

Centre Fédéré en Vérification

Technical Report number 2004.22

Formal Analysis of Multi-party Contract Signing

Rohit Chadha Steve Kremer Andre Scedrov



This work was partially supported by a FRFC grant: 2.4530.02

<http://www.ulb.ac.be/di/ssd/cfv>

Formal analysis of multi-party contract signing

Rohit Chadha^{1*}, Steve Kremer^{2**}, and Andre Scedrov^{3***}

¹ University of Sussex

² Université Libre de Bruxelles

³ University of Pennsylvania

Abstract. We analyze the multi-party contract-signing protocols of Garay and MacKenzie (GM) and of Baum and Waidner (BW). We use a finite-state tool, MOCHA, which allows specification of protocol properties in a branching-time temporal logic with game semantics. While our analysis does not reveal any errors in the BW protocol, in the GM protocol we discover serious problems with fairness for four signers and an oversight regarding abuse-freeness for three signers. We propose a complete revision of the GM subprotocols in order to restore fairness.

1 Introduction

The problem of digitally signing a contract over a network is more complicated than signing a contract by “pen and paper”. The problem arises because of an inherent asymmetry: no signer wants to be the first one to sign the contract because another signer could refuse to do so after having obtained the first signer’s contract.

A simple solution consists in using a trusted party (T) as an intermediary. Signers send their respective contracts to T , which first collects the contracts and then distributes them among the signers. An intermediary is known to be necessary [8]. However, because of the communication and computation bottleneck at T , this solution is inefficient. Other solutions include randomized protocols as well as protocols based on gradual information exchange. More recently, the so-called *optimistic* approach was introduced in [2, 4]. The idea is that T intervenes only when a problem arises, *e.g.*, a signer is trying to cheat or a network failure occurs at a crucial moment during the protocol. Such protocols generally consist of a main protocol and one or several subprotocols, each with a fixed number of messages. The main protocol is executed by the signers in order to exchange their signatures. The subprotocols are used to contact T in order to force a successful outcome or to abort the protocol.

A contract-signing protocol should respect several desirable properties. The first property is *fairness*. Intuitively, a contract-signing protocol is fair if at the end of the protocol either each signer obtains all the other signers’ contracts or no signer gets any valuable information. A second property, *timeliness*, ensures that signer has some recourse to prevent endless waiting. Both fairness and timeliness are standard properties that are also important in fair exchange, certified e-mail and fair non-repudiation protocols. A property that is specific to contract signing, *abuse-freeness*, was introduced in [9]. A protocol is abuse-free if no signer A is able to prove to an external observer that A has the power to choose between successfully concluding the protocol and aborting the protocol. A protocol that is not abuse-free gives an undesirable advantage to one signer, say Alice, who has the power to decide the outcome of the protocol and can prove this to an external observer. If, for instance, Alice wants to sell a house to Charlie, she could initiate a contract with Bob just to force Charlie to increase his offer.

* Partially supported by the ONR CIP/SW URI “Software Quality and Infrastructure Protection for Diffuse Computing” through ONR Grant N00014-01-1-0795 and by the NSF Grant CCR-0098096 and the EU Global Computing project “MIKADO”.

** This research was carried out while the author stayed at University of Pennsylvania funded by the “Communaut Franaise de Belgique”.

*** Partially supported by the ONR CIP/SW URI “Software Quality and Infrastructure Protection for Diffuse Computing” through ONR Grant N00014-01-1-0795 and by the NSF Grant CCR-0098096.

There have been several applications of formal methods to contract signing, so far only for the special case of two signers. The finite model-checker Mur ϕ is used in [13] to analyze two contract-signing protocols, discover subtle errors and suggest corrections. In [5] inductive methods are used to reason about contract-signing protocols specified in the multiset-rewriting framework, MSR. Protocol properties are expressed in terms of strategies, which provide a natural framework for the analysis. In [12] the finite model-checker MOCHA is used to analyze two contract-signing protocols. The advantage of using MOCHA rather than Mur ϕ is that MOCHA allows to specify protocol properties in ATL, a temporal logic with game semantics, which in turn allows reasoning about strategies. The most recent work on contract signing [6] introduces the notion of an *optimistic* signer, *i.e.*, a signer that prefers to wait for “some time” for messages from the other signers before contacting the trusted party. The main theorem of [6] is that, independently of a specific protocol, if any of the signers is optimistic, then the other signer will at some point of the protocol have the power to decide the outcome.

All the efforts just described consider only two-party protocols. In this paper we analyze multi-party contract-signing protocols [3, 10]. The protocol goal in that case is that each signer sends its signature on a previously agreed upon contract text to all other signers and that each signer receives all other signers’ signatures on this contract. In a multi-party framework, fairness, timeliness, and abuse-freeness should hold against *any* coalition of dishonest parties. Unlike in the two-party case, the complexity level of the multi-party protocols, especially [10], is such that a tool, e.g., a model-checker, is indispensable in the analysis. This partly comes about from an important difference between the two-party and the multi-party case, namely, in the multi-party case T has to be able to overturn its previous abort decisions [9]. As our analysis shows, this feature is particularly difficult to design correctly. We have discovered an essential obstacle in the GM protocol [10], which appears not to be removable without completely changing the subprotocols for T and which leads to the failure of fairness in the case of four signers. We present this attack in detail in the paper and propose a corrected version of the GM protocol, which has been validated by MOCHA. MOCHA did not find any problems with fairness in the BW protocol [3] nor in the original GM protocol with only three signers. In the latter case, MOCHA did find an amusing problem with abuse-freeness, but this problem is easily corrected. We believe that the main reason for robustness of the BW protocol is that overturning the aborts decisions has been designed correctly.

Outline of the paper. The rest of the paper is organized as follows. In section 2, we describe the BW and the GM protocols. In section 3, we present briefly the finite-state tool, MOCHA, the temporal logic ATL and its game semantics. Modeling of the protocols and protocol assumptions in the game semantics along with the modeling of fairness in ATL is briefly discussed. In section 4, we report on our analysis of the BW and GM protocols using MOCHA, present the fairness attack on four signers in detail and propose a corrected version of the protocol. We discuss briefly how to restore fairness and present the anomaly with respect to abuse-freeness for three signers. In order to detect this anomaly, we had to model optimistic signers and discuss this issue. We summarize our results and discuss directions for future work in section 5.

2 Protocol description

In this section we describe the multi-party contract-signing protocols proposed by Baum and Waidner in [3] and Garay and MacKenzie in [10]. Unlike two-party protocols, which generally have similar structures, the two multi-party protocols described below have fundamentally different structures. For this section and for the rest of the paper, we shall assume that each protocol participant has a private signing key and a corresponding public verification key. Each participant shall be identified with this private/public key pair, and if we say that “ A can ...”, we shall mean anyone that possesses the private key of A .

2.1 GM multi-party contract-signing protocol

The protocol allows n ($n \geq 2$) participants, say P_1, \dots, P_n , to exchange signatures with the help of a trusted party T on a preagreed contract text m . P_i is said to have a *contract* if it has everybody's signature on the text m . The order of the participants P_1, \dots, P_n , henceforth referred to as *signers*, and the identity of T are also agreed upon before the protocol begins. The preagreed contract text m contains an identifier that uniquely identifies each protocol instance. In [10], the communication amongst the participants is assumed to be over a network channel in control of a "Dolev-Yao intruder", while the communication between the participants and the trusted party is assumed to be over a private channel.

