

A Formal Definition of Online Abuse-freeness^{*}

Ralf Küsters¹, Henning Schnoor², and Tomasz Truderung¹

¹ Universität Trier, Germany,
{kuesters, truderung}@uni-trier.de

² Christian-Albrechts-Universität zu Kiel, Germany,
schnoor@ti.informatik.uni-kiel.de

Abstract. Abuse-freeness is an important security requirement for contract-signing protocols. In previous work, Kähler, Küsters, and Wilke proposed a definition for *offline* abuse-freeness. In this work, we generalize this definition to *online* abuse-freeness and apply it to two prominent contract-signing protocols. We demonstrate that online abuse-freeness is strictly stronger than offline abuse-freeness.

Key words: contract signing, cryptographic protocols, formal verification

1.1 Introduction

In a (two-party) contract-signing protocol (see, e.g., [5, 4, 12]), two parties, A (Alice) and B (Bob), aim to exchange signatures on a contractual text that they previously agreed upon. In this paper, we consider *optimistic* contract-signing protocols. In such protocols, a trusted third party T (TTP), serving as an impartial judge, is not involved in every protocol run, but in case of a problem only.

A central security property for optimistic contract-signing, introduced in [12], is abuse-freeness: This property (formulated for the case of the honest signer Alice) requires that there is no state in a protocol run in which dishonest Bob (the prover) can convince an outside party, Charlie (the verifier), that the protocol is in an unbalanced state, i.e., a state in which Bob has both (i) a strategy to prevent Alice from obtaining a valid contract and (ii) a strategy to obtain a valid contract himself. In other words, if a contract-signing protocol is not abuse-free, then Alice can be misused by Bob to get leverage for another contract (with Charlie). Obviously, abuse-freeness is a highly desirable security property.

In [15], Kähler, Küsters, and Wilke presented the first rigorous and protocol-independent definition of abuse-free for (two-party) optimistic contract-signing. However, their definition focusses on an *offline* setting: Charlie is not actively involved in the protocol run and may receive a single message from Bob only, based on which he has to make his decision.

^{*} This work was partially supported by *Deutsche Forschungsgemeinschaft* (DFG) under Grant KU 1434/5-1.

The goal of this work is to generalize the definition from [15] to the setting in which Charlie may be *online*, i.e., may be actively involved in the protocol run, and to apply the definition to prominent contract-signing protocols.

Contribution of this Work. We propose a definition for *online* abuse-freeness, generalizing the definition of Kähler et al., who considered *offline* abuse-freeness. As theirs, our definition is protocol-independent. More precisely, we define two variants of online abuse-freeness: *Weak* abuse-freeness requires that there is no way for dishonest Bob to convince the verifier that the protocol is *currently* in an unbalanced state. As will be explained in Section 1.3, in the setting of *offline* abuse-freeness, every contract-signing protocol is weakly abuse-free. *Strong* abuse-freeness requires that Bob cannot even prove to Charlie that the protocol was in an unbalanced state at some point of the run.

We apply our definitions to two prominent contract signing protocols: a protocol by Asokan, Shoup, and Waidner [4] (ASW protocol), and one by Garay, Jakobsson, and P. MacKenzie [12] (GJM protocol). The latter was explicitly designed with abuse-freeness in mind. Depending on whether the verifier is allowed to eavesdrop on the network connection between the signers or on the channel between the signers and the TTP, and whether the initiator or responder in the protocol is dishonest, the protocols behave differently: We show that if the verifier can read the messages between the signers and the TTP, and the initiator is dishonest, then the ASW protocol is vulnerable to a very strong attack, i.e. it is not even weakly abuse-free. In this attack, the online aspect of our definition plays a crucial role, as the verifier “dictates” parts of the messages sent by the dishonest signer. In all other situations, ASW is weakly, but not strongly abuse-free. The GJM protocol shows a stronger resistance against abuse: It is weakly abuse-free in all situations, and strongly abuse free if the verifier cannot eavesdrop on the network channel between the signers.

Related Work. As mentioned above, Kähler et al. [15] introduced the first rigorous and protocol-independent definition of offline abuse-freeness.

Kremer et al. [16] analyzed the ASW and GJM protocol w.r.t. abuse-freeness using a finite-state model checking tool. They explicitly needed to specify the behavior of dishonest principals and which states are the ones that are convincing to Charlie.

Chadha et al. [6] introduce a stronger notion than abuse-freeness, namely *balance*: A protocol is balanced, if unbalanced states (see above) do not occur at all. Obviously, a balanced protocol is abuse-free as well. However, balance is very difficult to achieve. In fact, as shown by Chadha et al. [8], if principals are optimistic, i.e., they are willing to wait for messages of other parties, balance is impossible to achieve.

Shmatikov and Mitchell [18] employ the finite-state model checker Mur φ to automatically analyze contract-signing protocols. They approximate the notion of abuse-freeness by a notion similar to balance.

A cryptographic definition of the balance property was presented by Cortier, Küsters, and Warinschi in [10].

Aizatulin, Schnoor, and Wilke [3] introduced a contract signing protocol which satisfies a probabilistic notion of balance. Wang [19] introduces an abuse-free contract signing protocol based on the RSA signature scheme.

Structure of the Paper. In Section 1.2, we introduce our protocol model. The definition of abuse-freeness is then given in Section 1.3, and applied to the ASW and GJM protocols in Sections 1.4 and 1.5, respectively.

1.2 Protocol Model

In this section, following [17], we present a quite abstract symbolic protocol model. In this model, processes are represented as functions that from a sequence of input messages (the messages received so far) produce output messages. This model is the basis of our definition of abuse-freeness provided in the next section. We note, however, that the details of the model are not essential for the definition. The main motivation for using this model is brevity of presentation. We could as well have used another protocol model, such as the applied pi calculus [2].

1.2.1 Terms and Messages

Let Σ be some signature for cryptographic primitives (including a possibly infinite set of constants for representing participant names, etc.), $X = \{x_1, x_2, \dots\}$ be a set of variables, and Nonce be an infinite set of *nonces*, where the sets Σ , X , and Nonce are pairwise disjoint. For $N \subseteq \text{Nonce}$, the set T_N of *terms* over $\Sigma \cup N$ and X is defined as usual. Ground terms, i.e., terms without variables, represent messages. We assume some fixed equational theory associated with Σ and denote by \equiv the congruence relation on terms induced by this theory. The exact definition of Σ and the equational theory will depend on the cryptographic primitives used in the protocol under consideration. A simple example of a signature Σ_{ex} and its associated equational theory is provided in Figure 1.1. A term of the form $\text{sig}(\text{sk}(k), m)$ represents a message m signed using the (private) key $\text{sk}(k)$. Checking validity of such a signature is modeled by equation (1.1). The fact that signatures do not necessarily hide the signed message is expressed by equation (1.2). A term of the form $\{x\}_{\text{pk}(k)}^r$ represents the ciphertext obtained by encrypting x under the public key $\text{pk}(k)$ using randomness r . Decryption of such a term using the corresponding private key $\text{sk}(k)$ is modeled by equation (1.3). A term of the form $\langle x, y \rangle$ models the pairing of terms x and y . The components x and y of $\langle x, y \rangle$ can be extracted by applying the operators $\text{first}(\cdot)$ and $\text{second}(\cdot)$, respectively, as modeled by the equations (1.4) and (1.5). A term of the form $\text{hash}(m)$ represents the result of applying a hash function to a message m . Note that $\text{hash}(\cdot)$ is a free symbol, i.e. there is no equation involving this symbol in the given equational theory. For example, let \equiv_{ex} denote the congruence relation induced by the equational theory in Figure 1.1, then we have that $\text{dec}(\{a\}_{\text{pk}(k)}^r, \text{first}(\langle \text{sk}(k), b \rangle)) \equiv_{ex} a$.

