

On Defendability of Security Properties

Wojciech Jamroga

Computer Science and Communication
& Interdisciplinary Centre for Security, Reliability, and Trust,
University of Luxembourg
wojtek.jamroga@uni.lu

Matthijs Melissen

Computer Science and Communication,
University of Luxembourg
& School of Computer Science,
University of Birmingham
m.melissen@cs.bham.ac.uk

Henning Schnoor

Arbeitsgruppe Theoretische Informatik,
University of Kiel
henning.schnoor@email.uni-kiel.de

We study the security of interaction protocols when incentives of participants are taken into account. We begin by formally defining correctness of a protocol, given a notion of rationality and utilities of participating agents. Based on that, we propose how to assess security when the precise incentives are unknown. Then, the security level can be defined in terms of *defender sets*, i.e., sets of participants who can effectively “defend” the security property as long as they are in favor of the property.

We present some theoretical characterizations of defendable protocols under Nash equilibrium, first for bijective games (a standard assumption in game theory), and then for games with non-injective outcomes that better correspond to interaction protocols. Finally, we apply our concepts to analyze fairness in the ASW contract-signing protocol.

1 Introduction

Interaction protocols are ubiquitous in multi-agent systems. Protocols can be modeled as games, since every participant in the protocol has several strategies that she can employ. From a game-theoretic perspective, protocols are an interesting class of games since they have a *goal*, i.e., a set of outcomes that are preferred by the designer of the protocol. *Security protocols* use cryptography to enforce their goals against any possible behavior of participants. Such a protocol is deemed correct with respect to its goal if the goal is achieved in all runs where a predefined subset of players follows the protocol.

We point out that this definition of correctness can be too strong, since violation of the goal may be achievable only by irrational responses from the other players. On the other hand, the definition may also prove too weak when the goal can be only achieved by an irrational strategy of agents supporting the goal, in other words: one that they should never choose to play. To describe and predict rational behavior of agents, game theory has proposed a number of *solution concepts* [13]. Each solution concept captures some notion of rationality which may be more or less applicable in different contexts. We do not fix a particular solution concept, but consider it to be a parameter of the problem.

Our main contributions are the following. First, in Section 3.1, we define a parametrized notion of *rational correctness* for security protocols, where the parameter is a suitable solution concept. Secondly, based on this notion, we define a concept of *defendability of security* in a protocol, where the security property is guaranteed under relatively weak assumptions (Section 3.3). Thirdly, in Section 4, we propose a *characterization* of defendable security properties when rationality of participants is based on Nash equilibrium. Finally, we consider the case of mixed strategies in Section 5, we generalize the results to

non-injective game models in Section 6, and apply our concepts to analyze fairness in the ASW contract-signing protocol in Section 7. Most of this paper (Sections 2–5) is a compressed version of the material already published in [9]. The novel contribution is presented in Sections 6 and 7.

We want to emphasize that our work does not focus on “classical” security protocols where most participants are assumed to be “honest”, i.e., to follow a typically deterministic sequence of actions. More appropriately, we should say that we study *interaction protocols* in general, where actions of participants may or may not be “honest”, and the actual set of available behaviors depends on the execution semantics of the protocol. We believe that the two kinds of assumptions (honesty vs. being in favor of the protocol objective) are largely orthogonal. A study of interplay between the two is left for future work.

