

Computational Semantics for Basic Protocol Logic – A Stochastic Approach

Gergei Bana*, Koji Hasebe, and Mitsuhiro Okada

¹ Dept of Computer Science, University of California at Davis, Davis, CA, USA

² Research Center for Verification and Semantics,

National Institute of Advanced Industrial Science and Technology Osaka, Japan

³ Department of Philosophy, Keio University, Tokyo, Japan

gebana@cs.ucdavis.edu k-hasebe@aist.go.jp

mitsu@abelard.flet.keio.ac.jp

Abstract. This paper is concerned about relating formal and computational models of cryptography in case of active adversaries when formal security analysis is done with first order logic. We first argue that the way Datta et al. defined computational semantics to their Protocol Composition Logic, gives rise to problems because of focusing on occurrences of bit-strings on individual traces instead of occurrences of probability distributions of bit-strings across the distribution of traces. We therefore introduce a new, fully probabilistic method to assign computational semantics to the syntax. We present this via considering a simple example of such a formal model, the Basic Protocol Logic by K. Hasebe and M. Okada [19], but the technique is suitable for extensions to more complex situations such as PCL. The idea is to make use of the usual mathematical treatment of stochastic processes, hence be able to treat arbitrary probability distributions, non-negligible probability of collision, causal dependence or independence, and so on. Along the way, we also point out some instances of the original syntax that had to be modified, as – although sound for formal semantics – they were not sound for computational semantics.

Keywords. cryptographic protocols, formal methods, first order logic, computational semantics

1 Introduction

In the past few years, linking the formal and computational models of cryptography has become of central interest. Several different methods have emerged for both active and passive adversaries. In this paper we would like to consider the relationship of the two models when formal security analysis is done with first order logic. In the formal approach of our interest, protocol correctness is analyzed by defining a syntax with axioms and inference rules and then proving some property. A logical proof then ensures that the property will be true in any formal model (semantics) of the syntax. The link to the computational world then is done by assigning a computational semantics (instead of formal) to the syntax, proving that the axioms and inference rules hold there, and hence a property correct in the syntax must be true in the computational model. However, as it turns out, it is not unambiguous how to define the computational semantics, and when a property should be deemed “true” computationally.

Recently, Datta et al. in [13] gave a computational semantics to the syntax of their Protocol Composition Logic of [15, 12] (cf. also [1] for a *protocol composition logic project overview*). In their treatment, every action by the honest participants is recorded on each execution trace (which

* Partially supported by a Packard Fellowship.

have identical probabilities), and bit strings emerging later are checked whether they were recorded earlier and to what action they corresponded (the only actions of the adversary that are recorded are send and receive). This way, they first define whether a property is true on a particular trace, and they say the property is true in the model if it is true on an overwhelming number of traces. This method however relies on negligible collision probabilities, because otherwise there would be a large probability of identifying bit strings with the wrong actions. Moreover, as the comparisons are done on each trace separately it is not possible to track correlations.

Our approach puts more emphasis on probabilities. Instead of defining what is true on each trace, we say that a property is true in the model if a “cross-section” of traces provides the right probabilities for computational realizations of the property in question. An underlying stochastic structure ensures that we can detect if something depends on the past or does not. It is not coincidences on traces that we look for, but correlations of probability distributions.

We introduce our method on a rather simple syntax, namely, a somewhat modified version of Basic Protocol Logic (or BPL, for short) by K. Hasebe and M. Okada [19] and leave extensions to more complex situations such as the Protocol Composition Logic to future work. The reason for this is partly the limited space, partly to avoid distraction by an elaborate formal model from the main ideas, but also that a complete axiomatization of the syntax used by Datta et al. for their computational PCL has not yet been published anywhere, only fragments are available. We would like to emphasize though that our point is not to give a computational semantics to BPL but to provide a technique that works well in much more general situations as well.

BPL is a logical inference system to prove correctness of a protocol. Originally, it included signatures as well, but for simplicity, we leave that out from this analysis. BPL was defined to give a simple formulation of a core part of the protocol logics of [15, 12, 11] for proving some aimed properties within the framework of first order logic.

We first give the axiomatic system in first-order predicate logic for proving the agreement properties. A message is represented by a first-order term that uses encryption and pairing symbols, an atomic formula is a sequence of primitive actions (as send, receive and generate) of principals on terms. We set some properties about nonces and cryptographic assumptions as non-logical axioms, and give a specific form of formulas, called *query form*, which has enough expressive power to specify our intended authentication properties.