$$\text{checksig}(\text{sig}(\text{sk}(k), m), \text{pk}(k)) = \text{T} \quad (1.1)$$

$$\text{extractmsg}(\text{sig}(\text{sk}(k), m)) = m \quad (1.2)$$

$$\text{dec}(\{x\}_{\text{pk}(k)}^r, \text{sk}(k)) = x \quad (1.3)$$

$$\text{first}(\langle x, y \rangle) = x, \quad (1.4)$$

$$\text{second}(\langle x, y \rangle) = y \quad (1.5)$$

Fig. 1.1. The equational theory associated with the signature $\Sigma_{ex} = \{\text{sig}(\cdot, \cdot), \langle \cdot, \cdot \rangle, \{\cdot\}, \text{T}, \text{checksig}(\cdot, \cdot), \text{extractmsg}(\cdot), \text{first}(\cdot), \text{second}(\cdot), \text{hash}(\cdot), \text{pk}(\cdot), \text{sk}(\cdot)\}$.

1.2.2 Event Sequences and Views

Let Ch be a set of *channels* (*channel names*). An *input/output event* is of the form $(c : m)$ and $(\bar{c} : m)$, respectively, for $c \in \text{Ch}$ and a message m (note that $\bar{c} \notin \text{Ch}$). A finite or infinite sequence of events is called an *event sequence*. For a sequence $\rho = (c_1 : m_1) (c_2 : m_2), \dots$ of input events, we denote by $\text{chan}(\rho)$ the sequence c_1, c_2, \dots of channels in this sequence. For $C \subseteq \text{Ch}$, we denote by $\rho|_C$ the subsequence of ρ containing only the events of the form $(c : m)$ with $c \in C$.

Let $\tau \in T_N$ be a term, which may contain variables x_1, x_2, \dots . Then, with ρ as above, we denote by $\tau[\rho]$ the message $\tau[m_1/x_1, m_2/x_2, \dots]$, where x_i is replaced by m_i . For example, assume that $\tau_{ex} = \text{dec}(x_1, \text{first}(x_2))$ and $\rho_{ex} = (c_1 : \{a\}_{\text{pk}(k)}^r), (c_2 : \langle \text{sk}(k), b \rangle)$. Then $\tau_{ex}[\rho_{ex}] = \text{dec}(\{a\}_{\text{pk}(k)}^r, \text{first}(\langle \text{sk}(k), b \rangle)) \equiv_{ex} a$.

Borrowing the notion of static equivalence from [2], we call two event sequences ρ and ρ' *statically equivalent w.r.t. a set $C \subseteq \text{Ch}$ of channels and a set $N \subseteq \text{Nonce}$ of nonces*, written $\rho \equiv_N^C \rho'$, if (i) $\text{chan}(\rho|_C) = \text{chan}(\rho'|_C)$ and (ii) for every $\tau_1, \tau_2 \in T_N$ we have that $\tau_1[\rho|_C] \equiv \tau_2[\rho|_C]$ iff $\tau_1[\rho'_|_C] \equiv \tau_2[\rho'_|_C]$. Intuitively, $\rho \equiv_N^C \rho'$ means that a party listening on channels C and a priori knowing the nonces in N cannot distinguish between the inputs received according to ρ and ρ' . We call the equivalence class of ρ w.r.t. \equiv_N^C , the (C, N) -*view* on ρ . For example, if a and b are different constants, k, k', r and r' are nonces, $C = \{c_1, c_2\}$, and $N = \emptyset$, then it is easy to see that $\rho_{ex}^1 = (c_1 : \{a\}_{\text{pk}(k)}^r), (c_2 : \langle \text{sk}(k'), b \rangle), (c_3 : \text{sk}(k))$ and $\rho_{ex}^2 = (c_1 : \{b\}_{\text{pk}(k)}^{r'}), (c_2 : \langle \text{sk}(k'), b \rangle)$ yield the same (C, N) -view w.r.t. \equiv_{ex} .

1.2.3 Processes

A process is, basically, a function that given a sequence of input events (representing the history so far) produces a sequence of output events. We require that a process behaves the same on inputs on which it has the same view. More precisely, a *process* is a tuple $\pi = (I, O, N, f)$ where

- (i) $I, O \subseteq \text{Ch}$ are finite sets of *input* and *output* channels, respectively,
- (ii) $N \subseteq \text{Nonce}$ is a set of *nonces used by π* ,

- (iii) f is a mapping which assigns a sequence $f(U) = (c_1 : \tau_1) \cdots (c_n : \tau_n)$ with $c_i \in O$ and $\tau_i \in T_N$ to each (I, N) -view U .

We note that (iii) guarantees that π performs the same computation on event sequences that are equivalent according to \equiv_N^I , and hence, on which π has the same view.

For an event sequence ρ , we write $\pi(\rho)$ for the output produced by π on input ρ . This output is $(c_1 : \tau_1[\rho']) \cdots (c_n : \tau_n[\rho'])$, where $\rho' = \rho|_I$ and $(c_1 : \tau_1) \cdots (c_n : \tau_n) = f(U)$ for the equivalence class U of ρ w.r.t. \equiv_N^I . For example, let $I = \{c_1, c_2\}$, $N = \emptyset$, U be the equivalence class of ρ_{ex}^1 , and assume that $f(U) = (c_4 : \langle x_1, \text{first}(x_2) \rangle)$. Then, $\pi(\rho_{ex}^1) = (c_4 : \langle \{a\}_{pk(k)}^r, \text{first}(\langle \text{sk}(k'), b \rangle) \rangle)$, which modulo \equiv_{ex} can be written equivalently as $(c_4 : \langle \{a\}_{pk(k)}^r, \text{sk}(k') \rangle)$ and $\pi(\rho_{ex}^2) = (c_4 : \langle \{b\}_{pk(k)}^{r'}, \text{first}(\langle \text{sk}(k'), b \rangle) \rangle)$, which modulo \equiv_{ex} can be equivalently written as $(c_4 : \langle \{b\}_{pk(k)}^{r'}, \text{sk}(k') \rangle)$. Note that since ρ_{ex}^1 and ρ_{ex}^2 yield the same (I, N) -view w.r.t. \equiv_{ex} , π performs the same transformation on ρ_{ex}^1 and ρ_{ex}^2 . We refer to I , O and N by I_π , O_π , and N_π , respectively. We note that the sets I_π and O_π do not have to be disjoint, i.e., π can send messages to itself. By $\text{Proc}(I, O, N)$ we denote the set of all processes π with $I_\pi \subseteq I$, $O_\pi \subseteq O$, and $N_\pi \subseteq N$.