1.1 Related Work

Researchers have considered protocol execution as a game with the very pessimistic assumption that the only goal of the other participants (“adversaries”) is to break the intended security property of the protocol. In this case, a protocol is correct if the “honest” participants have a strategy such that, for all strategies of the other agents, the goal of the protocol is satisfied (cf. e.g. [10]). Recently, protocols have been analyzed with respect to some game theoretic notions of rationality [7, 2] where preferences of participants are taken into account. An overview of connections between cryptography and game theory is given in [6]. Another survey [12] presents arguments suggesting that study of incentives in security applications is crucial. Buttyán, Hubaux and Čapkun [4] model protocols in a way similar to ours, and also use incentives to model the behavior of agents. However, they restrict their analysis to strongly Pareto-optimal Nash equilibria which is not necessarily a good solution concept for security protocols: First, it is unclear why agents would *individually* converge to a strongly Pareto-optimal play. Moreover, in many protocols it is unclear why agents would play a Nash equilibrium in the first place. Our method is more general, as we use the solution concept as a parameter to our analysis. Asharov et al. (2011) [2] use game theory to study gradual-release fair exchange protocols. They consider a protocol to be game-theoretically fair if the strategy that never aborts the protocol is a computational Nash-equilibrium. They prove that their analysis allows for solutions that are not admitted by the traditional cryptographic definition. Groce and Katz [8] show that if agents have a strict incentive to achieve fair exchange, then gradual-release fair exchange without trusted third party (TTP) is possible under the assumption that the other agents play rationally. Syverson [14] presents a *rational exchange* protocol for which he shows that “enlightened, self-interested parties” have no reason to cheat. Finally, Chatterjee & Raman [5] use assume-guarantee synthesis for synthesis of contract signing protocols.

In summary, rationality-based correctness of protocols has been studied in a number of papers, but usually with a particular notion of rationality in mind. In contrast, we define a concept of correctness where a game-theoretic solution concept is a parameter of the problem. Even more importantly, our concept of *defendability* of a security property is completely novel. The same applies to our characterizations of defendable properties under Nash equilibrium.

2 Protocols and Games

A protocol is a specification of how agents should interact. Protocols can contain *choice points* where several actions are available to the agents. An agent is *honest* if he follows the protocol specification, and *dishonest* otherwise, i.e., when he behaves in a way that is not allowed by the protocol. In the latter case, the agent is only restricted by the physical and logical actions that are available in the environment.

For instance, in a cryptographic protocol, dishonest agents can do anything that satisfies properties of the cryptographic primitives, assuming perfect cryptography (as in [11]). The protocol, together with a model of the environment of action, a subset of agents who are assumed to be honest, and the operational semantics of action execution, defines a multi-agent transition system that we call the *model* of the protocol. In the rest of the paper, we focus on protocol models, and abstract away from how they arise. We also do not treat the usual “network adversary” that can intercept, delay and forge messages, but essentially assume the existence of secure channels. The issue of the “network adversary” is of course highly relevant for security protocols, but orthogonal to the aspects we discuss in this paper. In the full version of this paper [9], we present contract signing protocols as a running example. In such a protocol, Alice and Bob want to sign a contract. Among the most relevant game-theoretic security properties of such protocols are fairness, balancedness, and abuse-freeness.

We use *normal-form games* as abstract models of interaction in a protocol.

Definition 2.1 (Frames and games). *A game frame is a tuple $\Gamma = (N, \Sigma)$, where $N = \{A_1, \dots, A_{|N|}\}$ is a finite set of agents, and $\Sigma = \Sigma_{A_1} \times \dots \times \Sigma_{A_{|N|}}$ is a set of strategy profiles.*

A normal-form (NF) game is a game frame plus a utility profile $u = \{u_1, \dots, u_{|N|}\}$ where $u_i : \Sigma \rightarrow \mathbb{R}$ is a utility function assigning utility values to strategy profiles.

Game theory uses *solution concepts* to define which strategy profiles capture rational interactions. Let \mathcal{G} be a class of games with the same strategy profiles Σ . Formally, a solution concept for \mathcal{G} is a function $SC : \mathcal{G} \rightarrow \mathcal{P}(\Sigma)$ that, given a game, returns a set of *rational* strategy profiles. Well-known solution concepts include e.g. Nash equilibrium (NE), dominant and undominated strategies, Stackelberg equilibrium, Pareto optimality etc.