The protocol uses zero-knowledge cryptographic primitives, *private contract signatures*, that were first introduced in [9]. The private contract signature of A for B on text m with respect to a trusted party T , denoted as $PCS_A(m, B, T)$ has the following properties:

- a) $PCS_A(m, B, T)$ can be created by A .
- b) $PCS_A(m, B, T)$ can be faked by B . Only A , B and T can tell difference between PCS and its simulation.
- c) $PCS_A(m, B, T)$ can be converted into a conventional universally-verifiable digital signature, $S_A(m)$, by both A and T . Only A and T can do this conversion.

The protocol itself consists of three subprotocols: *main*, *abort*, and *recovery* subprotocols. Usually signers try to achieve the exchange by executing the main subprotocol. They contact T using one of the other two subprotocols when they think something is amiss. Once a signer contacts T , it no longer takes part in the main subprotocol. T responds to a request with either an abort token or a signed contract. The decision whether to reply with an abort token or with a signed contract is based on a database maintained by T , which stores all the relevant information of the requests and its responses. Once T sends back a signed contract, it always replies with the signed contract. As discussed below, a decision to abort may, however, be overturned in order to maintain fairness. We discuss the subprotocols in some detail.

Main protocol. The main protocol for n signers is divided into n -levels, that can be described recursively. For each level of recursion, a different "strength" of promise is used. The strength of a promise is denoted by an integer "level", and an "i-level promise from signer A to signer B on a message m " is implemented using PCS : $PCS_A((m, i), B, T)$.

In level i , signers P_i through P_1 exchange i -level promises to sign the contract. The i -level protocol is triggered when P_i receives 1-level promises from P_{i+1}, \dots, P_n . After receiving these promises, P_i sends out its 1 level promise to signers P_{i-1}, \dots, P_1 and waits for $i - 1$ level to finish. At the end of the $i - 1$ level, P_1, \dots, P_{i-1} have exchanged $i - 1$ level promises and P_i receives a $i - 1$ level promise from each of the signers P_1, \dots, P_{i-1} . Now P_i, \dots, P_1 exchange i level promises, and close the higher levels.

In order to close level a where $a > i$, P_i sends an $(a - 1)$ -level promise to P_a and waits for a -level promises from signers P_{i+1}, \dots, P_a . After receiving these promises, P_i indicates its willingness to close the level a to signers P_1, \dots, P_{i-1} by sending them its a -level promise, and in return waits for a -level promises from them. Upon receiving these, P_i sends its a -level promises to P_{i+1}, \dots, P_a completing its obligation in the a -level protocol. P_i then proceeds to complete $a + 1$ level.

Once the n -levels are completed, each signer has a n -level promise from everybody else, and the contract exchange is ready to begin. In this exchange, each signer also sends a $n + 1$ -level promise to everybody along with its signature on the preagreed text. In order to complete the exchange, signer P_i waits for the contract and $n + 1$ -level promises from P_n, \dots, P_{i+1} . Upon receiving these, P_i sends its signature and $n + 1$ -level promises to everybody, and waits for the signatures and $n + 1$ -level promises from P_{i-1}, \dots, P_1 . Once these are received, the protocol ends for P_i , and P_i has the contract.

If some expected messages are not received, P_i may either quit the protocol or contact T . P_i may simply quit the protocol if it has not sent any promises or contact T if it has sent some promises. It may contact T with a request to abort if it has not received any promise from some signer. It may request T to recover the protocol if it has a promise from every other signer. A detailed description of the main protocol is given in table 1.

Table 1 GM multi-party contract-signing protocol—Main

Wait for all higher recursive levels to start
1. $P_j \rightarrow P_i: PCS_{P_j}((m, 1), P_i, T)$ ($n \geq j > i$)
If P_i does not receive 1-level promises from $P_n \dots P_{i+1}$ in a timely manner, P_i simply quits.
Start recursive level i
2. $P_i \rightarrow P_j: PCS_{P_i}((m, 1), P_j, T)$ ($i > j \geq 1$)
Wait for recursive level $i-1$ to finish
3. $P_j \rightarrow P_i: PCS_{P_j}((m, i-1), P_i, T)$ ($i > j \geq 1$)
If P_i does not receive ($i-1$)-level promises from $P_{i-1} \dots P_1$ in a timely manner, P_i aborts.
Send i -level promises to all lower-numbered signers
4. $P_i \rightarrow P_j: PCS_{P_i}((m, i), P_j, T)$ ($i > j \geq 1$)
Finish recursive level i when i -level promises are received
5. $P_j \rightarrow P_i: PCS_{P_j}((m, i), P_i, T)$ ($i > j \geq 1$)
If P_i does not receive i -level promises from $P_{i-1} \dots P_1$ in a timely manner, P_i recovers.
Complete all higher recursive levels
For $a = i + 1$ to n , P_i does the following:
6.1. $P_i \rightarrow P_a: PCS_{P_i}((m, a-1), P_a, T)$
6.2. $P_j \rightarrow P_i: PCS_{P_i}((m, a), P_j, T)$ ($a \geq j > i$)
If P_i does not receive a -level promises from $P_a \dots P_{i+1}$ in a timely manner, P_i recovers.
6.3. $P_i \rightarrow P_j: PCS_{P_i}((m, a), P_j, T)$ ($i > j \geq 1$)
6.4. $P_j \rightarrow P_i: PCS_{P_i}((m, a), P_j, T)$ ($i > j \geq 1$)
If P_i does not receive a -level promises from $P_{i-1} \dots P_1$ in a timely manner, P_i recovers.
6.5. $P_i \rightarrow P_j: PCS_{P_i}((m, a), P_j, T)$ ($a \geq j > i$)
Wait for signatures and $(n+1)$ -level promises from higher-numbered signers
7. $P_j \rightarrow P_i: PCS_{P_j}((m, n+1), P_i, T), S_{P_j}(m, 1)$ ($n \geq j > i$)
If P_i does not receive signatures and $(n+1)$ -level promises from $P_n \dots P_{i+1}$ in a timely manner, P_i recovers.
Send signatures and $(n+1)$ -level promises to signers
8. $P_i \rightarrow P_j: PCS_{P_i}((m, n+1), P_j, T), S_{P_i}(m, 1)$ ($j \neq i$)
Wait for signatures from lower-numbered signers
9. $P_j \rightarrow P_i: PCS_{P_j}((m, n+1), P_i, T), S_{P_j}(m, 1)$ ($i > j \geq 1$)
If P_i does not receive signatures and $(n+1)$ -level promises from $P_{i-1} \dots P_1$ in a timely manner, P_i recovers.

Abort protocol. T maintains two sets, S_m and F_m , that are used by T to make decisions when a signer contacts T . These sets are created when T is contacted for the first time for m and are initialized to be empty. The set S_m contains the indices of all signers that have contacted T and received an abort token from T in response. The intuitive meaning of the set F_m is not clearly stated in [10], but it contains some additional information that T uses in deciding when to overturn an abort decision that T has taken before.

The details of the abort protocol are given in table 2. Mainly, if T is contacted with a request to abort, then T checks its database. If this is the first request or if the protocol has not already been recovered, T sends back an abort token and updates the sets S_m and F_m . If the protocol has already been successfully recovered, T sends back a signed contract.

