation of the Prêt à Voter approach is to encode the vote using a randomised candidate list. This contrasts with earlier verifiable schemes that involved the voter inputting her selection to a device that then produces an encryption of the selection. Here what is encrypted is the candidate order which can be generated and committed in advance, and the voter simply marks her choice on the ballot in the traditional manner.

Suppose that our voter is called Anne. At the polling station, Anne registers and chooses at random a ballot form sealed in an envelope and saunters over to the booth. An example of such a form is shown in Figure 1a. In the booth, Anne extracts her ballot form from the envelope and makes her selection in the usual way by placing a cross in the right hand column against the candidate of her choice (for approval or ranked voting, she marks her selection or ranking against the candidates). Once her selection has been made, she separates the left and right hand strips and discards the left hand strip. She is left with the right hand strip which now constitutes her *privacy protected receipt*, as shown in Figure 1b.

Anne now exits the booth clutching her receipt, returns to the registration desk, and casts her receipt. Her receipt is placed over an optical reader or similar device that records the random value at the bottom of the strip and records in which cell her X is marked. Her original, paper receipt is digitally signed and franked and returned to her to keep and later check that her vote is correctly recorded on the bulletin board. The randomisation of the candidate list on each ballot form ensures that the receipt does not reveal the way she voted, so ensuring the secrecy of her vote. Incidentally, it also removes any bias towards the candidate at the top of the list that can occur with a fixed ordering.

The value printed on the bottom of the receipt is what enables extraction of the vote during the tabulation phase: buried cryptographically in this value is the information needed to reconstruct the candidate order and so extract the vote encoded on the receipt. This information is encrypted with secret keys shared across a number of tellers. Thus, only a threshold set of tellers acting together are able to interpret the vote encoded in the receipt. In practice, the value on the receipt will be the output of an agreed cryptographic hash of the ciphertext committed to the bulletin board.

After the election, voters can visit the secure Web Bulletin Board (WBB) and confirm that their receipts appear correctly. Once any discrepancies are resolved, the tellers take over and perform anonymising mixes and decryption of the receipts. At the end, the plaintext votes will be posted in secret shuffled order, or in the case of homomorphic tabulation, the final result is posted. All the processing of the votes can be made universally verifiable, i.e., any observer can check that no votes were manipulated. These are carefully designed so as not to impinge on ballot privacy.

Prêt à Voter brings several advantages in terms of privacy and dispute resolution. Firstly, avoidance of side channel leakage of the vote from the encryption device. Secondly, improved dispute resolution: ballot assurance is based on random audits of the ballot forms, which can be performed by the voter or independent observers. A ballot form is either well-formed, i.e. the plaintext order matches the encrypted order, or not. This is independent of the voter or her choice, hence there can be no dispute as to what choice the voter provided. Such disputes can arise in Benaloh challenges and similar cut-and-choose style audits. Furthermore, auditing ballots does not impinge on ballot privacy, as nothing about the voter or the vote can be revealed at this point.

## 4 Modelling Prêt à Voter in UPPAAL

In this section, we present how the components and participants of Prêt à Voter can be modelled in UPPAAL. To this end, we give a description of each module template, its elements, and their interactions. The templates represent the behavior of the following types of agents: *voters*, *coercers*, *mix tellers*, *decryption tellers*, *auditors*, and the *voting infrastructure*. For more than one module of a given type, an identifier  $i = 0, 1, \dots$  will be associated with each instance.

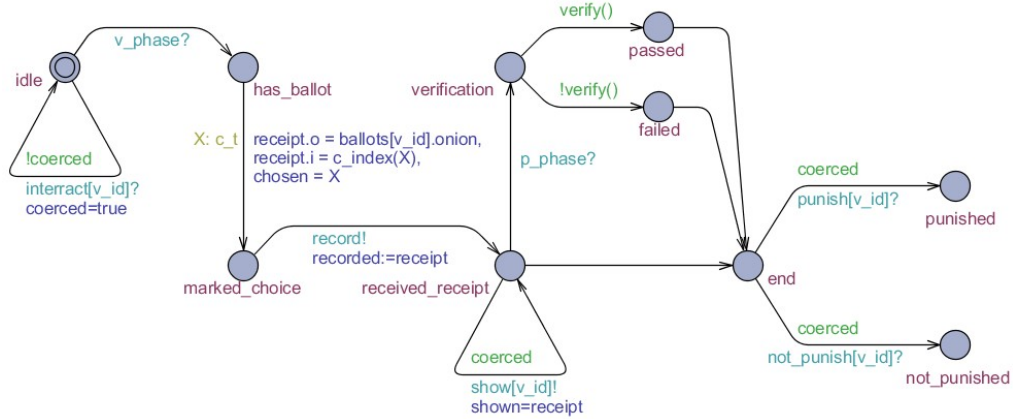


Figure 2: Voter template for the model of Prêt à Voter

The code of the model is available at <https://github.com/pretvsuppaal/model>. Here, we present in detail only the Voter template. The details of the other modules can be found in the extended version of the paper, available at <https://github.com/pretvsuppaal/extended>.