Protocols as Games. Let P be a model of a protocol. We will investigate properties of P through the game frame $\Gamma(P)$ in which strategies are *conditional plans* in P , i.e., functions that specify for each choice point which action to take. A set of strategies, one for each agent, uniquely determines a *run* of the protocol, i.e., a sequence of actions that the agents will take. $\Gamma(P)$ takes runs to be the outcomes in the game, and hence maps strategy profiles to runs.

Security protocols are designed to achieve one or more *security requirements* and/or *functionality requirements*. We only consider requirements that can be expressed in terms of single runs having a certain property. We model this by a subset of possible behaviors, called the *objective of the protocol*.

Definition 2.2. *Given a game frame $\Gamma = (N, \Sigma)$, an objective is a set $\gamma \subseteq \Sigma$. We call γ nontrivial in Γ iff γ is neither impossible nor guaranteed in Γ , i.e., $\emptyset \neq \gamma \neq \Sigma$.*

3 Incentive-Based Security Analysis

In this section, we give a definition of correctness of security protocols that takes into account rational decisions of agents, based on their incentives.

3.1 Incentive-Based Correctness

As we have pointed out, the requirement that all strategy profiles satisfy the objective might be too strong. Instead, we will require that all *rational* runs satisfy the objective. In case there are no rational runs, all behaviors are equally rational; then, we require that all strategy profiles must satisfy γ .

Definition 3.1. A protocol model represented as game frame $\Gamma = (N, \Sigma)$ with utility profile u is correct with respect to objective γ under solution concept SC , written $(\Gamma, u) \models_{SC} \gamma$, iff:

$$\begin{cases} SC(\Gamma, u) \subseteq \gamma & \text{if } SC(\Gamma, u) \neq \emptyset \\ \gamma = \Sigma & \text{otherwise.} \end{cases}$$

3.2 Unknown Incentives

Definition 3.1 applies to a protocol when a utility profile is given. However, the exact utility profiles are often unknown. One way out is to require the protocol to be correct for *all possible* utility profiles.

Definition 3.2. A protocol model represented by game frame Γ is valid with respect to objective γ under solution concept SC (written $\Gamma \models_{SC} \gamma$) iff $(\Gamma, u) \models_{SC} \gamma$ for all utility profiles u .

It turns out that, under some reasonable assumptions, protocols are only valid for trivial objectives.

Definition 3.3. Let $G = (N, \Sigma, (u_1, \dots, u_n))$. Let $\pi = (\pi_1, \dots, \pi_n)$, where for all $i \in N$, $\pi_i : \Sigma_i \rightarrow \Sigma_i$ is a permutation on Σ_i . We slightly abuse the notation by writing $\pi((s_1, \dots, s_n))$ for $(\pi_1(s_1), \dots, \pi_n(s_n))$. A solution concept is closed under permutation iff $s \in SC((N, \Sigma, (u'_1, \dots, u'_n)))$ if and only if $\pi(s) \in SC((N, \Sigma, (u'_1 \circ \pi_1^{-1}, \dots, u'_n \circ \pi_n^{-1})))$.

Theorem 3.4. If SC is closed under permutation, then $\Gamma \models_{SC} \gamma$ iff $\gamma = \Sigma$.¹

Thus, correctness for all distributions of incentives is equivalent to correctness in all possible runs.

3.3 Defendability of Protocols

Typical analysis of a protocol implicitly assumes some participants to be aligned with its purpose. E.g., one usually assumes that communicating parties are interested in exchanging a secret without the eavesdropper getting hold of it, that a bank wants to prevent web banking fraud etc. In this section, we formalize this idea by assuming a subset of agents, called the *defenders* of the protocol, to be in favor of its objective. Our new definition of correctness says that a protocol is correct with respect to some objective γ if and only if it is correct with respect to every utility profile in which the preferences of all defenders comply with γ .²

Definition 3.5. A group of agents $D \subseteq N$ supports the objective γ in game (N, Σ, u) iff for all $i \in D$, if $s \in \gamma$ and $s' \in \Sigma \setminus \gamma$ then $u_i(s) > u_i(s')$.