Recovery protocol. The details of the recovery protocol are given in table 3. For P_i to recover, it sends the message $S_{P_i}(\{PCS_{P_j}((m, k_j), P_i, T)\}_{j \in \{1, \dots, n\} \setminus \{i\}}, S_{P_i}((m, 1)))$, where

- if $j > i$, k_j is the maximum level of a promise received from P_j on m ,
- if $j < i$, k_j is the maximum level of promises received from all signers $P_{j'}$, with $j' < i$, *i.e.*, the min-max of the level of promises from signers with lower index. (*E.g.*, if the maximum level of the promises received by P_4 from P_3 and P_2 was 6, and the maximum level received by P_4 from P_1 was 5, then it would send the 5-level promises for P_1 , P_2 and P_3 .)

If T is contacted with a request to recover, then T checks its database. If this is the first request for m or if the protocol has already been recovered, T replies with a signed contract

Table 2 GM multi-party contract-signing protocol—Abort

1. $P_i \rightarrow T: S_{P_i}(m, P_i, (P_1, \dots, P_n), abort)$
if not validated(m) then
if $S_m = \emptyset$, T stores $S_T(S_{P_i}(m, P_i, (P_1, \dots, P_n), abort)); S_m = S_m \cup \{i\}$;
if i is larger than the maximum index in S_m , T clears F_m
 2. $T \rightarrow P_i: S_T(S_{P_j}(m, P_j, (P_1, \dots, P_n), abort), S_T(m, S_m, abort))$
 - else (validated(m)=true)
 3. $T \rightarrow P_i: \{S_{P_j}((m, k_j))\}_{j \in \{1, \dots, n\} \setminus \{i\}}$
where k_j is the level of the promise from P_j that was converted to a universally-verifiable signature during the recovery protocol.
-

which it obtains by converting the promises into conventional digital signatures. Otherwise, if the protocol has already been aborted, T must decide whether to maintain the abort or to overturn it. Overturning of the abort is necessary in order to maintain fairness. Indeed, consider the scenario in which a dishonest P_{n-1} contacts T with an abort request, receives an abort token and dishonestly continues the protocol. After the n -levels are completed, P_n sends its signature to others and waits for signatures from other signers. If P_{n-1} does not send back its signature, then P_n will be forced to contact T with a request to recover. Now, T must overturn its previous abort, otherwise P_n will not receive the signature of P_{n-1} . The decision whether to overturn is based on the contents of the sets S_m and F_m , as described in table 3.

2.2 BW multi-party contract-signing protocol

The protocol allows n ($n \geq 2$) participants or *signers*, say P_1, \dots, P_n , to exchange signatures with the help of a trusted party T on a preagreed contract text m . In our description, we suppose $n - 1$ potentially dishonest signers. The original protocol is actually parameterized with respect to a threshold t , the maximum number of possibly dishonest signers. In our analysis however we assume the worst possible scenario for an honest signer, namely that all the other signers are dishonest (*i.e.*, t is $n - 1$).

The protocol consists of two subprotocols: main and recovery. Usually signers try to achieve the exchange by executing the main subprotocol. They contact T using the recovery subprotocol when they think something is amiss.

Main protocol. A detailed description of the main protocol for each signer P_i is given in table 4. The protocol is composed of $n + 1$ rounds¹. The main protocol is symmetric for each signer, and in round i a signer sends an i -level promise to other signers. In [3], the i -level promise is implemented using a universally-verifiable digital signature and includes all the promises (through the vectors M and X , as defined in table 4) that the signer has received till round $i - 1$. Note that these promises are based on universally verifiable signatures and that the protocol does not intend to provide abuse-freeness. The $n + 1$ level promises from everybody is considered to be the signed contract. If any expected message is not received, a signer can launch the recovery protocol. Once the signer launches the recovery protocol, it is not allowed to continue the main subprotocol.

Recovery protocol. The details of the recovery protocol are given in table 5. To recover, P_i sends a recovery request. If the recovery request is launched in the first round, *i.e.* P_i did not receive a message from all the signers, the recovery request consists in the first level promise. Otherwise, if $r > 1$, the recovery request contains, via the vector $X_{r-1,i}$, the set of received messages.

T maintains for each contract for which it has been contacted a variable `recovered` indicating whether the given contract has already been recovered or not, a set `con`, containing the indexes of the signers that contacted T for the given contract and the set `abort_set` containing the indexes of the signers for whom T aborted the protocol.

¹ In [3], the protocol has $t + 2$ rounds.

Table 3 GM multi-party contract-signing protocol—Recovery

1. $P_i \rightarrow T: S_{P_i}(\{PCS_{P_j}((m, k_j), P_i, T)\}_{j \in \{1, \dots, n\} \setminus \{i\}}, S_{P_i}((m, 1)))$
if $i \in S_m$, T ignores the message
else if $\text{validated}(m)$
 2. $T \rightarrow P_i: \{S_{P_j}((m, k_j))\}_{j \in \{1, \dots, n\} \setminus \{i\}}$
where k_j is the level of the promise from P_j that was converted to a universally-verifiable signature.
else if $S_m = \emptyset$
 $\text{validated}(m) := \text{true}$
 3. $T \rightarrow P_i: \{S_{P_j}((m, k_j))\}_{j \in \{1, \dots, n\} \setminus \{i\}}$
else ($\text{validated}(m) = \text{false} \wedge S_m \neq \emptyset$)
 - a. if $i \notin F_m$, then
 - (i) if for any $\ell \in S_m$ there is a $j \in S_m$ such that $j > k_\ell$
 $S_m := S_m \cup \{i\}$
 let a be the maximum value in S_m
 if $a > i$ then $\forall j$, such that $k_j = a - 1 \cdot F_m := F_m \cup \{j\}$
 else $F_m := \emptyset$
 - 4.1.1. $T \rightarrow P_i: S_T(S_{P_j}(m, P_j, (P_1, \dots, P_n), \text{abort}), S_T(m, S_m, \text{abort}))$
 where $S_T(S_{P_j}(m, P_j, (P_1, \dots, P_n), \text{abort}))$ corresponds to the stored abort token
 - (ii) else
 $\text{validated}(m) := \text{true}$
 - 4.1.2. $T \rightarrow P_i: \{S_{P_j}((m, k_j))\}_{j \in \{1, \dots, n\} \setminus \{i\}}$
 - b. else ($i \in F_m$)
 let a be the maximum value in S_m
 - (i) if ($\forall j$, such that $i < j \leq a \cdot k_j < a$) \wedge ($\forall j < i \cdot k_j \geq a$)
 $\text{validated}(m) := \text{true}$
 - 4.2.1. $T \rightarrow P_i: \{S_{P_j}((m, k_j))\}_{j \in \{1, \dots, n\} \setminus \{i\}}$
 - (ii) else
 $S_m := S_m \cup \{i\}$
 if $a > i$ then $\forall j, k_j = a - 1 \cdot F_m := F_m \cup \{j\}$
 if $a = i$ then $F_m := \emptyset$
 - 4.2.2. $T \rightarrow P_i: S_T(S_{P_j}(m, P_j, (P_1, \dots, P_n), \text{abort}), S_T(m, S_m, \text{abort}))$
 where $S_T(S_{P_j}(m, P_j, (P_1, \dots, P_n), \text{abort}))$ corresponds to the stored abort token

When T receives a valid recovery request, it first verifies whether this signer already contacted the T or not. If so, T ignores the request. If not, it checks whether the contract has already been successfully recovered or not. A successful recovery is always maintained. Otherwise, we distinguish two cases. If the recovery request is sent in the first round, T must abort the protocol, as the request does not contain a proof that all signers actually started the protocol. If the recovery request is sent during any later round, T verifies whether all the requests that were aborted previously occurred at least two rounds before. If so T can be sure that the previous abort replies were sent to dishonest players, as they continued the protocol and sent out higher level promises. Hence, the previous abort decision is overturned. Otherwise, T replies with an abort token.