To facilitate readability and manageability of the model code, we define some data structures and type name aliases based on the configuration variables:

- **Ciphertext**: a pair  $(y_1, y_2)$ , representing some cipher-text.
- **Ballot**: a pair  $(\theta, cl)$  of onion  $\theta = E_{PK}(s, *)$  and candidate list  $cl = \pi(s)$ , where  $s$  is a seed associated with the ballot,  $\pi : \mathbb{R} \rightarrow Perm_C$  is a function that associates a seed with a permutation of the candidates. To implement basic absorption of marked cell index into the onion, we will only use cyclic shifts of base candidate order and use a Voter's id as a seed of associated with her ballot.
- **Receipt**: a pair  $(\theta, r)$  of onion  $\theta$  and an index  $r$  of marked cell. It can be used to verify if a term was recorded and if it was done correctly.
- **c\_t**: an integer with range  $[0, c\_total)$ , a candidate;
- **v\_t**: an integer with range  $[0, v\_total)$ , a voter;
- **z\_t**: an integer with range  $[0, z\_total)$ , an element of  $\mathbb{Z}_p^*$ .

#### 4.1 Voter Template

The structure of the Voter template is shown in Figure 2. The idea is that while the voter waits for the start of election she might be subject to coercion. When the ballots are ready, the voter selects a candidate, and transmits the receipt to the system. Then she decides if she wants to check how her vote has been recorded, and if she wants to show the receipt to the coercer. If coerced, she also waits for the Coercer's decision to punish her or refrain from punishment.

The module includes the following private variables:

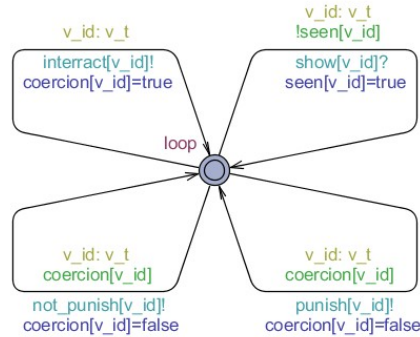


Figure 3: Coercer template

- **receipt**: instance of **Receipt**, obtained after casting a vote;
- **coerced**[=**false**]: a Boolean, indicating if coercer has established a contact;
- **chosen** - integer value of chosen candidate.

Moreover, the following procedures are included:

- **c\_index(target)**: returns an index, at which **target** can be found on the candidate list of a ballot;
- **verify()**: returns *true* if voter’s **receipt** can be found on the board list, else returns *false*.

States:

- *idle*: waiting for the election, might get contacted by coercer;
- *has\_ballot*: obtained the ballot form;
- *marked\_choice*: marked an index of chosen candidate (and destroyed left hand side with candidate list);
- *received\_receipt*: receipt is obtained and might be shown to the coercer;
- *verification*: decided to verify the receipt;
- *passed*: has a confirmation that the receipt appears correctly;
- *failed*: has an evidence that the receipt either does not appear or appears incorrectly (in case of index absorption both cases are the same);
- *end*: the end of the voting ceremony;
- *punished*: has been punished by the Coercer;
- *not\_punished*: has not been punished by the Coercer.

Transitions:

- *idle*→*idle*: if was not already coerced, enable transition; if taken, then set **coercion** to *true*;
- *idle*→*has\_ballot*: enabled transition; if taken, then voter acquired a ballot form;