A protocol model represented as game frame Γ is defended by agents D , written $\Gamma \models_{SC} [D]\gamma$, iff $(\Gamma, u) \models_{SC} \gamma$ for all utility profiles u such that D supports γ in game (Γ, u) .

Clearly, if there are no defenders, then defendability is equivalent to ordinary protocol validity:

Proposition 3.6. If Γ is a game frame and SC is a solution concept, we have that $\Gamma \models_{SC} [\emptyset]\gamma$ iff $\Gamma \models_{SC} \gamma$.

If all agents are defenders, any protocol is correct, as long as the solution concept does not select *strongly Pareto-dominated* strategy profiles, and there always is some strategy profile which is rational according to the solution concept.

Definition 3.7. A solution concept is weakly Pareto iff it never selects a strongly Pareto dominated outcome (i.e., such that there exists another outcome strictly preferred by all the players). It is efficient iff it never returns the empty set.

¹ For proofs of all theorems and definitions of auxiliary concepts, we refer to the original paper [9].

² There is an analogy of the concept to [1] where ‘‘robust’’ goals are studied, i.e., goals that are achieved as long as a selected subset of agents behaves correctly.

Theorem 3.8. *If Γ is a game frame and SC is an efficient weakly Pareto solution concept then $\Gamma \models_{SC} [N]\gamma$.*

Many solution concepts are both efficient and weakly Pareto, for example: Stackelberg equilibrium, maximum-perfect cooperative equilibrium, backward induction and subgame-perfect Nash equilibrium in perfect information games. On the other hand, Nash equilibrium is neither weakly Pareto nor efficient, and equilibrium in dominant strategies is weakly Pareto but not necessarily efficient.

Clearly, defendability of a protocol is monotonic with respect to the set of defenders. This justifies the following definition.

Definition 3.9. *The game-theoretic security level of protocol P is the antichain of minimal sets of defenders that make the protocol correct.*

4 Characterizing Defendability under Nash Equilibrium

In this section, we turn to properties that can be defended if agents' rationality is based on Nash equilibrium or Optimal Nash Equilibrium.

4.1 Defendability under Nash Equilibrium

From Theorem 3.4, we know that no protocol is valid under Nash equilibrium (NE) for any nontrivial objective, since NE is closed under permutation. Do things get better if we assume some agents to be in favor of the security objective? We now look at the extreme variant of the question, i.e., defendability by the grand coalition N . Note that, by monotonicity of defendability wrt the set of defenders D , nondefendability by N implies that the objective is not defendable by any coalition at all.

Our first result in this respect is negative: we show that in every game frame there are nontrivial objectives that are not defendable under NE.

Theorem 4.1. *Let Γ be a game frame with at least two players and at least two strategies per player. Moreover, let γ be a singleton objective, i.e., $\gamma = \{\omega\}$ for some $\omega \in \Sigma$. Then, $\Gamma \not\models_{NE} [N]\gamma$.*

In particular, the construction from the above proof shows that, as mentioned before, there are cases where the “defending” coalition has a strategy to achieve a goal γ , but there are still rational plays in which the goal is not achieved.

To present the general result that characterizes defendability of security objectives under Nash equilibrium, we need to introduce additional concepts. In what follows, we use $s[t_i/i]$ to denote $(s_1, \dots, s_{i-1}, t_i, s_{i+1}, \dots, s_N)$, i.e., the strategy profile that is obtained from s when player i changes her strategy to t_i .