Although BPL is sound with respect to purely formal semantics, in order to ensure soundness for computational semantics, some modifications of the original syntax of BPL were necessary:

1. Instead of denoting encryptions as $\{m\}_A$, which was used in the original version of the purely symbolic model-based BPL inference system, we indicate the random seed of the encryption as $\{m\}_A^r$ (as Datta et al. do) As it turned out, a consistent computational interpretation is much harder, if not impossible without the random seed in the syntax.

2. The original subterm and equiterm axioms were not all computationally sound so we just take a certain subset of those that are computationally sound. We are not taking all the sound term axioms, as it is not known how to give a complete characterization of them. ⁴

⁴ The uses of subterm and equiterm relations, such as $s \sqsubseteq t$ and $s = t$, are essential for correctness proofs of protocols in general, including BPL and PCL. Any symbolic term model, hence, should reflect the symbolic term structures, and such a term model maybe called a “standard” model with respect to subterm and equiterm relations. BPL’s symbolic semantics takes such standard term model which also satisfy certain properties for nonce-verifications, which are listed as non-logical axioms in our BPL syntax. Our result 2 above shows that only the truth of a certain useful subset of the subterm theory axioms is preserved under computational interpretation. As we will show, the nonce-verification axioms turn out to be sound.

The original BPL also proved completeness for formal semantics with the original set of axioms, however, we do not consider completeness in this paper. It is an open question whether anything about completeness can be said in the computational case.

We then define the computational semantics. This involves giving a stochastic structure that results when the protocol is executed. Principals output bit strings (as opposed to terms) with certain probability distributions. The bit strings are then recorded in a trace as being generated, sent or received by some principal. This provides a probability distribution of traces. We show how to answer whether a bit string corresponding to a term was sent around with high probability or not. For example a formal term $\{M\}_A^r$ was sent around in the computational model if a cross-section of all traces provides the correct probability distribution that corresponds to sending $\{M\}_A^r$. Or, a nonce N was generated, if another cross-section provides the right probabilities, and that distribution must be independent of everything that happened earlier. This way we define when a certain formula in the syntax is true in the computational semantics. We then analyze whether the axioms of the syntax are true in the semantics, and if they are, then we conclude that a formula that can be proved in the syntax is also true in the semantics.

The merits of our approach compared with the computational semantics of Datta et al.

We would like to point out the main aspects where problems arise in case of the treatment of Datta et al. and where we think our method works better than theirs:

1. They rely on counting equiprobable traces. Unequal probabilities may be dealt with by counting a trace more than once (although a priori it is not quite clear whether this will lead to problems), but their method certainly only applies to executions when the number of possible computational traces for a given security parameter is finite. Since probabilistic polynomial time processes are not limited to finitely many traces (only the *expected termination time* is polynomial), infinite number of traces should not be excluded. Our method works for infinite number of traces and arbitrary probability distributions.

2. As Datta et al. derive the validity of a formula in the model from validity of the formula on individual traces, they have to make sure that there are not too many accidental coincidences. As a result, their method only works when collision probability is negligible, but, more importantly, this results in a weaker set of syntactic axioms than what would otherwise be possible in our method. For example, they postulate that $\neg\text{Send}(\hat{X}, t)[b]_X \neg\text{Send}(\hat{X}, t)$ is an axiom whenever for all σ evaluation of variables by bit-strings, $\sigma(b) \neq \sigma(\text{Send}(\hat{X}, t))$. Here, \hat{X} is a principle, t is a term, $[b]_X$ is an action b carried out by principal \hat{X} in thread X assuming also that nothing else is carried out. In other words, it is an axiom that if \hat{X} did not send t before action b , then it did not send it even after action b as long as no σ evaluates b and $\text{Send}(\hat{X}, t)$ the same way. However, if there is even one coincidence in their evaluations, that prevents the axiom. We think this is an unnecessary restriction. As long as the probability distributions are different (up to negligibility) for any computational interpretation of b and $\text{Send}(\hat{X}, t)$, we can include $\neg\text{Send}(\hat{X}, t)[b]_X \neg\text{Send}(\hat{X}, t)$ in our axioms (in this paper we do not consider modal formulas, but this is not really a limitation as the actions of b can be included in the premise of the formula).