1.2.4 Systems and Runs

A *system* S is a finite set of processes with disjoint sets of input channels and sets of nonces, i.e., $I_\pi \cap I_{\pi'} = \emptyset$ and $N_\pi \cap N_{\pi'} = \emptyset$, for distinct $\pi, \pi' \in S$. We will write $\pi_1 \parallel \cdots \parallel \pi_n$ for the system $\{\pi_1, \dots, \pi_n\}$.

Given a system S and a finite sequence s_0 of output events, a *run* ρ of S *initiated by* s_0 is a finite or infinite sequence of input and output events which evolves from s_0 in a natural way: An output event is chosen non-deterministically (initial from s_0). Once an output event has been chosen, it will not be chosen anymore later on. By definition of systems, there exists at most one process, say π , in S with an input channel corresponding to the output event. Now, π (if any) is given the input event corresponding to the chosen output event, along with all previous input events on channels of π . Then, π produces a sequence of output events as described above. Now, from these or older output events an output event is chosen non-deterministically, and the computation continues as before.

We emphasize that s_0 may induce many runs, due to the non-deterministic delivery of messages. In what follows, we assume *fair* runs, i.e., every output event in a run will eventually be chosen. A run is *complete* if it is either infinite or else all output events have been chosen at some point. For runs ρ, ρ' , we write $\rho \leq \rho'$, if ρ' is an extension of ρ , i.e., is obtained by continuing the run ρ .

1.2.5 Protocols

A *protocol* is a tuple $P = (A, in, out, nonce, s_0, \Pi)$, where

- (i) A is a finite set of *agent names*. An agent $a \in A$ has access to his/her nonces $\text{nonce}(a)$, input and output channels $\text{in}(a), \text{out}(a) \subseteq \text{Ch}$, respectively, such that $\text{nonce}(a) \cap \text{nonce}(a') = \emptyset$ and $\text{in}(a) \cap \text{in}(a') = \emptyset$, for $a \neq a'$,
- (ii) s_0 is a finite sequence of output events, the *initial output sequence*, for initializing parties,
- (iii) for every $a \in A$, $\Pi(a) \subseteq \text{Proc}(\text{in}(a), \text{out}(a), \text{nonce}(a))$ is the set of *programs* or *processes* of a . We will write $P(a)$ for $\Pi(a)$.

If $A = \{a_1, \dots, a_n\}$ and $\pi_i \in \Pi(a_i)$, then the system $(\pi_1 \parallel \dots \parallel \pi_n)$ is an *instance* of P . A *run* of P is a fair run of some instance of P initiated by s_0 . A *property* γ of P is a subset of runs of P .

We note that our model allows to express nondeterminism: To make a nondeterministic choice, a program can simply send two (or more) messages to itself, and change its behaviour depending on which message arrives first.

1.3 Online Abuse-Freeness

We define (online) abuse-freeness of a protocol $P = (A, \text{in}, \text{out}, \text{nonce}, s_0, \Pi)$ with respect to two distinct agents of P : the prover $\mathfrak{p} \in A$ and the verifier $\mathfrak{v} \in A$. Both agents are considered to be dishonest, and hence, the sets of programs of these agents will typically contain all possible processes, i.e., $P(\mathfrak{p}) = \text{Proc}(\text{in}(\mathfrak{p}), \text{out}(\mathfrak{p}), \text{nonce}(\mathfrak{p}))$ and $P(\mathfrak{v}) = \text{Proc}(\text{in}(\mathfrak{v}), \text{out}(\mathfrak{v}), \text{nonce}(\mathfrak{v}))$; these processes are only limited by their network interfaces, i.e., the set of input/output channels available to them.

Moreover, we define abuse-freeness of P with respect to two properties of P : γ^+ and γ^- . The property γ^+ is supposed to contain all the runs of P in which \mathfrak{p} obtains a valid contract from an honest signer a and γ^- is supposed to contain all runs where the honest signer a is prevented from obtaining a valid contract from \mathfrak{p} .

To define abuse-freeness, we first need to formalize the notion of an unbalanced run. Intuitively, a run of an instance of a protocol P is unbalanced with respect to the properties γ^+ and γ^- if \mathfrak{p} has both a strategy to achieve γ^+ (i.e., enforce a continuation of the run so that the overall run belongs to γ^+) and a strategy to achieve γ^- . In other words, in an unbalanced state, the prover can unilaterally determine the outcome of the protocol: i) obtain a signed contract from the honest signer a or ii) prevent a from obtaining a signed contract from \mathfrak{p} .

To model the choice made by the prover to either achieve γ^+ or γ^- , we introduce the following notation. We assume that the prover \mathfrak{p} has a distinct input channel $\text{ch}_{\text{choice}}$ which is not an output channel of any agent in the protocol P . Moreover, we assume that the events $(\overline{\text{ch}}_{\text{choice}} : 0)$ and $(\overline{\text{ch}}_{\text{choice}} : 1)$ belong to the initial event sequence s_0 of P . Intuitively, if in a run \mathfrak{p} receives 1 on $\text{ch}_{\text{choice}}$, then \mathfrak{p} will try to achieve γ^+ . If \mathfrak{p} receives 0 on $\text{ch}_{\text{choice}}$, then \mathfrak{p} will try to achieve γ^- . More precisely, a run ρ in which neither $(\text{ch}_{\text{choice}} : 0)$ nor $(\text{ch}_{\text{choice}} : 1)$ has been delivered is called *open*; intuitively, in such a run the prover has not yet

made a decision. Otherwise the run is called *closed*. In such a run, p tries to achieve γ^+ or γ^- depending on the message received; note that only the first message received on $\text{ch}_{\text{choice}}$ will set p 's goal.

In the following definition, given a finite open run ρ , we denote by $\rho^{(\text{ch}_{\text{choice}}:0)}$ the run obtained from ρ by delivering 0 on channel $\text{ch}_{\text{choice}}$ to p , i.e., in $\rho^{(\text{ch}_{\text{choice}}:0)}$ the prover p is now determined to achieve γ^- . The run $\rho^{(\text{ch}_{\text{choice}}:1)}$ is defined analogously. We say that a run ρ' is a *complete extension* of ρ if ρ' is an extension of ρ and is complete.

We are now ready to formally define unbalanced runs.