3 Model

ATS, ATL and MOCHA. The desired properties of contract signing are easily described using games, and hence we chose a game-variant of Kripke structures, alternating transition systems (ATS) [1], to model the protocols. An ATS is composed of a set of players Σ , a set of states Q that represents all possible game configurations, a set $Q_0 \subseteq Q$ of initial states, a finite set of propositions Π , a labeling function $\pi : Q \rightarrow 2^\Pi$ that labels states with propositions, and a game transition function $\delta : Q \times \Sigma \rightarrow 2^{2^Q}$. For a player a and a state q , $\delta(q, a)$ is the set of choices that a can make in the state q . A choice is a set of possible next states. One step of the game at a state q is played as follows : each player $a \in \Sigma$ makes its choice and the next state of the game q' is the intersection (required to be a singleton) of the choices made by all the players of Σ , *i.e.*,

Table 4 Baum-Waidner multi-party contract signing protocol—Main

- $r := 1$
1. $P_i \rightarrow P_j: m_{1,i} = S_{P_i}(c, 1, \text{prev_round_ok})(j \neq i)$
 2. $P_j \rightarrow P_i: m_{1,j} = S_{P_j}(c, 1, \text{prev_round_ok})(j \neq i)$
if P_i times out then `recovery(1)`
 P_i computes vectors $M_{1,i} := (m_{1,1}, \dots, m_{1,n})$ and $X_{1,i} := M_{1,i}$
for $r := 2$ to $n + 1$ do
 3. $P_i \rightarrow P_j: m_{r,i} = S_{P_i}(M_{r-1,i}, r, \text{vec_ok}), S_{P_i}(c, r, \text{prev_round_ok})(i \neq j)$
 4. $P_j \rightarrow P_i: m_{r,j} = S_{P_j}(M_{r-1,j}, r, \text{vec_ok}), S_{P_j}(c, r, \text{prev_round_ok})(j \neq i)$
if P_i times out then `recovery(r)`
 P_i computes vectors

$$M_{r,i} := (S_{P_1}(c, r, \text{prev_round_ok}), \dots, S_{P_n}(c, r, \text{prev_round_ok}))$$

and

$$X_{1,i} := (S_{P_1}(M_{r-1,1}, r, \text{vec_ok}), \dots, S_{P_n}(M_{r-1,n}, r, \text{vec_ok}))$$

$\{q'\} = \cap_{a \in \Sigma} \delta(q, a)$. A computation is an infinite sequence $\lambda = q_0 q_1 \dots q_n \dots$ of states obtained by starting the game in q_0 , where $q_0 \in Q_0$.

In order to reason about ATS, we use alternating-time temporal logic (ATL) [1]. For a given set of players $A \subseteq \Sigma$, a set of computations Λ , and a state q , consider the following game between a protagonist and an antagonist starting in q . At each step, to determine the next state, the protagonist selects the choices controlled by the players in the set A , while the antagonist selects the remaining choices. If the resulting infinite computation belongs to the set Λ , then the protagonist wins. If the protagonist has a winning strategy, we say that the ATL formula $\langle\langle A \rangle\rangle \Lambda$ is satisfied in state q . Here, $\langle\langle A \rangle\rangle$ is a path quantifier, parameterized by the set A of players, which ranges over all computations that the players in A can force the game into, irrespective of how the players in $\Sigma \setminus A$ proceed. The set Λ is defined using temporal logic formulas. For those familiar with branching time temporal logics, the parameterized path quantifier $\langle\langle A \rangle\rangle$ can be seen as a generalization of the CTL path quantifiers: the existential path quantifier \exists corresponds to $\langle\langle \Sigma \rangle\rangle$ and the universal path quantifier \forall corresponds to $\langle\langle \emptyset \rangle\rangle$.

We now illustrate the expressive power of ATL that allows to model both cooperative as well as adversarial behavior amongst the players. Consider the set of players $\Sigma = \{a, b, c\}$ and the following formulas with their verbal reading:

- $\langle\langle a \rangle\rangle \diamond p$, player a has a strategy against players b and c to eventually reach a state where the proposition p is true;
- $\neg \langle\langle b, c \rangle\rangle \square p$, the coalition of players b and c does not have a strategy against a to make p true forever;
- $\langle\langle a, b \rangle\rangle \circ (p \wedge \neg \langle\langle c \rangle\rangle \square p)$, a and b can cooperate so that the next state satisfies p and from there c does not have a strategy to impose p forever.

The details of ATS and ATL can be found in [1]. Instead of modeling protocols directly with ATS we use a more user-oriented notation: a guarded command language a la Dijkstra. The details about the syntax and semantics of this language (given in terms of ATS) can be found in [11]. Intuitively, each player $a \in \Sigma$ disposes of a set of guarded commands of the form $\text{guard}_\xi \rightarrow \text{update}_\xi$. A computation-step is defined as follows: each player $a \in \Sigma$ chooses one of its commands whose boolean guard evaluates to true, and the next state is obtained by taking the conjunction of the effects of each update part of the commands selected by the players. Given an ATS described in terms of guarded commands, the finite state tool MOCHA automates the model-checking of ATL formulae over the specified ATS.

Modeling protocols. Unlike the classical security protocols aiming at secrecy and authentication, optimistic contract-signing protocols usually consist of subprotocols that can be invoked at specified moments. Running a protocol at a time not foreseen by the designer, may have unexpected

Table 5 Baum-Waidner multi-party contract signing protocol—Recovery

1. $P_i \rightarrow T$: $\text{resolve}_{r,i}$
where
$$\text{resolve}_{r,i} = \begin{cases} (1, i, S_{P_i}(m_{1,i}, \text{resolve})) & \text{if } r = 1 \\ (r, i, S_{P_i}(X_{r-1,i}, \text{resolve})) & \text{otherwise} \end{cases}$$

if $i \in \text{con}$ then stop
else if recovered
con:=con $\cup\{i\}$
 2. $T \rightarrow P_i$: $\text{signed}_{r,i}$
where $\text{signed}_{r,i} = \text{first_signed}$
else (\neg recovered)
 - a. if $r = 1$
abort_set:=abort_set $\cup\{(r, i)\}$
con:=con $\cup\{i\}$
 3. $T \rightarrow P_i$: $\text{aborted}_{r,i}$
where $\text{aborted}_{r,i} = S_T(c, r, i, \text{aborted})$
 - b. else ($r > 1$)
 - (i) if $\forall (s, k) \in \text{abort_set} \cdot s < r - 1$
if con= \emptyset then T sets $\text{first_signed} := (\text{resolve}_{r,i}, S_T(c, r, i, \text{recovered}))$
recovered:=true
con:=con $\cup\{i\}$
 4. $T \rightarrow P_i$: $\text{signed}_{r,i}$
where $\text{signed}_{r,i} = \text{first_signed}$
 - (ii) else
abort_set:=abort_set $\cup\{(r, i)\}$
con:=con $\cup\{i\}$
 5. $T \rightarrow P_i$: $\text{aborted}_{r,i}$
where $\text{aborted}_{r,i} = S_T(c, r, i, \text{aborted})$
-

side-effects. This may be used by a signer to gain an advantage over other signers. We believe that such concurrency issues are a major source of problems. Therefore, and since the high number of messages would create a serious state explosion, we only analyze the *structure* of the protocols and concentrate only on single protocol instance. We now discuss our model in detail.