Definition 4.2. *Let γ be a set of strategy profiles in Γ . The deviation closure of γ is defined as $Cl(\gamma) = \{s \in \Sigma \mid \exists i \in N, t_i \in \Sigma_i . s[t_i/i] \in \gamma\}$.*

$Cl(\gamma)$ extends γ with the strategy profiles that are reachable by unilateral deviations from γ . Thus, $Cl(\gamma)$ can be seen as the closure of γ with the behaviors that are relevant for Nash equilibrium. Moreover, the following notion captures strategy profiles that can be used to construct sequences of unilateral deviations ending up in a cycle.

Definition 4.3. *A strategic knot in γ is a subset of strategy profiles $S \subseteq \gamma$ such that there is a permutation (s^1, \dots, s^k) of S where: (a) for all $1 \leq j < k$, $s^{j+1} = s^j[s_i^{j+1}/i]$ for some $i \in N$, and (b) $s^j = s^k[s_i^j/i]$ for some $i \in N, j < k$.*

Essentially, this means that every strategy s^{j+1} is obtained from s^j by a unilateral deviation of a single agent. If these deviations are rational (i.e., increase the utility of the deviating agent), then the knot represents a possible endless loop of rational, unilateral deviations which precludes a group of agents from reaching a stable joint strategy. We now state the main result of this section.

Theorem 4.4. *Let Γ be a finite game frame and γ a nontrivial objective in Γ . Then, $\Gamma \models_{\text{NE}} [N]\gamma$ iff $Cl(\gamma) = \Sigma$ and there is a strategy profile in γ that belongs to no strategic knots in γ .*

4.2 Optimal Nash Equilibria

Nash equilibrium is a natural solution concept for a game played repeatedly until the behavior of all players converges to a stable point. For a one-shot game, NE possibly captures convergence of the process of deliberation. It can be argued that, among the available solutions, no player should contemplate those which are strictly worse for everybody when compared to another stable point. This gives rise to the following refinement of Nash equilibrium: $\text{OptNE}(\Gamma, u)$ is the set of *optimal Nash equilibria* in game (Γ, u) , defined as those equilibria *that are not strongly Pareto-dominated by another Nash equilibrium*. Defendability by the grand coalition under OptNE has the following simple characterization.

Theorem 4.5. *Let Γ be a finite game frame and γ a nontrivial objective in Γ . Then, $\Gamma \models_{\text{OptNE}} [N]\gamma$ iff there is a strategy profile in γ that belongs to no strategic knots in γ .*

5 Defendability in Mixed Strategies

So far, we considered only deterministic (pure) strategies. It is well known that for many games and solution concepts, rational strategies exist only when taking mixed strategies into account. We now extend our definition of correctness to mixed strategies, i.e., randomized conditional plans represented by probability distributions over pure strategies from Σ_{A_i} . Let $\text{dom}(s)$ be the support (domain) of a mixed strategy profile s , i.e., the set of pure strategy profiles that have nonzero probability in s . We extend the notion to sets of mixed strategy profiles in the obvious way. By SC^m we denote the variant of SC in mixed strategy profiles. A protocol is correct in mixed strategies iff all the possible behaviors resulting from a rational (mixed) strategy profile satisfy the goal γ ; formally: $\Gamma, u \models_{SC^m} \gamma$ iff $\text{dom}(SC^m(\Gamma, u)) \subseteq \gamma$ when $SC^m(\Gamma, u) \neq \emptyset$ and $\gamma = \Sigma_\Gamma$ otherwise. The definitions of protocol validity and defendability in mixed strategies ($\Gamma \models_{SC^m} \gamma$ and $\Gamma \models_{SC^m} [D]\gamma$) are analogous. For defendability in mixed strategies under Nash equilibrium, we have the following, rather pessimistic result.

Theorem 5.1. *Let Γ be a finite game frame, and γ an objective in it. Then, $\Gamma, u \models_{\text{NE}}^m [N]\gamma$ iff $\gamma = \Sigma$.*

On the other hand, it turns out that *optimal Nash equilibrium* yields a simple and appealing characteristics of N -defendable properties. In the following, γ is closed under convex combination of strategies iff every combination of strategies that appear in some profile in γ again is an element of γ .