Definition 1. *Given an instance S of a protocol P as above with a prover p and two properties γ^+ and γ^- , we say that a finite open run ρ of S is unbalanced, if the following two conditions hold true:*

- (i) γ^- holds in every complete extension of $\rho^{(\text{ch}_{\text{choice}}:0)}$.
- (ii) γ^+ holds in every complete extension of $\rho^{(\text{ch}_{\text{choice}}:1)}$.

Now, intuitively, a protocol is abusive if a prover p can convince the verifier v that the current run is unbalanced. In other words, p can convince v that in the current run he, the prover, has a strategy to obtain a valid contract from the honest signer (and hence, close the deal) and a strategy to prevent the honest signer from obtaining a valid contract (and hence, cancel the deal). This may convince v to agree into a deal with p that for p is more profitable than the one with the honest signer. Thus, in an abusive protocol, p can take advantage of the honest signer.

Since we consider *online* abuse-freeness in this paper, we allow v to be actively involved in the protocol run. In particular, p and v can freely exchange messages during a run. For example, v could dictate (parts of the) messages p is supposed to send to the honest signer, and v could request to receive the private keys of p . The verifier v may even control some of the network traffic. However, this is not hard-wired in our definition. The power of p and v can be modeled in a flexible way in terms of the programs p and v may run and the network interface they have.

We will consider two forms of abuse-freeness, namely strong and weak abuse-freeness. In the strong form, p merely needs to convince v that the run was unbalanced at some point. In contrast, for the weak form, p needs to convince v that the run is unbalanced in the *current state* of the run. Since in the latter case, the task of p is harder, the latter form of abuse-freeness is weaker. It is desirable that a protocol is abuse-free in the strong sense since the fact that a run was and potentially still is unbalanced might already be sufficient incentive for v to agree into a deal with p .

In the formal definition of (online) abuse-freeness, we assume that the verifier v can *accept* a run by sending the message `accept` on the designated channel $\text{ch}_{\text{accept}}$, indicating that v is convinced that the run is/was unbalanced. We say that a finite run is *freshly accepted*, if the message `accept` is sent by v in the last step of this run.

We also use the following notation in the definition of abuse-freeness: Let $P = (A, in, out, nonce, s_0, \Pi)$ be a protocol. For a program $v \in \Pi(\mathbf{v})$ of the verifier, we write $P|_v$ for the protocol that coincides with P except that the set $\Pi(\mathbf{v})$ of programs of \mathbf{v} is restricted to $\{v\}$. In particular, in every instance of $P|_v$ the verifier runs the program v .

We are now ready to define (online) abuse-freeness. We start with the strong form of abuse-freeness.

Definition 2. Let $P = (A, in, out, nonce, s_0, \Pi)$ with $\mathbf{p}, \mathbf{v} \in A$. Let γ^+ and γ^- be properties of P . Then, the protocol P is called (γ^+, γ^-) -*abusive* w.r.t. the prover \mathbf{p} and the verifier \mathbf{v} , if there is a program $v \in \Pi(\mathbf{v})$ of \mathbf{v} such that the following conditions are satisfied:

- (i) If an open run ρ of $P|_v$ is accepted by \mathbf{v} , then there is an unbalanced run ρ' with $\rho' \leq \rho$.
- (ii) There exists an open, freshly accepted, unbalanced run ρ of $P|_v$.

The protocol P is (*strongly*) (γ^+, γ^-) -*abuse-free* w.r.t. \mathbf{p} and \mathbf{v} , if P is not (γ^+, γ^-) -abusive w.r.t. \mathbf{p} and \mathbf{v} .

Condition (i) in the above definition says that if \mathbf{v} accepts a run, i.e., is convinced that the run was unbalanced at some point, then this is in fact the case. Note that according to the definition of unbalanced runs, \mathbf{v} may help \mathbf{p} to achieve his goals (γ^+ or γ^-). One could as well consider a variant where \mathbf{p} has to achieve these goals against \mathbf{v} (and in fact, our negative results, presented in Sections 1.4 and 1.5, use a prover that works without the help of the verifier). However, this would make the definition only weaker. We note that it would not make sense to consider closed runs in Condition (i): The definition of unbalanced runs only applies to open runs. Moreover, the restriction to open runs does not limit the power of any agent.

While Condition (i) is the core of the above definition, it would not make sense without Condition (ii): A verifier who never accepts a run would satisfy Condition (i) trivially. Moreover, a verifier who only accepts runs which are not unbalanced anymore would potentially also suffice to meet Condition (i). By Condition (ii) we require that the strategy of the verifier for accepting a run is reasonable in the sense that there is at least one run which is accepted and which is still unbalanced.

Altogether the above definition says that a protocol is abuse-free if there is no program a verifier could run which i) reliably tells, for any dishonest prover (which the verifier does not trust), whether a run was unbalanced and ii) accepts an actual unbalanced run.

We note that the above definition is possibilistic. It does, for example, not take into account the probability with which unbalanced runs occur or a verifier accepts an (unbalanced) run. As mentioned in the introduction, a cryptographic, in particular probabilistic definition of the balance property, which is stronger than abuse-freeness in that it does not require a dishonest signer to convince an outside party of the fact that he is in an unbalanced state, was presented in [10].

Given Definition 2, it is now straightforward to define weak abuse-freeness:

Definition 3. Let $P, \mathbf{p}, \mathbf{v}, \gamma^+$, and γ^- be given as in Definition 2. Then protocol P is *strongly* (γ^+, γ^-) -*abusive* w.r.t. \mathbf{p} and \mathbf{v} , if there is a program v such that the following conditions are satisfied:

- (i) If an open run ρ of $P|_v$ is freshly accepted, then ρ is unbalanced,
- (ii) There exists an open, freshly accepted, unbalanced run ρ of $P|_v$.

The protocol P is *weakly* (γ^+, γ^-) -*abuse-free* w.r.t. \mathbf{p} and \mathbf{v} , if P is not strongly (γ^+, γ^-) -abusive w.r.t. \mathbf{p} and \mathbf{v} .

This notion differs from the (strong) abuse-freeness only in Condition (i): Now we require that the accepted run *is* unbalanced, not only that it was unbalanced at some previous point. Clearly, (strong) abuse-freeness implies weak abuse-freeness.

Note that a notion like weak abuse-freeness does not make sense in the offline setting considered in [15]: If the verifier receives only a single message from the prover, this message can only prove that the protocol was in an unbalanced state at some point during the protocol run; since the prover may withhold that evidence for as long as he wishes, it does not prove that the current state is unbalanced.

1.4 The ASW Protocol

In this section, we study abuse-freeness of the contract-signing protocol proposed by Asokan, Shoup, and Waidner (ASW protocol) in [4].

In [15], it has been shown (in a *synchronous* communication model without optimistic honest parties, see below) that the ASW protocol is *offline* abusive. Not surprisingly, the protocol is also abusive in the online setting. More precisely, we show that the protocol is weakly abusive. Interestingly, we can show that the protocol is, in some cases, even *strongly abusive*. For this attack, it is crucial that the verifier is online, i.e., can interact with the prover. In fact, the verifier will dictate part of the message the prover sends to the honest signer.