The protocol instance is modeled as an ATS and each protocol participant is modeled as a player in the ATS using the above introduced guarded command language. Besides, the branching aspect, another notable difference with more classical secrecy and authentication protocols, is that contract-signing protocols must be secure against malicious signer, rather than an external intruder. In order to model this we have two process for each signer, one describing it's honest behavior and the other the dishonest one. Communication is modeled using shared variables. Each protocol message is modeled using a boolean variable, initialized to *false* and set to *true* when it is sent. Sending of a message is modeled using guarded commands, where the guard depends on previously sent out messages. When modeling the honest behavior of a participant, we ensure that a given message is sent out only when specified by the protocol. In contrary, the guards are relaxed in the malicious version of the signer so that each message can be sent out, as soon as possible, *i.e.*, as soon as all messages needed to compose the given message are received. We do not explicitly model any cryptographic primitives, but only the fact that protocol messages can be sent out of order. Hence, a dishonest signer can send messages out of order and continue the protocol, even if it is supposed to stop. We manually decide which messages must be known in order to send some other message. Moreover, the communication between any two signers is assumed to be on private channels and we do not model the possibility to spy other channels. The trusted party is modeled to be always honest.

As an example, consider the short extract of the modeling of the three-party GM protocol depicted in figure 1. In the extract, the integer variable Pr_{i-j}_L models the promises that P_i has sent to P_j , and $\text{Pr}_{i-j}_L = k$ means that P_i has sent out up to k -level promises to P_j . (For

efficiency reasons, we use the logarithmic encoding of a ranged integer variable, rather than having one boolean variable for each level of promise). In the extract, the first rule of honest P_1 says that P_1 may quit the protocol, if it has not contacted the trusted party, and has neither received nor sent any promises. The corresponding modeling of dishonest behavior of P_1 states that P_1 may quit the protocol at any moment. The second rule of honest P_1 gives the exact condition when the first level promise has to be sent to P_2 . The corresponding dishonest rule, merely requires that P_1 has not quit the protocol before sending the promise.

<p>Extract of honest modeling of P_1 for the three-party GM protocol:</p> <pre> [] ~P1_stop & ~P1_contacted_T & Pr_1_3_L=0 & Pr_1_2_L=0 & ~(Pr_3_1_L>0 & Pr_2_1_L>0) -> P1_stop' := true [] ~P1_stop & ~P1_contacted_T & Pr_1_3_L=0 & Pr_1_2_L=0 & Pr_2_1_L>0 & Pr_3_1_L>0 -> Pr_1_2_L' := 1 </pre> <p>The corresponding actions of a dishonest modeling:</p> <pre> [] ~P1_stop -> P1_stop' := true [] ~P1_stop & Pr_1_2_L < 1 -> Pr_1_2_L' := 1 </pre>

Fig. 1: Extract of the three-party GM protocol modeling

As we are unable to verify parametric systems with MOCHA, we simplify our task and verify the protocols only for a given number n of signers. In order to avoid encoding each instance of the protocol using guarded commands, we have written a dedicated C++ program for each protocol which takes the number n of signers as a parameter and generates the protocol specification. Although our model is restricted with regard to several aspects, the model seems to be of interest as several unknown anomalies have been revealed.

Modeling properties. We express the desired security guarantees using ATL. We concentrate on modeling of fairness. Consider an instance of the protocol with n signers, which we denote as P_1, \dots, P_n . In the following, we assume that only one of the signers, say P_1 is honest, and the other dishonest signers are colluding to cheat the honest signer.

A protocol is fair for an honest P_1 , if at the end of the protocol, either P_1 receives signed contracts from all the other signers or it is not possible for any other signer to obtain P_1 's signed contract. One possible ATL formula for modeling this says that if any signer receives P_1 's signed contract, then P_1 has a strategy to get the signed contracts from other signers:

$$\forall \square ((P_2.S_{P_1}(m) \vee \dots \vee P_n.S_{P_1}(m)) \rightarrow \langle\langle P_1 \rangle\rangle \diamond (P_1.S_{P_2}(m) \wedge \dots \wedge P_1.S_{P_n}(m))) \quad (1)$$

where $P_i.S_{P_j}(m)$ denotes that player P_i received P_j 's signature on the contract text m .

It can of course be argued that P_1 having a strategy to receive the signed contracts is not a sound modeling of fairness: P_1 may have this strategy but if it is ignorant or if it mistakenly does not follow this strategy, then the protocol may end in an unfair state. Therefore one could require the following stronger property: in whatever way P_1 resolves the remaining choices specified by the protocol, P_1 receives all the signed contract.

$$\forall \square ((P_2.S_{P_1}(m) \vee \dots \vee P_n.S_{P_1}(m)) \rightarrow \forall \diamond (P_1.S_{P_2}(m) \wedge \dots \wedge P_1.S_{P_n}(m))) \quad (2)$$

In the same vein, a third formulation only requires that there exists a path where P_1 receives the signed contracts.

$$\forall \square ((P_2.S_{P_1}(m) \vee \dots \vee P_n.S_{P_1}(m)) \rightarrow \exists \diamond (P_1.S_{P_2}(m) \wedge \dots \wedge P_1.S_{P_n}(m))) \quad (3)$$

Formula (2) implies formula (1), which in turn implies formula (3). We concentrate on the last, weakest version of fairness. As we show that fairness is violated even in this weakest version, the other stronger versions are also violated. Now, we are ready to discuss our analysis.

It is also possible to express abuse-freeness in our formalism. We are going to discuss this property when analysing the GM protocol

4 Analysis

We have verified the BW protocol with two and up to five signers, but the model-checker MOCHA did not find any flaw. Due to lack of space we do not go into the details of our analysis of the BW protocol. However, the design of the BW protocol is much simpler than the GM protocol and the decision to overturn an abort is based on the following argument: T only aborts if it can infer that no previous abort reply has been sent to a honest principal. This seems to ensure the robustness of the protocol. Note that we only analyzed the structure of the protocol. Hence, our results only prove that the protocol is correct in the given model. Nevertheless, we believe that the protocol would also be correct in a more general model as all messages are signed and a unique contract identifier is used.