Theorem 5.2. $\Gamma \models_{\text{OptNE}}^m [N]\gamma$ iff $\gamma = \text{Conv}(\gamma)$, i.e., γ is closed under convex combination of strategies.

Corollary 5.3. $\Gamma \models_{\text{OptNE}}^m [N]\gamma$ iff there exist subsets of individual strategies $\chi_1 \subseteq \Sigma_1, \dots, \chi_{|N|} \subseteq \Sigma_{|N|}$ such that $\gamma = \chi_1 \times \dots \times \chi_{|N|}$.

That is, security property γ is defendable by the grand coalition in Γ iff γ can be decomposed into constraints on individual behavior of particular agents.

6 Defendability in Non-Injective Games

Normal game frames are usually defined in the literature as $\Gamma = (N, \Sigma, \Omega, o)$, where N, Σ are as before, Ω is the set of (abstract) *outcomes* of the game, and $o : \Sigma \rightarrow \Omega$ maps strategy profiles to outcomes. Our analysis so far has been based on the standard assumption that o is a bijection. In other words, there




Figure 1: Naive contract-signing: from protocol to EF game to NF game

is a 1-1 relationship between joint behaviors of agents and the outcomes of those behaviors. Then, we can identify outcomes with strategy profiles, and omit the former from the game model. However, the standard construction of a game model from a protocol assumes the outcomes to be *runs* of the protocol. In that case, the assumption does *not* hold; in particular, the mapping is not injective.

Example 6.1. Consider the naive contract signing protocol in Figure 1. Alice sends her signature to Bob, who responds with his signature. Alice and Bob can stop the protocol at any moment (thereby deviating from the protocol). If we assume runs of the protocol to be the outcomes, this gives rise to an Extensive Form game frame, which can be then transformed to an NF game frame by the canonical construction. Clearly, the mapping between strategy profiles and outcomes is not injective.

In general NF games, utility functions assign utility values to *outcomes* rather than strategy profiles. That is, $u_i : \Omega \rightarrow \mathbb{R}$. Moreover, an objective is assumed to select a *subset of outcomes*. This follows from the methodological assumption that an outcome encapsulates every relevant aspect of the play that has occurred. We observe that the definitions in Section 3 can be lifted to the general case by changing the types of u_i and γ accordingly. However, the results in Sections 4–5 cannot be lifted that easily. Games with non-injective outcome functions require a more general treatment, which we present below.

Definition 6.2. Given a game frame Γ , we define the deviation graph of Γ ($Dev(\Gamma)$) to be the undirected graph where outcomes from Γ are vertices, and edges connect outcomes that are obtained from strategy profiles which differ only in 1 individual strategy (thus corresponding to a potential unilateral deviation).

Moreover, for an objective $\gamma \subseteq \Omega$, we will use $Dev_\gamma(\Gamma)$ to denote the subgraph of $Dev(\Gamma)$ consisting only of the vertices from γ and the edges between them.

It is easy to see that the construction of $Dev(\Gamma)$ and $Dev_\gamma(\Gamma)$ from Γ, γ is straightforward. Let V be a subset of nodes in a graph. We define the *neighborhood of V* , denoted $Neighb(V)$, as V together with all the nodes adjacent to V . We observe that $Neighb(V)$ “implements” the deviation closure of V in $Dev(\Gamma)$. Moreover, ω does not lie on a strategic knot iff its connected component does not include a cycle. This leads to the following, more general, characterizations of defendability (we omit the proofs due to lack of space). Again, we assume that γ is nontrivial, i.e., $\emptyset \neq \gamma \neq \Omega$.

Theorem 6.3. γ is defended by the grand coalition in Γ under Nash equilibrium iff:

1. The neighborhood of γ in $Dev(\Gamma)$ covers the whole graph ($Neighb(\gamma) = \Omega$), and
2. $Dev_\gamma(\Gamma)$ includes at least one connected component with no cycles.