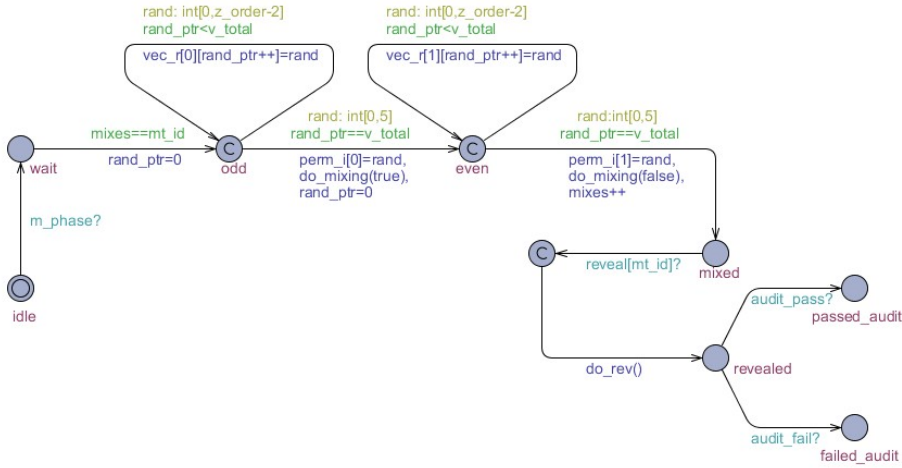


Figure 4: Mteller template

- *has\_ballot*→*marked\_choice*: mark a cell with chosen candidate;
- *marked\_choice*→*received\_receipt*: send **receipt** to the System over channel **record** using shared variable **recorded**;
- *received\_receipt*→*received\_receipt*: if was coerced, enable transition; if taken, then pass the **receipt** to Coercer using shared variable **shown**;
- *received\_receipt*→*verification*: enabled transition; if taken, then Voter decided to verify whether receipt appears on board;
- (*received\_receipt* || *passed* || *failed*)→*end*: end of voting ceremony for voter;
- *end*→*punished*: if was coerced, enable transition; if taken, then Voter was punished by Coercer;
- *end*→*not\_punished*: if was coerced, enable transition; if taken, then Voter was not punished by Coercer.

## 4.2 Coercer

The coercer can be thought of as a party that can influence voters by forcing them to obey certain instructions. To enforce this, the Coercer can punish the voter. The structure of the Coercer is presented in Figure 3; see the extended version of the paper at <https://github.com/pretvsuppaal/extended> for the technical details.

## 4.3 Mix Teller (Mteller)

Once the mixing phase starts, each mix teller performs two re-encryption mixes. The order of turns is ascending and determined by their identifiers. The randomization factors and permutation of each mix are selected in a nondeterministic way and stored for a possible audit of re-encryption mixes. When audited, the



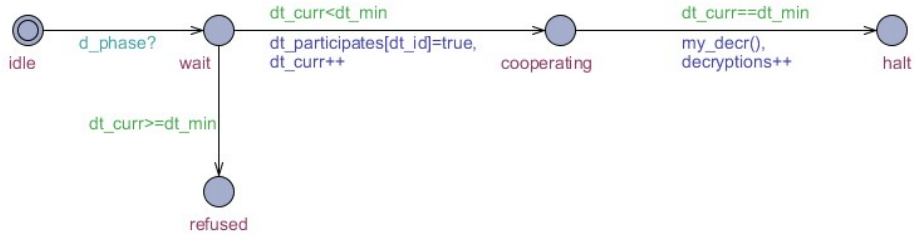


Figure 5: Dteller template

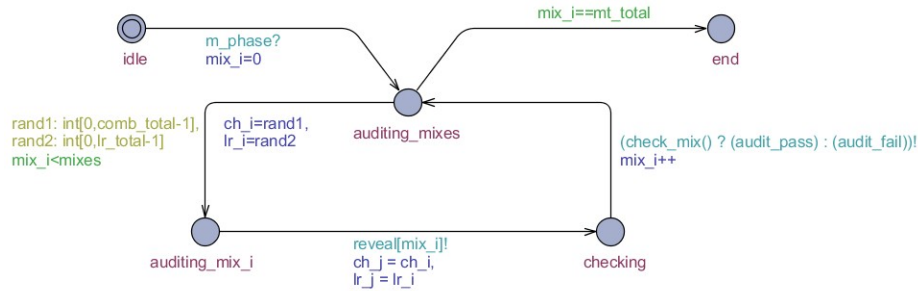


Figure 6: Auditor template

mix teller reveals the requested links and the associated factors, thus allowing Auditor to verify that the input ciphertext maps to the output. The structure of the mix teller is shown in Figure 4.

#### 4.4 Decryption Teller (Dteller)

In this module, after the re-encryption mixes are done, a decryption teller non-deterministically chooses a subset of tellers. Note that it is possible that two or more of decryption tellers refuse to cooperate, which would lead to a deadlock. In order to avoid that, only subsets with cardinality of 2 will be considered.