1.4.1 Description of the Protocol

The ASW protocol assumes the following scenario: Alice and Bob want to sign a contract and a TTP is present. The following two types of messages, the *standard contract* (SC) and the *replacement contract* (RC), will be recognized as valid contracts between Alice and Bob with contractual text text : $SC = \langle me_1, N_A, me_2, N_B \rangle$ and $RC = \text{sig}(\text{sk}(k_t), \langle me_1, me_2 \rangle)$ where N_A and N_B stand for nonces, $me_1 = \text{sig}(\text{sk}(k_a), \langle A, B, \text{text}, \text{hash}(N_A) \rangle)$, and $me_2 = \text{sig}(\text{sk}(k_b), \langle me_1, \text{hash}(N_B) \rangle)$, with $\text{sk}(k_t)$, $\text{sk}(k_a)$, and $\text{sk}(k_b)$ denoting the private keys of the TTP, Alice, and Bob, respectively. In addition to SC and RC , the variants of SC and RC which one obtains by exchanging the roles of A and B are regarded as valid contracts.

There are three interdependent parts to the protocol: an exchange protocol, an abort protocol, and a resolve protocol. The *exchange protocol* consists of four steps, which, in Alice-Bob notation, are displayed in Fig. 2. The first two messages, me_1 and me_2 , serve as respective *promises* of Alice and Bob to sign the contract, and N_A and N_B serve as *contract authenticators*: After they have been revealed, Alice and Bob can compose the standard contract, SC .

The *abort protocol* is run between Alice and the TTP and is used by Alice to abort the contract signing process when she does not receive Bob's promise. Alice will obtain (from the TTP) an abort receipt or, if the protocol instance has already been resolved (see below), a replacement contract. The first step is $A \rightarrow T: ma_1$, where $ma_1 = \text{sig}(\text{sk}(k_a), \langle \text{aborted}, me_1 \rangle)$ is Alice's *abort request*; the second step is the TTP's reply, which is either $\text{sig}(\text{sk}(k_t), \langle \text{aborted}, ma_1 \rangle)$, the *abort receipt*, if the protocol has not been resolved, or the replacement contract, RC .

The *resolve protocol* can be used by Alice and Bob to resolve the protocol, which either results in a replacement contract or, if the protocol has already been aborted, in an abort receipt. When Bob runs the protocol (because Alice has not sent her contract authenticator yet), the first step is $B \rightarrow T: \langle me_1, me_2 \rangle$; the second step is the TTP's reply, which is either the abort receipt $\text{sig}(\text{sk}(k_t), \langle \text{aborted}, ma_1 \rangle)$, if the protocol has already been aborted, or the replacement contract, RC . The same protocol (with roles of A and B exchanged) is also used by Alice.

1.4.2 Modeling

Our modeling of the ASW protocol uses the equational theory presented in Section 1.2 (however, without encryption, which is not used in the protocol). We consider, besides the regular protocol participants of the protocol—Alice, Bob, and the trusted third party—two additional parties, the verifier and a key distribution center. We will consider four cases depending on (a) which signer (Alice or Bob) is dishonest and plays the role of the prover and (b) which part of the network is controlled by the verifier.

In each case we assume that the honest signer is *optimistic* in the sense that he/she only contacts the TTP if the dishonest signer allows the honest signer to do so. In other words, the dishonest signer can buy himself as much time as he needs, before the honest signer contacts the TTP. This assumption, also made in [8], seems realistic. In any case, it only makes the dishonest party more powerful, and hence, strengthens our positive results.

Let $P_{ASW-Net}^A$ denote the specification of the ASW protocol, as a protocol in the sense of our definition (see Section 1.2.5), with dishonest Alice and honest Bob, where the verifier can eavesdrop on (but not block) the network traffic between Alice and Bob. Analogously, $P_{ASW-Net}^B$ denotes the protocol with dishonest Bob and honest Alice, where again the verifier can eavesdrop on the network traffic between Alice and Bob. Let $P_{ASW-TTP}^A$ ($P_{ASW-TTP}^B$) be the protocols with

$A \rightarrow B : me_1$
 $B \rightarrow A : me_2$
 $A \rightarrow B : N_A$
 $B \rightarrow A : N_B$

Fig. 1.2: ASW exchange protocol

dishonest Alice (Bob) and honest Bob (Alice), where the verifier can eavesdrop on (but not block) the communication between the signers and the TTP.

In the modeling of these protocols (see below), we allow an honest signer to *not* be engaged in the protocol run. This is of course realistic; also, otherwise the initial state of the protocol would already be unbalanced before a signer has committed to the contract. To model this, we assume that an honest signer decides nondeterministically (see end of Section 1.2.5) as to whether he/she will participate in the protocol run.

More formally, the set of programs of the protocol participants are defined as follows:

Key Distribution Center. The set of programs for this party consists of exactly one program, which generates key pairs (using its set of nonces) for all other parties. Private keys, modeled as terms of the form $\text{sk}(k)$, where k is a nonce, are sent, via dedicated channels, only to the respective parties. Public keys, modeled as terms of the form $\text{pk}(k)$, are distributed to all parties, including the verifier. Honest parties will first wait to receive their public/private key pair and the public keys of the other protocol participants. In the specification of honest parties below this is assumed implicitly.

Dishonest parties (prover and verifier). The sets of programs of dishonest parties contain all the possible processes, only constrained by the network configuration and, possibly, some additional constraints, as described below. We allow the prover and verifier to communicate directly with each other via a direct (asynchronous) channel.

Network configuration. In the protocols $P_{ASW-TTP}^A$ and $P_{ASW-TTP}^B$ (in which the verifier can eavesdrop on messages between the signers and the TTP) we assume that the messages that Alice and Bob want to send to the TTP are routed through the verifier. We require the verifier to forward these messages to the recipient, i.e. we restrict the set of program of the verifier to those programs which comply with this constraint. However, we assume direct (asynchronous) channels between Alice and Bob.

Similarly, in $P_{ASW-TTP}^A$ and $P_{ASW-TTP}^B$ messages between Alice and Bob are routed through the verifier, who, as above, can only eavesdrop on these messages. Message between the signers and the TTP can be sent via direct (asynchronous) channels.

TTP. The set of programs of TTP consists of only one program (process), namely the one that performs exactly the steps defined by the protocol as described in Section 1.4.1.

Honest Alice or Bob. The set of programs of Alice in $P_{ASW-TTP}^A$ and $P_{ASW-Net}^A$ consists of only one program, namely the one described in Section 1.4.1. As mentioned above, at the beginning Alice first nondeterministically chooses whether to participate in the contract signing. Also, she contacts the TTP only if she receives a message from Bob that she is allowed to contact the TTP. The case of honest Bob is analogous.