3. A further problem, that even makes the soundness proofs of Datta et al. questionable is the following: They define a formula (e.g. $\text{Send}(\hat{X}, t)$) to be true in the model if it holds on all traces except for some with negligible probability. They ignore the fact that the position of $\text{Send}(\hat{X}, t)$ on the traces may vary badly from trace to trace, for example, may depend on the future of the trace. A simple example of such a situation is when on two traces, which coincide up to step t_0 , say, $\text{Send}(\hat{X}, t)$ is chosen on one trace for $t_1 < t_0$, but on the other trace it is chosen somewhere else. Since the two traces coincide at step t_1 , if this time is picked on one trace, it must be picked on the other trace too. Maybe it is possible to prove that if there is a bad choice of the positions then there is a good choice as well, but we see no indication of such concerns in the papers of Datta

et al. As we suggest to use the standard tool of *filtration* in stochastic processes, this problem is taken care of in our semantics.

4. Finally, ignoring probability distributions and correlations give rise to pathologies like this one, putting further doubts at the correctness of their soundness proofs: Suppose that the encryption scheme is such that for any n_1, n_2 bit-strings generated randomly as nonces, any public key bit-string k_2 and any random seed r_2 for the encryption, there is another public key bit-string k_1 and a random seed r_1 such that $\{n_1\}_{k_1}^{r_1} = \{n_2\}_{k_2}^{r_2}$. This does not contradict CCA-2 security. Suppose principal A generates randomly nonce n_1 , and then principal B receives $\{n_2\}_{k_2}^{r_2}$ from the adversary. In such a case, it will be true according to the semantics of Datta et al., that $\exists N \exists R \exists K. \text{New}(A, N) \wedge \text{Receive}(B, \{N\}_K^R)$. This is however pathologic, and is a consequence of ignoring the fact that k_1 , if created by the adversary, cannot correlate with n_1 , which was not yet sent around. Furthermore, this seems to contradict their axiom saying that $\text{FirstSend}(X, t, t') \wedge a(Y, t'') \rightarrow \text{Send}(X, t') < a(Y, t'')$ where $X \neq Y$ and t subterm of t'' (meaning in our case that the first send action of A sending N had to occur before B could do anything with N). This problem persists even if such a coincidence cannot be efficiently computed.

Related Work. Formal methods emerged from the seminal work of Dolev and Yao [14], whereas computational cryptography grew out of the work of Goldwasser and Micali [16]. The first to link the two methods were Abadi and Rogaway in [3] "soundness" for passive adversaries in case of so-called type-0 security. A number of other papers for passive adversaries followed, proving "completeness" [21, 5], generalizing for weaker, more realistic encryptions schemes [5], considering purely probabilistic encryptions [18, 5], including limited models for active adversaries [20], addressing the issue of forbidding key-cycles [4], considering algebraic operations and static equivalence [7, 2]. Other approaches including active adversaries are considered by Backes et al. and Canetti in their *reactive simulatability* [6] and *universal composability* [10] frameworks, respectively. Non trace properties were investigated elsewhere too, however, not in the context of first order logic.

Organization of this paper. In Section 2, we outline the syntax of Basic Protocol Logic. In Section 3, we give a computational semantics to Basic Protocol Logic, and discuss soundness. Finally, in Section 4, we conclude and present directions for future work.

Special Thanks. We would like to thank Stéphane Glondu, Jesus Almansa, Arnab Roy and John Mitchell for the valuable discussions on the topic as well as Matthew Franklin.

2 Basic Protocol Logic

In this section, we present the syntax of Basic Protocol Logic modified to be suitable for computational interpretation. For the original BPL, please consult [19].

2.1 Language

Sorts and terms. Our language is order-sorted, with sorts `coin`, `name`, `nonce` and `message` such that terms of sorts `name` and `nonce` are terms of sort `message`. Let $\mathcal{C}_{\text{name}}$ be a finite set of constants of sort `name` (which represent principal names), and $\mathcal{C}_{\text{nonce}}$ a finite set of constants of sort `nonce`. For each $A \in \mathcal{C}_{\text{name}}$ let `coinA` be a sort such that any term of sort `coinA` is of sort `coin`, and let $\mathcal{C}_{\text{coin}_A}$ be a finite set of constants of sort `coinA`. Let $\mathcal{C}_{\text{coin}} := \bigcup_{A \in \mathcal{C}_{\text{name}}} \mathcal{C}_{\text{coin}_A}$. We require countably infinite free variables and countably infinite bound variables for each sort. We will use $A, B, \dots, A_1, A_2, \dots$ ($Q, Q', \dots, Q_1, Q_2, \dots$, resp.) to denote constants (variables, resp.) of sort `name`, $N, N', \dots, N_1, N_2, \dots$ ($n, n', \dots, n_1, n_2, \dots$, resp.) denote constants (variables, resp.) of sort `nonce`, $r^A, \dots, r_1^A, r_2^A, \dots$ ($s^A, \dots, s_1^A, s_2^A, \dots$, or $s, s', \dots, s_1, s_2, \dots$, resp.) denote