#### 4.5 Auditor

In order to confirm that the mix tellers performed their actions correctly, the Auditor conducts an audit based on the randomized partial checking technique, RPC in short [14]. To this end, he requests each mix teller to reveal the factors for the selected half of an odd-mix batch, and verify whether the input corresponds to the output. The control flow of the Auditor is presented in Figure 6.

#### 4.6 Voting Infrastructure Module (Sys)

This module represents the behavior of the election authority that prepares the ballot forms, monitors the current phase, signals the progress of the voting

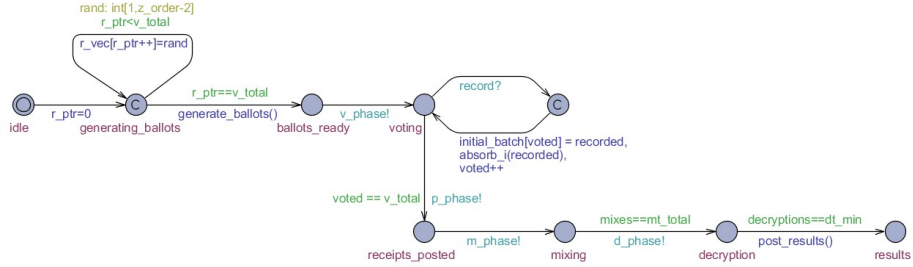


Figure 7: Sys Module

procedure to the other components, and posts the results of the election at the end. In addition, the module plays the role of a server that receives receipts and transfers them to the database throughout the election.

## 5 Verification

As we already mentioned, we chose UPPAAL for this study mainly because of its modeling functionality. The requirement specification capabilities of the tool were of secondary importance. In fact, they are quite limited. First, UPPAAL admits only a fragment of **CTL**: it excludes the “next” and “until” modalities, and does not allow for nesting of operators (with one exception that we describe below). Thus, the supported properties fall into the following categories: simple *reachability* ( $E\Diamond p$ ), *liveness* ( $A\Diamond p$ ), and *safety* ( $A\Box p$  and  $E\Box p$ ). The only allowed nested formulas come in the form of the *p leads to q* property, written  $p \rightsquigarrow q$ , and being a shorthand for  $A\Box(p \rightarrow A\Diamond q)$ .

Still, UPPAAL allows to model-check simple properties of Prêt à Voter, as we show in Section 5.1. Moreover, by tweaking models and formulas in a clever way, one can also verify some more sophisticated requirements, see Section 5.2.

### 5.1 Model Checking Temporal Requirements

The model in Section 4 allows us to verify properties of Prêt à Voter for different configurations of participants. In particular, we have analyzed variants of the model, generated by different numbers of instances for the Voter template. The values should be chosen carefully, to avoid state-space explosion. We have considered the following properties for verification:

1.  $E\Diamond \text{failed\_audit}_0$ : the first mix teller might eventually fail an audit.
2.  $A\Box \neg \text{punished}_i$ : voter  $i$  will never be punished by the coercer.
3.  $\text{has\_ballot}_i \rightsquigarrow \text{marked\_choice}_i$ : on all paths, whenever voter  $i$  gets a ballot form, she will eventually mark her choice.

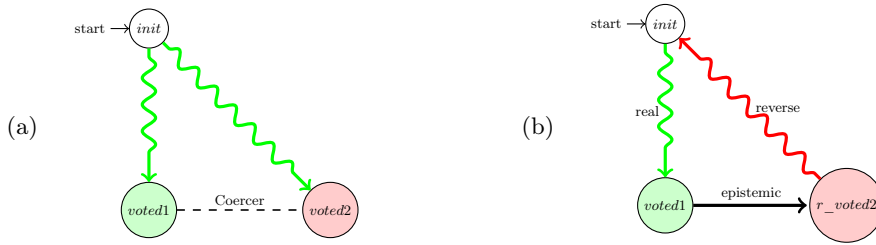


Figure 8: (a) Epistemic bisimulation triangle; (b) turning the triangle into a cycle by reversing the transition relation

We verified each formula on several configurations differing by the number of voters ranging from 1 to 5. For the first property, the UPPAAL verifier returns ‘Property is satisfied’ for the configurations with 1, 2, 3 and 4 voters. In case of 5 voters, we get ‘Out of memory.’ Formula (2) produces the answer ‘Property is not satisfied’ and pastes a counter-example into the simulator (if the diagnostic trace was set in the option panel) for all the five configurations. Finally, formula (3) ends with ‘Out of memory’ regardless of the number of voters.