Remark 1. One could also study the case where the verifier can eavesdrop on all channels. However in this case, both the ASW and GJM protocols clearly are strongly abusive, since the verifier always knows the exact stage of the protocol run.

1.4.3 Security Analysis

We define γ^+ as the set of all runs where the prover is able to construct the standard contract or has received the replacement contract. Analogously, γ^- consists of those runs in which the honest signer is not able to construct the standard contract and has not received the replacement contract. For the ASW protocol, we prove the following results.

Theorem 1. *The protocol $P_{ASW-TTP}^A$ is not weakly (γ^+, γ^-) -abuse-free (and hence also not abuse-free). $P_{ASW-TTP}^B$, $P_{ASW-Net}^A$, and $P_{ASW-Net}^B$ are weakly (γ^+, γ^-) -abuse-free but not abuse-free.*

A proof sketch of this theorem is presented in Appendix A. We note that the first result exhibits a particularly devastating attack, which makes heavy use of the fact that the verifier is an online agent. This result shows that under certain conditions the ASW protocol is not even weakly abuse-free. The above results also show that weak abuse-freeness is a much weaker security property than strong abuse-freeness.

Remark 2. The proofs for $P_{ASW-Net}^A$ and $P_{ASW-Net}^B$ easily carry over to the case when the verifier not only eavesdrops on the channels between Alice and Bob, but also controls these channels.

1.5 The GJM Protocol

In [15], it has been shown that, in a *synchronous* communication model, the GJM protocol is *offline* abuse-free. In this section, we show that whether it is *online* abuse-free depends on assumptions about what part of the network the verifier can eavesdrop on. In particular, we show that in some cases the GJM protocol is not online abuse-free, which, again, illustrates the fact that online abuse-freeness is stronger than offline abuse-freeness.

1.5.1 Informal Description and Model of the Protocol

The structure of the GJM protocol is the same as the one of the ASW protocol. However, the actual messages exchanged are different. In particular, the exchange protocol of the GJM protocol the first two messages are so-called private contract signatures (PCS) [12] and the last two messages are actual signatures (obtained by converting the private contract signatures into universally verifiable signatures).

$$\text{pcsver}(w, \text{pk}(x), \text{pk}(y), \text{pk}(z), \text{pcs}(u, \text{sk}(x), w, \text{pk}(y), \text{pk}(z))) = \text{pcsok}, \quad (1.6)$$

$$\text{pcsver}(w, \text{pk}(x), \text{pk}(y), \text{pk}(z), \text{fake}(u, \text{sk}(y), w, \text{pk}(x), \text{pk}(z))) = \text{pcsok}, \quad (1.7)$$

$$\text{sver}(w, \text{pk}(x), \text{pk}(z), \text{sconv}(u, \text{sk}(x), \text{pcs}(v, \text{sk}(x), w, \text{pk}(y), \text{pk}(z)))) = \text{sok}, \quad (1.8)$$

$$\text{tpver}(w, \text{pk}(x), \text{pk}(z), \text{tpconv}(u, \text{sk}(z), \text{pcs}(v, \text{sk}(x), w, \text{pk}(y), \text{pk}(z)))) = \text{tpok}. \quad (1.9)$$

Fig. 1.3. Equations for private contract signatures.

For the GJM protocol we consider the signature $\Sigma_{GJM} = \{ \text{sig}(\cdot, \cdot, \cdot), \text{sigcheck}(\cdot, \cdot, \cdot), \text{pk}(\cdot), \text{sk}(\cdot), \text{fake}(\cdot, \cdot, \cdot, \cdot, \cdot), \text{pcs}(\cdot, \cdot, \cdot, \cdot, \cdot), \text{pcsver}(\cdot, \cdot, \cdot, \cdot, \cdot), \text{sconv}(\cdot, \cdot, \cdot), \text{tpconv}(\cdot, \cdot, \cdot), \text{sver}(\cdot, \cdot, \cdot, \cdot), \text{tpver}(\cdot, \cdot, \cdot, \cdot), \langle \cdot, \cdot \rangle, \text{first}(\cdot), \text{second}(\cdot), A, B, T, \text{text}, \text{initiator}, \text{responder}, \text{ok}, \text{pcsok}, \text{sok}, \text{tpok}, \text{aborted} \}$.

The equational theory for GJM contains, in addition to the equations for pairing and signatures, equations for modeling private contract signatures, as depicted in Figure 1.3. A term of the form $\text{pcs}(u, \text{sk}(x), w, \text{pk}(y), \text{pk}(z))$ stands for a PCS computed by x (with $\text{sk}(x)$) involving the text w , the party y , and the TTP z , while u models the random coins used to compute the PCS. Everybody can verify the PCS with the public keys involved (equation (1.6)), but cannot determine whether the PCS was computed by x or y (equation (1.7)): instead of x computing the “real” PCS, y could have computed a “fake” PCS which would also pass the verification with pcsver . Using sconv and tpconv , see (1.8) and (1.9), a “real” PCS can be converted by x and the TTP z , respectively, into a universally verifiable signature (verifiable by everyone who possesses $\text{pk}(x)$ and $\text{pk}(z)$).

We study the version of the GJM protocol with the modification proposed in [18] to obtain fairness. In the protocol, the following messages are exchanged: The initial messages containing the private contract signatures are $me_1 = \text{pcs}(u, \text{sk}(A), \text{contract}, \text{pk}(B), \text{pk}(TTP))$ and $me_2 = \text{pcs}(u', \text{sk}(B), \text{contract}, \text{pk}(A), \text{pk}(TTP))$, where $\text{sk}(A)$, $\text{pk}(A)$, $\text{sk}(B)$, $\text{pk}(B)$, and $\text{pk}(TTP)$ are the private and public keys of Alice, Bob, and the TTP. The abort request sent by Alice is of the form $ma_1 = \text{sig}(w, \text{sk}(A), \langle \text{contract}, A, B, \text{aborted} \rangle)$, where w are random coins (for the GJM protocol, we consider randomized signatures). The resolve request sent by Alice is $\langle me_1, me_2 \rangle$, the resolve request from Bob is $\langle me_2, me_1 \rangle$. As mentioned earlier, the structure of the protocol is the same as for the ASW protocol (see Section 1.4).

1.5.2 Security Analysis

We study the cases $P_{GJM-TTP}^A$, $P_{GJM-TTP}^B$, $P_{GJM-Net}^A$ and $P_{GJM-Net}^B$, which are defined analogously to the case of ASW (see Section 1.4.2). The properties γ^+ and γ^- are also defined analogously to the case of the ASW protocol (see Section 1.4.3).

Theorem 2. 1. $P_{GJM-TTP}^A$ and $P_{GJM-TTP}^B$ are (γ^+, γ^-) -abuse-free.
2. $P_{GJM-Net}^A$ and $P_{GJM-Net}^B$ are weakly (γ^+, γ^-) -abuse-free but not (γ^+, γ^-) -abuse-free.