We now report in more detail on our the analysis of the GM contract-signing protocol. The protocol has several peculiarities. The most notable one is that the protocol changes with the number of signers, *e.g.*, the protocol specification of P_1 differs when the value of n changes. The number of protocol messages increases considerably with the number of signers. For instance, if we have $n = 3$, the main protocol has 20 messages and there are 14 different recovery requests. When $n = 4$, the corresponding numbers are 41 and 36. Moreover, the protocol is not symmetric for the signers: the protocol specification for P_i is different from that for P_j , for all $i \neq j$. For instance, when $n = 4$, P_1 can launch 18 different recovery requests and P_4 only 2.

As mentioned in section 3, we have written a dedicated C++ program that takes the number of signers, n , as a parameter and generates the protocol specification. Our analysis revealed problems with fairness, when n is 4. Although, we did not discover any fairness problems when $n = 3$, we did find an amusing problem with abuse-freeness. We did not discover any problems with timeliness in the protocol. All these anomalies are novel and the protocol was believed to be secure since it was first published. We discuss our results in detail. The sources codes of our analysis of both protocols are also available at the following website <http://www.ulb.ac.be/di/scsi/skremer/MPCS/>.

Fairness. We did not discover any problems with fairness when $n = 3$. The formulas representing fairness for P_1 , P_2 and P_3 , introduced in section 3, are validated by MOCHA. However, as we use a restricted model and consider single runs, we can only conclude that the protocol does not present any *structural* weakness for $n = 3$. Indeed, if we relax the assumption of the private channels, the anomaly presented by Shmatikov and Mitchell in [13] on the two-party GJM protocol can be adapted to the multi-party version. In this scenario, a malicious signer eavesdrops on the channel between the honest signer and T , and succeeds in compromising fairness. With our present modeling, we do not find such flaws, as this requires to eavesdrop channels and to decompose messages. The fix proposed by Shmatikov and Mitchell applies to the multi-party protocol too. However, we should emphasize that the authors of the GM protocol require the channels to T to be private and hence this scenario does not represent a valid attack on the protocol.

We discovered several scenarios that compromised fairness when $n = 4$. The first scenario was discovered by hand, when we found an error in the proof of correctness given in the original paper [10]. A detailed analysis using MOCHA detected seven other scenarios. An analysis of these revealed that the proof also did not cover a case. In each scenario, an honest signer is cheated by the coalition of three malicious signers. These scenarios follow the general outline:

1. A dishonest signer contacts the trusted party, T , at the beginning of the protocol, gets an abort token, and dishonestly continues participating in the main protocol.
2. A second dishonest signer tries to recover at some later point. It does not succeed, but manages to put the honest signer in the list F_m . It dishonestly continues the main protocol.
3. The honest signer is forced to recover, but is not successful in getting the abort decision overturned since it is in the list F_m .

4. The third dishonest signer contacts T and manages to overturn the decision. Hence, while the honest signer does not get any signed contract, the honest signer's contract is obtained.

For lack of space, we just describe one of these scenarios in detail. In this scenario, P_1 , P_3 and P_4 collude to cheat P_2 . The scenario proceeds as follows:

- At the beginning of the protocol, P_3 aborts the protocol and T updates $S_m = \{3\}$. However, unlike specified by the protocol, dishonest P_3 continues the main protocol execution.
- As soon as P_1 receives the second level promise from P_2 , it asks T to recover by sending

$$S_{P_1}(\{PCS_{P_2}((m, 2), P_1, T), PCS_{P_3}((m, 1), P_1, T), PCS_{P_4}((m, 1), P_1, T)\}, S_{P_1}((m, 1))).$$

T refuses this request, answers with an abort message and updates $S_m = \{1, 3\}$ and $F_m = \{2\}$. As P_3 did before, P_1 also continues the protocol.

- The main protocol is executed normally until signer P_2 reaches point 6.2. (see table 1) of the protocol with $a = 4$. At that point P_2 has sent out the set of message

$$\{PCS_{P_2}((m, 1), P_1, T), PCS_{P_2}((m, 2), P_1, T), \\ PCS_{P_2}((m, 2), P_3, T), PCS_{P_2}((m, 3), P_1, T), \\ PCS_{P_2}((m, 3), P_3, T), PCS_{P_2}((m, 3), P_4, T)\}$$

and has received the set of messages

$$\{PCS_{P_4}((m, 1), P_2, T), PCS_{P_3}((m, 1), P_2, T), \\ PCS_{P_1}((m, 1), P_2, T), PCS_{P_1}((m, 2), P_2, T), \\ PCS_{P_3}((m, 3), P_2, T), PCS_{P_1}((m, 3), P_2, T)\}.$$

P_2 is at position 6.2. with $a = 4$ and is waiting for 4-level promises from P_3 and P_4 . P_3 and P_4 do not reply and P_2 is forced to send the following recovery request to T .

$$S_{P_2}(\{PCS_{P_2}((m, 3), P_2, T), PCS_{P_3}((m, 3), P_2, T), PCS_{P_4}((m, 1), P_2, T)\}, S_{P_2}((m, 1))).$$

P_2 is in F_m , the tests in the protocol description (see table 3) indicate that T refuses the request, updates $S_m = \{1, 2, 3\}$ and replies with an abort message. F_m remains unchanged.

- P_4 launches a resolve request, sending

$$S_{P_4}(\{PCS_{P_1}((m, 3), P_4, T), PCS_{P_2}((m, 3), P_4, T), PCS_{P_3}((m, 3), P_4, T)\}, S_{P_4}((m, 1))).$$

This request overturns the previous aborts, and hence violates fairness as T sends the signed contract back to P_4 .

In other scenarios, we discovered that protocol is unfair for signers P_1 , P_2 and P_3 when $n = 4$. We did not find any scenario that compromised fairness for P_4 . This is probably because the tests indicate that P_4 can never be added to F_m , when $n = 4$.

Correcting the Garay-MacKenzie Protocol. In order to restore fairness in the Garay-MacKenzie protocol, we had to do major revisions in the recovery protocol. We were unsuccessful to restore fairness with minor changes, and we believe that this is because the meaning of the list F_m is not clear in the protocol. The central idea behind the revision is that T , when presented with a recovery request, overturns its abort decision if and only if T can infer dishonesty on the part of each of the signer that contacted T in the past. This is also the main idea behind the recovery protocol in [3].

The main protocol remains the same. Major changes are in the recovery protocol. The recovery messages are designed so that T can infer the promises that an *honest* signer would have sent when it launched the recovery protocol (note that a signer may have dishonestly sent other promises). For P_i to recover, it sends the message

$$S_{P_i}(\{PCS_{P_j}((m, k_j), P_i, T)\}_{j \in \{1, \dots, n\} \setminus \{i\}}, S_{P_i}((m, 1)))$$

where k_j is computed as following:

Table 6 Revised GM multi-party contract-signing protocol—Abort

1. $P_i \rightarrow T: S_{P_i}(m, P_i, (P_1, \dots, P_n), abort)$ if not validated(m) then $S_m = S_m \cup \{i\};$ if $S_m = \emptyset$, T stores $S_T(S_{P_i}(m, P_i, (P_1, \dots, P_n), abort));$ 2. $T \rightarrow P_i: S_T(S_{P_j}(m, P_j, (P_1, \dots, P_n), abort), S_T(m, S_m, abort))$ else (validated(m)=true) 3. $T \rightarrow P_i: \{S_{P_j}((m, k_j))\}_{j \in \{1, \dots, n\} \setminus \{i\}}$ where k_j is the level of the promise from P_j that was converted to a universally-verifiable signature during the recovery protocol.