Theorem 6.4. γ is defended by the grand coalition in Γ under optimal Nash equilibrium iff $Dev_\gamma(\Gamma)$ includes at least one connected component with no cycles.

Theorem 6.5. γ is defended in mixed strategies by the grand coalition in Γ under optimal Nash equilibrium iff γ is obtained by a convex combination of strategies.

7 Example: The ASW contract-signing protocol

A contract-signing protocol is used by two participants, usually called Alice and Bob, to sign a contract over an asymmetric medium as the internet. The central security properties are *fairness* (Alice should get a signed copy of the contract if and only if Bob gets one), *balancedness* (there is no point in the protocol run where Bob alone can decide whether the contract will be signed or not, i.e., Alice cannot abort the signing anymore but Bob still can abort) and *abuse-freeness* (if balance cannot be achieved, then at least Bob should not be able to prove the fact that he has the above-mentioned strong position in the current state of the protocol to an outsider). The contract-signing protocol P_{ASW} , introduced in [3], uses *commitments*, which are legally binding “declarations of intent” by Alice and Bob to sign the contract. The protocol operates as follows: (1) Alice sends a commitment cm_A to Bob; (2) Bob sends his commitment cm_B to Alice; (3) Alice sends the contract sc_A , digitally signed with her signature, to Bob; (4) Bob sends the contract sc_B , signed with his signature, to Alice.

If one of these messages is not sent by the corresponding signer, the other party may contact the TTP:

- If Alice does not receive a commitment from Bob, she can contact the TTP with an *abort request*, which instructs the TTP to mark this session of the protocol as aborted;
- If Bob does not receive Alice’s signature, but has her commitment, he can send a *resolve request* to the TTP, who then issues a *replacement contract* (a document that is legally equivalent to the contract signed by Alice), unless Alice has sent an abort request earlier,
- If Alice does not receive Bob’s signature, but has his commitment, she can send a *resolve request* to the TTP as well, which allows her to receive a replacement contract.

It can be shown that the protocol is fair if the TTP is reliable (it will never stop the protocol on its own). It is also balanced if neither Alice nor Bob can drop or delay messages from the other signer to the TTP. Let us denote outcomes by sets of agents who have obtained the signature of the other player. Thus, \emptyset represents the situation where nobody got a signed contract, $\{\text{sign}_A\}$ the situation where Alice obtained Bob’s signature but note vice versa, etc. Applying the definitions in Section 3.3, one can show the following. If SC is either *Nash equilibrium* or *undominated strategies*, we have:

1. $P_{ASW} \models_{SC} [\{\text{Bob}\}]\{\emptyset, \{\text{sign}_B\}, \{\text{sign}_A, \text{sign}_B\}\}$,
2. $P_{ASW} \models_{SC} [\{\text{Alice}\}]\{\emptyset, \{\text{sign}_A\}, \{\text{sign}_A, \text{sign}_B\}\}$.

We now consider the case where TTP is not necessarily reliable. If the TTP can stop the protocol at any time, then the protocol does not guarantee fairness anymore. On the other hand, if Bob wants the protocol to be fair, then he can ensure fairness by simply sending a signed contract to Alice as soon as he receives her signature. Clearly, Alice alone (without an honest TTP to assist her) cannot achieve fairness. Hence the game-theoretic security level of the ASW protocol without reliable TTP is the set $\{\{\text{Bob}\}, \{\text{TTP}\}\}$. This holds for both Nash equilibrium and undominated strategies.

8 Conclusions

We propose a framework for analyzing security protocols (and other interaction protocols), that takes into account the incentives of agents. In particular, we consider a novel notion of *defendability* that guarantees that all the runs of the protocol are correct as long as a given subset of the participants (the “defenders”) is in favor of the security property. We have obtained some characterization results for defendability under Nash equilibria and optimal Nash equilibria. In the original paper [9], we also address the computational complexity of the corresponding decision problems, both in the generic case and in some special cases.