## 5.2 How to Make Model Checker Do More Than It Is Supposed To

Many important properties of voting refer to the knowledge of its participants. For example, receipt-freeness expresses that the coercer should never know how the voter has voted. Or, better still, that the coercer will never know if the voter voted as instructed. Similarly, voter-verifiability says that the voter will eventually know whether her vote has been registered and tallied correctly (assuming that she follows the verification steps).

A clear disadvantage of UPPAAL is that its language for specification of requirements is restricted to purely temporal properties. We show that, with some care, one can use it to embed the verification of more sophisticated properties. In particular, we show how to enable verification of some knowledge-related requirements by a technical reconstruction of models and formulas. The construction has been inspired by the reduction of epistemic properties to temporal properties, proposed in [12,15]. Consequently, UPPAAL and similar tools can be used to model check formulas of **CTLK** (i.e., **CTL** + Knowledge) that express variants of receipt-freeness and voter-verifiability.

In order to simulate the knowledge operator  $K_a$  under the **CTL** semantics, the model needs to be modified. The first step is to understand how the formula  $\neg K_c \neg \text{voted}_{i,j}$  (saying that the coercer doesn’t know that the voter hasn’t voted for candidate  $j$ ) is interpreted. Namely, if there is a reachable state in which  $\text{voted}_{i,j}$  is true, there must also exist another reachable state, which is indistinguishable from the current one, and in which  $\neg \text{voted}_{i,j}$  holds. The idea is shown in Figure 8a. We observe that to simulate the epistemic relation we need

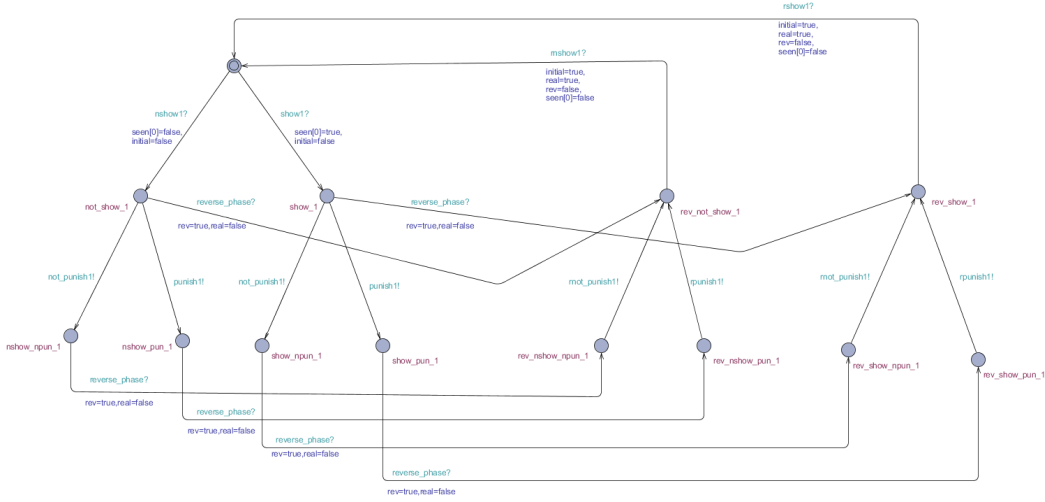


Figure 9: Coercer module augmented with the converse transition relation

to create copies of the states in the model (the “real” states). We will refer to those copies as the *reverse states*. They are the same as the real states, but with reversed transition relation. Then, we add transitions from the real states to their corresponding reverse states, that simulate the epistemic relation between the states. This is shown in Figure 8b.

To illustrate how the reconstruction of the model works on a concrete example, we depict the augmented Coercer template in Figure 9.

The next step is the reconstruction of formulas. Let us take the formula for the weak variant of receipt-freeness from Section 2.2, i.e.,  $E\Diamond(\text{results} \wedge \neg \text{voted}_{i,j} \wedge \neg K_c \neg \text{voted}_{i,j})$ . In order to verify the formula in UPPAAL, we need to replace the knowledge operator according to our model reconstruction method (see Figure 8 again). This means that the verifier should find a path that closes the cycle: from the initial state, going through the real states of the voting procedure to the vote publication phase, and then back to the initial state through the reversed states. In order to “remember” the relevant facts along the path, we use persistent Boolean variables  $\text{voted}_{i,j}$  and  $\text{negvoted}_{i,j}$ : once set to true they always remain true. We also introduce a new persistent variable  $\text{epist\_voted}_{i,j}$  to refer to the value of the vote after an epistemic transition. Once we have all that, we can propose the reconstructed formula:  $E\Diamond(\text{results} \wedge \text{negvoted}_{i,j} \wedge \text{epist\_voted}_{i,j} \wedge \text{initial})$ . UPPAAL reports that the formula holds in the model.