1. If P_i runs the resolve protocol in step 5 of the main protocol (see table 1), then $k_j = 1$ for $j > i$ and $k_j = i - 1$ for $j < i$.
2. In step 6.2 of the main protocol, $k_j = a - 1$ for $1 < j \leq a - 1, j \neq i$ and $k_j = 1$ for $j > a - 1$.
3. In step 6.4 of the main protocol, $k_j = a - 1$ for $j < i$, $k_j = a$ for $i < j \leq a$ and $k_j = 1$ for $j > a$.
4. In step 7 of the main protocol, $k_j = n$ for all j .
5. In step 9 of the main protocol, $k_j = n$ for all $j < i$ and $k_j = n + 1$ for all $j > i$.

k_j may alternately be computed as:

- If $j < i$, k_j is the maximum level of promises received from all signers $P_{j'}$, with $j' < i$, *i.e.* the min-max of the promises from signers with lower index. (For example, if the maximum level of the promises received by P_4 from P_3 and P_2 was 6, and the maximum level received by P_4 from P_1 was 5, then it would send the 5-level promises for P_1, P_2 and P_3 .)
- Let l be the maximum value l' such that P_i has l' level promises from P_j for all $i \leq j \leq l'$. If no such l' exists then let l be 0. If $l = 0$, then let $k_j = 1$ for all $j > i$. If $l \neq 0$, then let $k_j = l$ for all $i \leq j \leq l$ and $k_j = 1$ for all $j > l$. (For example, if P_2 has received level 1 promise from P_6 , level 5 and 1 promises from P_5 , level 5, 4 and 1 promises from P_4 , and level 4, 3 and 1 promises from P_3 then $k_6 = 1, k_5 = 1, k_4 = 4, k_3 = 4$.)

T maintains the set S_m of indices of signers that contacted T in the past and received an abort token. For each signer P_i in the set S_m , T also maintains two integer variables $h_i(m)$ and $l_i(m)$. Intuitively, h_i is the highest level promise an honest P_i could have sent to any higher indexed signer before it contacted T . l_i is the highest level promise an honest P_i could have sent to a lower indexed signer before it contacted T . The protocol for T works as follows:

- If T ever replies with a signed contract for m , then T responds with the contract for any further request.
- If the first request to T is a resolve request, then T sends back a signed contract.
- If the first request is an abort request, then T aborts the contract. T may overturn this decision in the future if it can deduce that all the signers in S_m have behaved dishonestly. T deduces that a signer P_i in S_m is dishonest when contacted by P_j if
 1. $j > i$ and P_j presents to T a k -level promise from P_i such that $k > h_i(m)$, or
 2. $j < i$ and P_j presents to T a k -level promise from P_i such that $k > l_i(m)$.

We describe the abort and recovery protocols in detail in table 6, respectively 7.

We analyzed the revised protocol for both 3 and 4 signers and MOCHA did not detect any errors in the revised protocol. Please note that this should not be construed as proof of correctness since we are using a restricted communication model and are modeling a single run. Nevertheless, we believe that the revised protocol would be fair in a more general setting, and for an arbitrary number of signers.

- the abort reply contains the set $S_m = \{1, 2\}$ and is different from the one P_2 received,
- if P_1 receives an abort reply from T , it is always the answer to a recovery request,
- a recovery request always includes a promise from each signer which is verified by T .

From these remarks, we can conclude that if P_1 shows the abort reply to Charlie, then Charlie will be convinced that P_3 has started the protocol even though Charlie is unable to verify the PCS from P_3 . In other words, we can say that T has verified the PCS for Charlie. Also note that at this point P_1 and P_2 can force the exchange to abort by simply quitting the protocol: P_3 has no promises from P_1 and P_2 . P_1 and P_2 can also force a successful completion of the contract exchange by simply (dishonestly) engaging P_3 in the main protocol. Hence the protocol is not abuse-free for P_3 .

This vulnerability can be easily addressed by excluding the set S_m from the abort reply. In this case, the abort messages from P_3 and P_2 are exactly similar and can be obtained by P_2 without P_3 's participation. Hence, an abort reply does not prove P_3 's participation in the protocol. This rather amusing scenario illustrates that sometimes additional information may be harmful. While explicitness is often considered a good engineering practice (and we do not attempt to criticize such thumb rules), care should be taken when applying these principles. In personal communication with the authors of the protocol, they propose a different fix in letting P_1 abort the protocol rather than just quitting.

Abuse-freeness for P_3 is naturally expressed in ATL as follows:

$$\neg\exists\Diamond(T.\text{send}(\text{abort})\text{ to }P_1 \wedge \langle\langle P_1, P_2 \rangle\rangle\Box(\neg P_3.S_{P_1}(m) \vee \neg P_3.S_{P_2}(m)) \wedge \langle\langle P_1, P_2 \rangle\rangle\Box(P_3.\text{stop} \rightarrow (P_1.S_{P_3}(m) \wedge P_2.S_{P_3}(m)))$$

The boolean variable $T.\text{send}(\text{abort})\text{ to }P_1$ is set to true when P_1 receives the the abort token $S_T(S_{P_2}(m, P_2, (P_1, P_2, P_3), \text{abort}))$ with $S_m = \{1, 2\}$. As discussed before, this serves as a proof of P_3 's participation. The variables $P_i.S_{P_j}(m)$ reflect that player i has received player j 's signature on the contract. More precisely, the formula requires that it is possible to reach a point where

1. P_1 and P_2 can prove to Charlie that the protocol was started by P_3 ($T.\text{send}(\text{abort})\text{ to }P_1$ is true),
2. P_1 and P_2 have a strategy to choose an unsuccessful outcome, *i.e.*, P_3 cannot get some signer's contract ($\langle\langle P_1, P_2 \rangle\rangle\Box(\neg P_3.S_{P_1}(m) \vee \neg P_3.S_{P_2}(m))$ is true), and
3. P_1 and P_2 have a strategy to choose a successful outcome, *i.e.*, when honest P_3 stops, the dishonest players must have obtained P_3 's contract ($\langle\langle P_1, P_2 \rangle\rangle\Box(P_3.\text{stop} \rightarrow (P_1.S_{P_3}(m) \wedge P_2.S_{P_3}(m)))$ is true).

The requirement of P_3 stopping in condition 3 is to prevent it from idling forever. Even though as discussed before, the protocol is not abuse-free if P_3 is optimistic, the above formula is validated if P_3 is honest. An honest P_3 may contact T non-deterministically as permitted by the protocol. Indeed, in the scenario discussed above, an honest P_3 could prevent P_1 and P_2 from getting P_3 's signature if it contacts T . Therefore, in order to capture the scenario described above, we needed to model optimistic signers.

Following [6], we implement an optimistic signer by adding *signals* that the signer uses to decide when to quit waiting for messages from other signers and contact T . P_3 uses 3 signals for this: one to decide when to ask T to abort and 2 to decide when to contact T for the two recovery protocols that P_3 can launch. These signals are controlled by a new player, $P3TimeOuts$, that is added to the model.