As made precise by this theorem, abuse-freeness of the GJM protocol in the online setting depends on the assumptions about what part of the network the verifier can eavesdrop on. In the offline case, the verifier was not allowed to eavesdrop on any part of the network (and of course, was also not allowed to be actively involved in the protocol run). Therefore, and just as in the case of the ASW protocol, our positive results are stronger than those shown for offline abuse-freeness. Conversely, our negative results exhibit the extra power of online verifiers. A proof sketch of Theorem 2 is presented in Appendix B.

References

1. M. Abadi and V. Cortier. Deciding knowledge in security protocols under equational theories. In *ICALP 2004*, volume 3142 of *LNCS*, pages 46–58. Springer, 2004.
2. M. Abadi and C. Fournet. Mobile Values, New Names, and Secure Communication. In *POPL 2001*, pages 104–115. ACM Press, 2001.
3. M. Aizatulin, and H. Schnoor, and Th. Wilke. Computationally Sound Analysis of a Probabilistic Contract Signing Protocol. In *ESORICS 2009*, volume 5789 of *LNCS*, pages 571–586, Springer, 2009.
4. N. Asokan, V. Shoup, and M. Waidner. Asynchronous protocols for optimistic fair exchange. In *IEEE Symposium on Research in Security and Privacy*, pages 86–99, 1998.
5. M. Ben-Or, O. Goldreich, S. Micali, and R.L. Rivest. Fair protocol for signing contracts. *IEEE Transactions on Information Theory*, 36(1):40–46, 1990.
6. R. Chadha, M.I. Kanovich, and A. Scedrov. Inductive methods and contract-signing protocols. In *CCS 2001*, pages 176–185. ACM Press, 2001.
7. R. Chadha, S. Kremer, and A. Scedrov. Formal analysis of multi-party contract signing. In *CSFW 2004*, pages 266–279. IEEE Computer Society Press, 2004.
8. R. Chadha, J.C. Mitchell, A. Scedrov, and V. Shmatikov. Contract Signing, Optimism, and Advantage. In *CONCUR 2003*, volume 2761 of *LNCS*, pages 361–377. Springer, 2003.
9. Y. Chevalier and M. Rusinowitch. Combining Intruder Theories. In *ICALP 2005*, volume 3580 of *LNCS*, pages 639–651. Springer, 2005.
10. V. Cortier, R. Küsters, and B. Warinschi. A Cryptographic Model for Branching Time Security Properties - The Case of Contract Signing Protocols. In *ESORICS 2007*, volume 4734 of *LNCS*, pages 422–437, Springer, 2007.
11. D. Dolev and A.C. Yao. On the Security of Public-Key Protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
12. J.A. Garay, M. Jakobsson, and P. MacKenzie. Abuse-free optimistic contract signing. In *CRYPTO'99*, volume 1666 of *LNCS*, pages 449–466. Springer-Verlag, 1999.
13. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems. In *STOC 1985*, pages 291–304. ACM Press, 1985.
14. D. Kähler, R. Küsters, and Th. Wilke. Deciding Properties of Contract-Signing Protocols. In *STACS 2005*, volume 3404 of *LNCS*, pages 158–169. Springer-Verlag, 2005.
15. D. Kähler, R. Küsters, and Th. Wilke. A Dolev-Yao-based Definition of Abuse-free Protocols. Technical report, IFI 0607, CAU Kiel, Germany, 2006. Available from <http://www.informatik.uni-kiel.de/reports/2006/0607.html>

16. S. Kremer and J.-F. Raskin. Game analysis of abuse-free contract signing. In *CSFW 2002*, pages 206–220. IEEE Computer Society, 2002.
17. R. Küsters and T. Truderung. An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In *Security and Privacy 2009*, pages 251–266. IEEE Computer Society, 2009.
18. V. Shmatikov and J.C. Mitchell. Finite-state analysis of two contract signing protocols. *Theoretical Computer Science (TCS), special issue on Theoretical Foundations of Security Analysis and Design*, 283(2):419–450, 2002.
19. G. Wang. An Abuse-Free Fair Contract-Signing Protocol Based on the RSA Signature. *IEEE Transactions on Information Forensics and Security*, 5(1):158–168, 2010.

A Proof Sketch of Theorem 1 (ASW Protocol)

Proof for $P_{ASW-TTP}^A$: In this case, Alice is the prover p , and the verifier v is able to observe the network traffic between the signers and the TTP. We describe the program v for the verifier v that satisfies the definition of strong abusiveness. This program is defined as follows:

- (1) Choose a nonce N_v .
- (2) Send $\text{hash}(N_v)$ to A .
- (3) Wait for the message $me'_2 = \text{sig}(\text{sk}(k_b), \langle me'_1, \text{hash}(N_B) \rangle)$, where $me'_1 = \text{sig}(\text{sk}(k_a), \langle A, B, \text{text}, \text{hash}(N_v) \rangle)$ (me'_2 is like the message me_2 , as specified by the protocol, but with N_v instead of N_A).
- (4) Accept if and only if neither Alice nor Bob has contacted the TTP (recall that the verifier can see this).

We claim that, for this program v , conditions (i) and (ii) of Definition 3 are satisfied. For showing that condition (i) holds, let us assume that ρ is a freshly accepted open run of P_v . We need to show that ρ is unbalanced. For this we give strategies for Alice to achieve γ^- and γ^+ , starting from ρ . The strategy for γ^- is to run the abort subprotocol with the TTP, the strategy for γ^+ is to run the resolve subprotocol. Since ρ is accepted, Alice has received me'_2 and therefore she can run these subprotocols. Furthermore, because the verifier accepted the protocol run, we know that neither Alice nor Bob has contacted the TTP. Therefore the abort/resolve subprotocol run by Alice will succeed, yielding γ^- and γ^+ , respectively, in every complete run obtained with these strategies (note that, in this case, since Bob is optimistic, he will not run neither the abort nor the resolve subprotocol).

For showing that condition (ii) holds, it is enough to show that there exists an accepted run of P_v , this run is then unbalanced due to condition (i) as just proved. As the programs of Bob and TTP are determined by the protocol, and the program for the verifier is specified above, it remains to define the program for Alice.

This program, having received the nonce N_v from the verifier, sends the message me'_1 , as defined above, to Bob, and then forwards to the prover Bob's response (if Bob does indeed reply, i.e., decides to participate in the protocol). For this instance, we consider a run in which Bob decides to perform the protocol.

Since Bob follows the protocol, he replies with the message me'_2 , which Alice then forwards to the verifier. Since Alice and Bob did not contact the TTP (Alice by construction, and Bob since he is optimistic), the verifier will accept the protocol run.