A stronger variant of receipt-freeness is expressed by another formula from Section 2.2, i.e.,  $A\Box(\text{results} \rightarrow \neg K_c \neg \text{voted}_{i,j})$ . Again, the formula needs to be rewritten to a pure **CTL** formula. As before, the model checker should find a cycle from the initial state, “scoring” the relevant propositions on the way. More precisely, it needs to check if, for every real state in which election has ended,

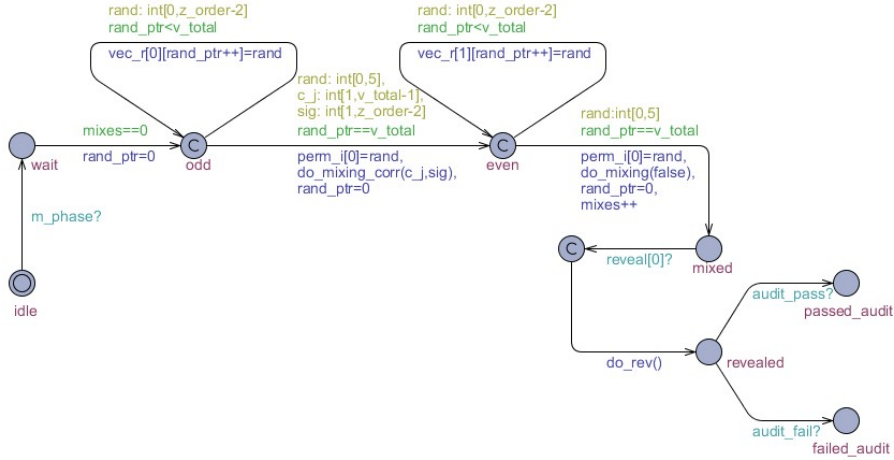


Figure 10: Corrupted Mix Teller module

there exist a path going back to the initial state through a reverse state in which the voter has voted for the selected candidate. This can be captured by the following formula:  $A\Box((\text{results} \wedge \text{real}) \rightarrow E\Diamond(\text{voted}_{i,j} \wedge \text{init}))$ . Unfortunately, this formula cannot be verified in UPPAAL, as UPPAAL does not allow for nested path quantifiers. In the future, we plan to run the verification of this formula using another model checker LTSmin [17] that accepts UPPAAL models as input, but allows for more expressive requirement specifications.

## 6 Replicating Pfitzmann’s Attack

A version of *Pfitzmann’s attack* is known to compromise mix-nets with randomized partial checking [18]. It can be used to break the privacy of a given vote with probability 1/2 of being undetected during RPC. The leaked information may differ depending on both the implementation of the attack and the voting protocol.

The main idea is that the first mix teller, who is corrupted, targets a ciphertext  $c_i$  from the odd mix input, and replaces some output term  $c_j$  with  $c_i^\delta$ , where  $\delta$  is a properly chosen integer (the criteria depend on the implementation). After the decryption results are posted, a pair of decrypted messages  $m$  and  $m'$  satisfying equation  $m' = m^\delta$  can be used to deduce sensitive information.

Clearly, the model presented in Section 4 is too basic to allow for detection of the attack. Still, this can be obtained by a simple extension of the model. For that, we change the Mteller template as shown in Figure 10. The only difference lies in how the first re-encryption mix is done: the corrupted mix teller targets  $c_0$ , chooses a random non-zero  $\delta$ , and uses  $c_0^\delta$  instead of some other output term. In all other respects, the teller behaves honestly.

Using the UPPAAL model checking functionality, it can be verified that there exist executions where the corrupted mix teller has failed the audit, as well as paths where he has passed. That is, both  $E\Diamond\text{failed\_audit}_0$  and  $E\Diamond\text{passed\_audit}_0$  produce ‘Property satisfied’ as the output. We note that, in order to successfully verify those properties in our model of Prêt à Voter, the search order option in UPPAAL should be changed from the (default) Breadth First to either Depth First or Random Depth First.