In the future, we plan to combine our framework with results for protocol verification using game logics (such as ATL), especially for those solution concepts that can be expressed in that kind of logics.

Acknowledgements. We thank the SR2014 reviewers for their extremely useful remarks. Addressing the fundamental ones was not possible in this extended abstract due to space and time constraints, but we will use them in the journal version of the paper (in preparation).

Wojciech Jamroga acknowledges support of the National Research Fund Luxembourg (FNR) under project GaLOT – INTER/DFG/12/06.

References

- [1] T. Ågotnes, W. van der Hoek & M. Wooldridge (2010): *Robust normative systems and a logic of norm compliance*. *Logic Journal of the IGPL* 18(1), pp. 4–30, doi:10.1093/jigpal/jzp070.
- [2] G. Asharov, R. Canetti & C. Hazay (2011): *Towards a Game Theoretic View of Secure Computation*. In: *Proceedings of EUROCRYPT*, pp. 426–445, doi:10.1007/978-3-642-20465-4_24.
- [3] N. Asokan, V. Shoup & M. Waidner (1998): *Asynchronous protocols for optimistic fair exchange*. In: *Proceedings of the IEEE Symposium on Research in Security and Privacy*, IEEE Computer Society Press, pp. 86–99, doi:10.1109/SECPRI.1998.674826.
- [4] L. Buttyán, J. Hubaux & S. Čapkun (2004): *A formal model of rational exchange and its application to the analysis of Syverson’s protocol*. *Journal of Computer Security* 12(3,4), pp. 551–587.
- [5] K. Chatterjee & V. Raman (2010): *Assume-Guarantee Synthesis for Digital Contract Signing*. CoRR abs/1004.2697.
- [6] Y. Dodis & T. Rabin (2007): *Cryptography and Game Theory*. In: *Algorithmic Game Theory*, chapter 8, pp. 181–208, doi:10.1017/CB09780511800481.010.
- [7] G. Fuchsbauer, J. Katz & D. Naccache (2010): *Efficient Rational Secret Sharing in Standard Communication Networks*. In: *Proceedings of TCC*, pp. 419–436, doi:10.1007/978-3-642-11799-2_25.
- [8] A. Groce & J. Katz (2012): *Fair Computation with Rational Players*. In: *Proceedings of EUROCRYPT*, pp. 81–98, doi:10.1007/978-3-642-29011-4_7.
- [9] W. Jamroga, M. Melissen & H. Schnoor (2013): *Defendable Security in Interaction Protocols*. In: *Proceedings of the 16th International Conference on Principles and Practice of Multi-Agent Systems PRIMA 2013*, LNCS 8291, Springer, pp. 132–148, doi:10.1007/978-3-642-44927-7_10.
- [10] S. Kremer & J. Raskin (2002): *Game Analysis of Abuse-Free Contract Signing*. In: *Proceedings of the 15th IEEE Computer Security Foundations Workshop (CSFW’02)*, IEEE Computer Society Press, pp. 206–220, doi:10.1109/CSFW.2002.1021817.
- [11] S. Kremer & J.-F. Raskin (2003): *A game-based verification of non-repudiation and fair exchange protocols*. *Journal of Computer Security* 11(3), doi:10.1007/3-540-44685-0_37.
- [12] T. Moore & R. Anderson (2011): *Economics and Internet Security: a Survey of Recent Analytical, Empirical and Behavioral Research*. Technical Report TR-03-11, Computer Science Group, Harvard University.
- [13] M. Osborne & A. Rubinstein (1994): *A Course in Game Theory*. MIT Press.
- [14] P. Syverson (1998): *Weakly Secret Bit Commitment: Applications to Lotteries and Fair Exchange*. In: *Proceedings of CSFW*, pp. 2–13, doi:10.1109/CSFW.1998.683149.