Proof for $P_{ASW-TTP}^B$: This case corresponds to the situation where Bob is the prover, and the verifier can eavesdrop on the channel between the signers and the TTP. We first show that $P_{ASW-TTP}^B$ is weakly abuse-free, i.e., not strongly abusive. Hence assume that the protocol is strongly abusive, and let v be a corresponding program of the verifier v . By definition, there is a freshly accepted unbalanced run ρ of $P|_v$. We can assume w.l.o.g. that Alice, in this run, nondeterministically chose to sign the contract and initiated the exchange by sending me_1 to Bob (if v accepted a run in which Alice had not initiated this exchange, then v clearly would also accept runs that are not unbalanced).

We will construct a run ρ' of $P|_v$ which is not unbalanced, but, from the verifier's point of view, is indistinguishable from ρ , and therefore also accepted. In ρ' , Bob runs the program which, having obtained me_1 , completes the standard protocol exchange with Alice (which means that the run is not unbalanced; note that this exchange cannot be noticed by the TTP) and then behaves exactly like the program used in ρ (in particular, all the messages that Bob sends to the verifier are exactly the same in both runs ρ and ρ'). Hence the protocol is weakly abuse-free.

We now show that the protocol is not abuse-free, i.e., that there is a program v of the verifier v satisfying the necessary conditions.

The program v simply expects a message conforming to the format of me_1 , i.e., a message of the form $\text{sig}(\text{sk}(k_a), \langle A, B, \text{text}, \text{hash}(t) \rangle)$ for some term t , and then accepts. Obviously, there is a freshly accepted protocol run that is unbalanced: When Alice starts the protocol run, and Bob forwards the message me_1 to the verifier, the run is accepted. It is also unbalanced, since Bob can obtain a contract by performing the resolve protocol with the TTP, and abort the protocol by simply stopping to send messages.

Now let ρ be an open protocol run accepted by the verifier. Since ρ is accepted, the prover did send a message to the verifier that contains Alice's signature. Since Alice follows the protocol run, this means that Alice did send the message me_1 , and the prover Bob received it. Hence at one point during the protocol run, Bob was able to send a resolve request to the TTP to obtain a contract (by which he achieves γ^+), as well as to abort by simply not sending any messages (by which he enforces γ^-). Hence there is an unbalanced prefix of ρ .

Proof for $P_{ASW-Net}^A$: Here we consider the situation where Alice is the prover, and the verifier can eavesdrop on the network channel between Alice and Bob, but does not have access to the communication between the signers and the TTP. We first show that the protocol is abusive. For this, let v be the program of the verifier that accepts any protocol run in which Bob did send a message to Alice (note that the verifier has read-access to the channel between Alice and Bob). Obviously, there is a freshly accepted run that is unbalanced, namely the run

in which the prover Alice sends a message of the correct form me_1 to Bob, and Bob responds with a well-formed me_2 .

Now let ρ be an accepted open run. By definition of v , Bob did send a message to Alice. Since Bob follows the protocol, this message is of the form $me'_2 = \text{sig}(\text{sk}(k_b), \langle me'_1, \text{hash}(N_B) \rangle)$, where me'_1 is the message that, from the point of view of Bob (and thus also from the point of view of the TTP), conforms to the protocol specification, i.e., $me'_1 = \text{sig}(\text{sk}(k_a), \langle A, B, \text{text}, m \rangle)$ for some message m . Hence after receiving me'_2 , Alice is able to form a resolve request to the TTP, and thus has a strategy to obtain a contract. She also has an abort-strategy by simply quitting the protocol. This strategy will be successful since Bob is optimistic. Hence there is an unbalanced prefix of ρ .

It remains to show that $P_{ASW-Net}^A$ is weakly abuse-free, i.e., not strongly abusive. Hence assume that there is a program v of the verifier satisfying the definition of strong abusiveness. Again, then there is an instance S of P_v and an unbalanced, freshly accepted protocol run ρ . Without loss of generality, in the run ρ , Bob sent a message to Alice, which v can read since he has access to the corresponding channel (if v accepted a run in which Bob had not sent a message, then v clearly also would accept runs that are not unbalanced). In particular, Alice has access to messages me_1 and me_2 , that allow her to form a resolve request for the TTP. Hence a run ρ' of an instance S' obtained by letting Alice, additionally, send a resolve request for the TTP is indistinguishable for v , and thus he accepts this run which clearly is not unbalanced.

Proof for $P_{ASW-Net}^B$: In this case Bob is the prover, and the verifier again has read access to the network channel between Alice and Bob. The proof of abusiveness is almost the same as in case (3): Since the verifier observes network traffic between Alice and Bob, he knows whether Alice has initiated the protocol run by sending me_1 , which leads to an unbalanced prefix of the protocol run. Also, the proof for weak abuse-freeness is very similar to the one in case (3) above: Again, without loss of generality, we can assume that a program v of the verifier v only accepts runs in which Alice did send a message to Bob, which allows Bob to send a resolve request to the TTP without the verifier noticing this, hence if there is an accepted run, there also is one that is not unbalanced. \square

B Proof Sketch of Theorem 2 (GJM Protocol)

We first show that $P_{GJM-TTP}^A$ and $P_{GJM-TTP}^B$ are abuse-free. In this case the verifier has read-access to the channels between the signers and the TTP. Assume that the protocols are not abuse-free, and let v be a program for the verifier that satisfies the definition. Then there must be an open protocol run ρ that is freshly accepted and unbalanced. In particular, since the run is unbalanced, no request has been sent to the TTP, and the prover received a message from the honest signer, but did not complete the exchange protocol with the honest signer. Additionally, since the verifier accepted the protocol run, the prover must have sent to v the private contract signature produced by the honest signer (obviously, the

verifier cannot accept a run, if he is not provided with any evidence that the honest signer has committed to the contract). Now, from the equational theory for private contract signatures it follows that if the prover sends to the verifier the fake signature instead, the verifier accepts as well. Since the latter message can be constructed by the prover without participation of the honest signer, this means that there is an accepted open protocol run in which the honest signer did not even start a protocol session. This shows that there is an accepted run that does not have an unbalanced prefix. Therefore, the protocol is not weakly abusive, i.e., it is strongly abuse-free if the verifier can read network traffic between the signers and the TTP but not the traffic between the signers.

We now consider the cases $P_{GJM-Net}^A$ and $P_{GJM-Net}^B$, which are very similar to their ASW counterparts (see Theorem 1). Since the verifier can observe network traffic between the signers, he knows, in particular, whether the honest signer has sent his private contract signature. Hence if the verifier simply accepts any protocol run in which the honest signer did send a message to the prover, then there is a freshly accepted open protocol run that is unbalanced, and every accepted protocol run has an unbalanced prefix (note again that the honest signer is optimistic and, thus, does not send an abort request, if not instructed by the prover), hence the protocol is abusive.

To prove that these cases do not lead to strong abusiveness, we use the same argument as in ASW: There is no way for the prover to convince the verifier that he has not performed the resolve protocol with the TTP, hence the verifier cannot be sure that the *current* state of the protocol run is unbalanced.