## 7 Related Work

Over the years, the properties of *ballot secrecy*, *receipt-freeness*, *coercion resistance*, and *voter-verifiability* were recognized as important for an election to work properly, see [22] for an overview. A number of secure and voter-verifiable schemes have been proposed, most notably the Prêt à Voter protocol for supervised elections [24], the Pretty Good Democracy approach to internet voting [26], and Selene, an enhanced form of tracking number-based scheme [25]. Nowadays, such schemes are starting to move out of the laboratory and into use in real elections. For example, Prêt à Voter has been successfully used in one of the state elections in Australia [7] while the Scantegrity II system [8] was used in municipal elections in the Takoma Park county, Maryland. This strongly emphasizes the need for extensive analysis and validation of such systems.

Formal analysis of selected voting protocols, based on theorem proving in first-order logic or linear logic, includes attempts at verification of vote counting in [1,23]. The Coq theorem prover [4] was used to implement the STV counting scheme in a provably correct way [11], and to produce a provably voter-verifiable variant of the Helios protocol [13]. Moreover, Tamarin [21] was used to verify receipt-freeness in Selene [6] and Electryo [27]. Approaches based on model checking are fewer and include the analysis of risk-limiting audits [2] with the CBMC model checker [9]. Moreover, [16] proposed and verified a simple multi-agent model of Selene using MCMAS [20]. Related research includes the use of multi-agent methodologies to specify and verify properties of authentication and key-establishment protocols [19,5] with MCMAS.

In all the above cases, the focus is on the verification itself. Indeed, all the tools mentioned above provide only a text-based interface for specification of the system. As a result, their model specifications closely resemble programming code, and insufficiently protect from the usual pitfalls of programming: unreadability of the code, lack of modularity, and opaque control structure. In this paper, we draw attention to tools that promote modular design of the model, emphasize its control structure, and facilitate inspection and validation.

## 8 Conclusions

Formal methods are well established in proving (and disproving) the correctness of cryptographic protocols. What makes voting protocols special is that they prominently feature human and social aspects. In consequence, an accurate

specification of the behaviors admitted by the protocol is far from straightforward. An environment that supports the creation of modular, compact, and – most of all – readable specifications can be an invaluable help.

In this context, the UPPAAL model checker has a number of advantages. Its modeling language encourages modular specification of the system behavior. It provides flexible data structures, and allows for parameterized specification of states and transitions. Last but not least, it has a user-friendly GUI. Clearly, a good graphical model helps to understand how the voting procedure works, and allows for a preliminary validation of the system specification just by looking at the graphs. Anybody who ever inspected a text-based system specification or the programming code itself will know what we mean.

In this paper, we try to demonstrate the advantages of UPPAAL through a case study based on the Prêt à Voter protocol. The models that we have obtained are neat, easy to read, and easy to modify. On the other hand, UPPAAL has not performed very well with the verification itself. This was largely due to the fact that its requirement specification language turned out to be very limited – much more than it seemed at the first glance. We managed to partly overcome the limitations by a smart reconstruction of models and formulas. In the long run, however, a more promising path is to extend the implementation of verification algorithms in UPPAAL so that they handle nested path quantifiers and knowledge modalities, given explicitly in the formula.

## References

1. B. Beckert, R. Goré, and C. Schürmann. Analysing vote counting algorithms via logic - and its application to the CADE election scheme. In *Proceedings of CADE*, volume 7898 of *Lecture Notes in Computer Science*, pages 135–144. Springer, 2013.
2. B. Beckert, M. Kirsten, V. Klebanov, and C. Schürmann. Automatic margin computation for risk-limiting audits. In *Proceedings of E-Vote-ID*, volume 10141 of *Lecture Notes in Computer Science*, pages 18–35. Springer, 2016.
3. G. Behrmann, A. David, and K. Larsen. A tutorial on UPPAAL. In *Formal Methods for the Design of Real-Time Systems: SFM-RT*, number 3185 in LNCS, pages 200–236. Springer, 2004.
4. Y. Bertot, P. Casteran, G. Huet, and C. Paulin-Mohring. *Interactive Theorem Proving and Program Development. Coq’Art: The Calculus of Inductive Constructions*. Springer, 2004.
5. I. Boureanu, P. Kouvaros, and A. Lomuscio. Verifying security properties in unbounded multiagent systems. In *Proceedings of International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1209–1217, 2016.
6. A. Bruni, E. Drewsen, and C. Schürmann. Towards a mechanized proof of selene receipt-freeness and vote-privacy. In *Proceedings of E-Vote-ID*, volume 10615 of *Lecture Notes in Computer Science*, pages 110–126. Springer, 2017.
7. C. Burton, C. Culnane, J. Heather, T. Peacock, P. Ryan, S. Schneider, V. Teague, R. Wen, Z. Xia, and S. Srinivasan. Using Prêt à Voter in victoria state elections. In *Proceedings of EVT/WOTE*. USENIX, 2012.
8. D. Chaum, R. Carback, J. Clark, A. Essex, S. Popoveniuc, R. Rivest, P. Ryan, E. Shen, A. Sherman, and P. Vora. Scantegrity II: end-to-end verifiability by