The decision to abort is modeled by 2 boolean variables: *setTimeoutAbort* and *expiredTimeoutAbort*. While P_3 changes the value of *setTimeoutAbort*, *expiredTimeoutAbort* is changed by $P3TimeOuts$. When P_3 sends level 1 promise to P_1 and P_2 , it sets the value of *setTimeoutAbort* to *true*, and then waits for level 2 promises from them. $P3TimeOuts$ may set *expiredTimeoutAbort* to *true* once *setTimeoutAbort* is set to *true* by P_3 . If the promises arrive before *expiredTimeoutAbort*

is *true*, then P_3 continues with the main protocol, otherwise P_3 may contact T with an abort request. The decision to send recovery requests are modeled similarly.

Following [6], abuse-freeness is modeled by having a coalition of $P3TimeOuts$, P_1 and P_2 . This coalition can choose a sufficiently "long time" to keep P_3 from contacting T , while allowing P_1 and P_2 to schedule its messages in order to get the desired result. Abuse-freeness can then be expressed as

$$\begin{aligned} & \neg\exists\Diamond(T.send(abort) \text{ to } P_1 \wedge \\ & \quad \langle\langle P_1, P_2, P3TimeOuts \rangle\rangle\Box(\neg P_3.S_{P_1}(m) \vee \neg P_3.S_{P_2}(m)) \wedge \\ & \quad \langle\langle P_1, P_2, P3TimeOuts \rangle\rangle\Box(P_3.stop \rightarrow (P_1.S_{P_3}(m) \wedge P_2.S_{P_3}(m))) \\ &) \end{aligned}$$

Please note that even an optimistic P_3 should eventually be allowed to contact T , otherwise P_3 may be stuck forever. Hence, $P3TimeOuts$ must eventually set *expiredTimeOutAbort* and other signals to *true*. Ideally, this should be set non-deterministically. However, ensuring that a variable changes its value in MOCHA slows down verification considerably. In order to make the verification feasible, we put a maximum limit, *tick*, on the number of computation steps after which the value must change, and vary this limit manually. Please note that the signals may change before this limit is reached. This modeling is sound in the sense that if the formula is violated for some value of *tick* then abuse-freeness must be violated: P_3 just needs to wait for sufficiently "long time" to allow P_1 and P_2 to schedule its messages. Indeed, the above property is violated when *tick* is set to 3 giving us the attack on abuse-freeness. As expected, if we drop the player $P3TimeOuts$ in the formula, then the property is not violated: P_1 and P_2 are not able to schedule their messages ahead of P_3 .

However, if P_3 is not optimistic, the above given formulation of abuse-freeness is not violated. In a non-optimistic setting, the vulnerability can be detected by weakening the third requirement of the formula in order to say that there is a trace in which P_3 does not contact T , and P_1 and P_2 get P_3 's signed contract. The ATL formula which captures this is

$$\begin{aligned} & \neg\exists\Diamond(T.send(abort) \text{ to } P_1 \wedge \\ & \quad \langle\langle P_1, P_2 \rangle\rangle\Box(\neg P_3.S_{P_1}(m) \vee \neg P_3.S_{P_2}(m)) \wedge \\ & \quad \exists\Box(P_3.stop \rightarrow (P_1.S_{P_3}(m) \wedge P_2.S_{P_3}(m))) \\ &) \end{aligned}$$

This property is violated even in the original non-optimistic model. MOCHA also detected the violation of a stronger version of abuse-freeness proposed in [12]. This formulation requires that as soon as P_2 and P_3 can prove to Charlie that P_1 started the protocol, then they may not achieve an unsuccessful outcome anymore.

5 Conclusions and Future Work

We have studied two multi-party contract-signing protocols [10, 3] using a finite-state tool, MOCHA, that allows specification of properties in a branching-time temporal logic with game semantics. In order to make this analysis feasible, we model single runs and assume a restricted communication model. Our analysis did not find any errors in the BW protocol [3]. We did encounter problems with fairness in the case of four signers in the GM protocol [10]. It appears that fairness cannot be restored without completely rewriting the subprotocols. The revised subprotocols are inspired by the BW protocol. We also discovered a rather amusing problem with abuse-freeness in the GM protocol with three signers that occurs because abort messages from the trusted party reveal who have contacted it in the past. This problem is easily addressed by ensuring that the trusted party does not send this extra information. We had to implement optimistic signers to demonstrate this problem using MOCHA.

We plan to verify the protocols without fixing the number of signers. One major challenge in such a parametric verification is that the protocol descriptions change fundamentally with the

number of signers in that the protocol for n signers is not merely putting n identical processes in parallel. We hope to prove the correctness of these protocols in a more general setting which accounts for cryptography, multiple concurrent sessions, and relaxes the communication model. We plan to use, at least partially, abstraction techniques such as proposed by Das and Dill [7] to achieve this.

References

1. Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. In *38th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 1997.
2. N. Asokan, Matthias Schunter, and Michael Waidner. Optimistic protocols for fair exchange. In *4th ACM Conference on Computer and Communications Security*, Zurich, Switzerland, April 1997. ACM Press.
3. Birgit Baum-Waidner and Michael Waidner. Round-optimal and abuse free optimistic multi-party contract signing. In *Automata, Languages and Programming — ICALP 2000*, volume 1853 of *Lecture Notes in Computer Science*, pages 524–535, Geneva, Switzerland, July 2000. Springer-Verlag.
4. H. Burk and A. Pfitzmann. Value exchange systems enabling security and unobservability. In *Computers and Security*, 9(8):715–721, 1990.
5. Rohit Chadha, Max Kanovich, and Andre Scedrov. Inductive methods and contract-signing protocols. In *8th ACM Conference on Computer and Communications Security*, Philadelphia, PA, USA, November 2001. ACM Press.
6. Rohit Chadha, John C. Mitchell, Andre Scedrov, and Vitaly Shmatikov. Contract signing, optimism, and advantage. In *CONCUR 2003 — Concurrency Theory*, volume 2761 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
7. Satyaki Das and David L. Dill. Successive approximation of abstract transition relations. In *Sixteenth Annual IEEE Symposium on Logic in Computer Science*, 2001.
8. Shimon Even and Yacov Yacobi. Relations among public key signature systems. Technical Report 175, Technion, Haifa, Israel, March 1980.
9. Juan A. Garay, Markus Jakobsson, and Philip D. MacKenzie. Abuse-free optimistic contract signing. In *Advances in Cryptology—Crypto 1999*, volume 1666 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
10. Juan A. Garay and Philip D. MacKenzie. Abuse-free multi-party contract signing. In *International Symposium on Distributed Computing*, volume 1693 of *Lecture Notes in Computer Science*, Bratislava, Slavak Republic, September 1999. Springer-Verlag.
11. Thomas A. Henzinger, Rupak Manjundar, Freddy Y.C. Mang, and Jean-François Raskin. Abstract interpretation of game properties. In *SAS 2000: Intertional Symposium on Static Analysis*, Lecture Notes in Computer Science. Springer-Verlag, 2000.
12. Steve Kremer and Jean-François Raskin. Game analysis of abuse-free contract signing. In *15th IEEE Computer Security Foundations Workshop*, Cape Breton, Canada, June 2002. IEEE Computer Society Press.
13. Vitaly Shmatikov and John Mitchell. Finite-state analysis of two contract signing protocols. *Theoretical Computer Science, special issue on Theoretical Foundations of Security Analysis and Design*, 283(2):419–450, 2002.