constants of sort coin_A (variables of sort coin_A , or variables of sort coin , resp.). The symbols $m, m', \dots, m_1, m_2, \dots$ are used to denote variables of sort message and $M, M', \dots, M_1, M_2, \dots$ to denote constants of sort message (that is, either name or nonce). Let $P, P', \dots, P_1, P_2, \dots$ denote any term of sort name , and let $\rho, \rho', \dots, \rho_1, \rho_2, \dots$ denote anything of sort coin . Compound terms of sort message are built from constants and free variables, and are defined by the grammar:

$$t ::= M \mid m \mid \langle t, t \rangle \mid \{t\}_P^\rho.$$

Where again, $M \in \mathcal{C}_{\text{name}} \cup \mathcal{C}_{\text{nonce}}$, m is any free variable of sort message , P is any constant or free variable of sort name , and ρ is any constant or free variable of sort coin . Hence, for example, $\langle \langle A_1, \{ \langle n, A_2 \rangle_Q^{r^A} \}, m \rangle$ is a term. We will use the shorter $\{n, A_2\}_Q^{r^A}$ instead of $\{ \langle n, A_2 \rangle_Q^{r^A} \}$. We will use the meta-symbols $t, t', \dots, t_1, t_2, \dots$ to denote terms, and ν, ν', \dots to denote any term of sort nonce .

Formulas. We introduce five binary predicate symbols: P generates ν , P receives t , P sends t , $t = t'$ and $t \sqsubseteq t'$, which represent “ P generates a fresh value ν as a nonce”, “ P receives a message of the form t ”, “ P sends a message of the form t ”, and equality and subterm relation for “ t is identical with t' ” and “ t is a subterm of t' ”, respectively. The first three are called *action predicates*, and the meta expression *acts* is used to denote one of the action predicates: *generates*, *receives* and *sends*. Atomic formulas are either of the form $P_1 \text{ acts}_1 t_1; P_2 \text{ acts}_2 t_2; \dots; P_k \text{ acts}_k t_k$, or $t = t'$, or $t \sqsubseteq t'$. The first one is called *trace formula*. A trace formula is used to represent a sequence of the principals’ actions such as “ P sends a message m , and after that, Q receives a message m' ”. We also use $\alpha_1; \dots; \alpha_k$ (or α , for short) to denote $P_1 \text{ acts}_1 t_1; \dots; P_k \text{ acts}_k t_k$ (where k indicates the *length* of α). When every P_i is identical with P for $1 \leq i \leq k$, then α^P denotes such a trace formula. For $\alpha (\equiv \alpha_1; \dots; \alpha_m)$ and $\beta (\equiv \beta_1; \dots; \beta_n)$, we say β includes α (denoted by $\alpha \subseteq \beta$), if there is a one-to-one, increasing function $j : \{1, \dots, m\} \rightarrow \{1, \dots, n\}$ such that $\alpha_i \equiv \beta_{j(i)}$. Formulas are defined by

$$\varphi ::= \alpha \mid t_1 = t_2 \mid t_1 \sqsubseteq t_2 \mid \neg \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \forall m \varphi' \mid \exists m \varphi'$$

where m is some bound variable, and φ' is obtained from φ by substituting m for every occurrence in φ of a free variable m' of the same sort as m . We use the meta expression $\varphi[\mathbf{m}]$ to indicate the list of all variables \mathbf{m} occurring in φ . Substitutions are represented in terms of this notation.

Finally, we introduce the notion of (*strict*) *order-preserving merge of trace formulas* α and β : An order-preserving merge of $\alpha (\equiv \alpha_1; \dots; \alpha_l)$ and $\beta (\equiv \beta_1; \dots; \beta_m)$ is a trace formula $\delta (\equiv \delta_1; \dots; \delta_n)$ if there are one-to-one increasing functions $j^\alpha : \{1, \dots, l\} \rightarrow \{1, \dots, n\}$, $j^\beta : \{1, \dots, m\} \rightarrow \{1, \dots, n\}$ such that $\alpha_i \equiv \delta_{j^\alpha(i)}$, $\beta_i \equiv \delta_{j^\beta(i)}$, and the union of the ranges of j^α and j^β cover $\{1, \dots, n\}$. δ is called a *strict order-preserving merge* if, furthermore, the ranges of j^α and j^β are disjoint.