- voters of optical scan elections through confirmation codes. *Trans. Info. For. Sec.*, 4(4):611–627, 2009.
9. E. Clarke, D. Kroening, and F. Lerda. A tool for checking ANSI-C programs. In *Proceedings of TACAS*, volume 2988 of *Lecture Notes in Computer Science*, pages 168–176. Springer, 2004.
  10. E. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 995–1072. Elsevier, 1990.
  11. M. Ghale, R. Goré, D. Pattinson, and M. Tiwari. Modular formalisation and verification of STV algorithms. In *Proceedings of E-Vote-ID*, volume 11143 of *Lecture Notes in Computer Science*, pages 51–66. Springer, 2018.
  12. V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
  13. T. Haines, R. Goré, and M. Tiwari. Verified verifiers for verifying elections. In *Proceedings of CCS*, pages 685–702. ACM, 2019.
  14. M. Jakobsson, A. Juels, and R. Rivest. Making mix nets robust for electronic voting by randomized partial checking. In *USENIX Security Symposium*, pages 339–353, 2002.
  15. W. Jamroga. Knowledge and strategic ability for model checking: A refined approach. In *Proceedings of MATES’08*, volume 5244 of *Lecture Notes in Computer Science*, pages 99–110, 2008.
  16. W. Jamroga, M. Knapik, and D. Kurpiewski. Model checking the SELENE e-voting protocol in multi-agent logics. In *Proceedings of E-VOTE-ID*, volume 11143 of *Lecture Notes in Computer Science*, pages 100–116. Springer, 2018.
  17. G. Kant, A. Laarman, J. Meijer, J. van de Pol, S. Blom, and T. van Dijk. LTSmin: High-performance language-independent model checking. In *Tools and Algorithms for the Construction and Analysis of Systems. Proceedings of TACAS*, volume 9035 of *Lecture Notes in Computer Science*, pages 692–707. Springer, 2015.
  18. S. Khazaei and D. Wikstroem. Randomized partial checking revisited. In *Topics in Cryptology – CT-RSA 2013*, volume 7779 of *Lecture Notes in Computer Science*, pages 115–128. Springer, 2013.
  19. A. Lomuscio and W. Penczek. LDYIS: a framework for model checking security protocols. *Fundamenta Informaticae*, 85(1-4):359–375, 2008.
  20. A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: An open-source model checker for the verification of multi-agent systems. *International Journal on Software Tools for Technology Transfer*, 19(1):9–30, 2017.
  21. S. Meier, B. Schmidt, C. Cremers, and D. Basin. The TAMARIN prover for the symbolic analysis of security protocols. In *Computer Aided Verification, Proceedings of CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 696–701. Springer, 2013.
  22. B. Meng. A critical review of receipt-freeness and coercion-resistance. *Information Technology Journal*, 8(7):934–964, 2009.
  23. D. Pattinson and C. Schürmann. Vote counting as mathematical proof. In *Advances in Artificial Intelligence, Proceedings of AI*, volume 9457 of *Lecture Notes in Computer Science*, pages 464–475. Springer, 2015.
  24. P. Ryan. The computer ate my vote. In *Formal Methods: State of the Art and New Directions*, pages 147–184. Springer, 2010.
  25. P. Ryan, P. Rønne, and V. Iovino. Selene: Voting with transparent verifiability and coercion-mitigation. In *Financial Cryptography and Data Security: Proceedings of FC 2016. Revised Selected Papers*, volume 9604 of *Lecture Notes in Computer Science*, pages 176–192. Springer, 2016.



26. P. Ryan and V. Teague. Pretty good democracy. In *Security Protocols XVII*, volume 7028 of *Lecture Notes in Computer Science*, pages 111–130. Springer Berlin Heidelberg, 2013.
27. M. Zollinger and P. R. P. Ryan. Mechanized proofs of verifiability and privacy in a paper-based e-voting scheme. In *Proceedings of 5th Workshop on Advances in Secure Electronic Voting*, 2020.