Roles. A *protocol* is a set of *roles*, and each role for a principal (say, Q) is described as a trace formula of the form $\alpha^Q \equiv Q \text{ acts}_1 t_1; \dots; Q \text{ acts}_k t_k$.

2.2 The Axioms of Basic Protocol Logic

We extend the usual first-order predicate logic with equality by adding the following axioms (I), (II) and (III).⁵ The axioms we present here were sufficient for the protocols we have checked, but

⁵ The set of axioms here is not the same as it was in the original formulation of BPL, as those axioms were not all computationally sound. For example, the original axioms included that $\{m_1\}_B^s$ and $\{m_2\}_Q^{r^A}$ are equal only if $s = r^A$, $Q = B$, and $m_1 = m_2$. However, if the principal Q did not generate its public key properly, and if the randomization of s is not honest, then the interpretations of these terms may turn out to be equal (the whole distribution).

other protocols may need additional axioms as much more can be defined that are also computationally sound.

(I) Term axioms. Consider any set \bar{C} of countably infinitely many elements of each of sort `name`, sort `nonce` and sort `coin` such that it includes all elements of C_{name} , C_{nonce} and C_{coin} . Let \bar{A} be the free algebra constructed from \bar{C} via $\langle \cdot, \cdot \rangle$ and $\{ \cdot \}$: (with the appropriate sorts in the indexes of the encryption terms). The elements of \bar{A} are of sort `message`. We postulate the following axioms for $=$ and \sqsubseteq . Let m be all variables occurring in the corresponding terms. We require these for all $A, B \in C_{\text{name}}$:

- (a) If $t = t'$ is true in \bar{A} , then $\forall m t = t'$ is axiom. If $t \sqsubseteq t'$ is true in \bar{A} , then $\forall m t \sqsubseteq t'$ is axiom.
- (b) $\forall m (t = t), \forall m (t_1 = t_2 \rightarrow t_2 = t_1), \forall m (t_1 = t_2 \wedge t_2 = t_3 \rightarrow t_1 = t_3), \forall m (t_1 = t_2 \rightarrow t_1 \sqsubseteq t_2), \forall m (t_1 \sqsubseteq t_2 \wedge t_2 \sqsubseteq t_3 \rightarrow t_1 \sqsubseteq t_3)$
- (c) $\forall m Q s s' (\{t_1\}_Q^s = \{t_2\}_{Q'}^{s'} \rightarrow t_1 = t_2)$
- (d) If t_1 and t_2 are built by constants, and $C, D \in C_{\text{name}}$ then $\{t_1\}_A^C = \{t_2\}_B^D \rightarrow A = B \wedge C = D$.
- (e) $\forall m (\langle t_1, t_2 \rangle = \langle t_3, t_4 \rangle \rightarrow t_1 = t_3 \wedge t_2 = t_4)$
- (f) $\forall m Q s (\{t\}_Q^s \neq \langle t_1, t_2 \rangle), \forall m Q s n (\{t\}_Q^s \neq n), \forall m Q Q' s (\{t\}_Q^s \neq Q')$
- (g) $\forall m n (\langle t_1, t_2 \rangle \neq n), \forall m Q (\langle t_1, t_2 \rangle \neq Q)$
- (h) $\forall m (t \sqsubseteq \langle t_1, t_2 \rangle \rightarrow t \sqsubseteq t_1 \vee t \sqsubseteq t_2 \vee t = \langle t_1, t_2 \rangle), \forall m Q s (t_1 \sqsubseteq \{t_2\}_Q^s \rightarrow t_1 = \{t_2\}_Q^s \vee \exists m Q' s' (\{t_2\}_Q^s = \{m\}_{Q'}^{s'} \wedge t_1 \sqsubseteq m))$
- (i) $\forall m n (m \sqsubseteq n \rightarrow m = n), \forall m Q (m \sqsubseteq Q \rightarrow m = Q)$

(II) Rules for trace formulas. We postulate that $\beta \rightarrow \alpha$ for $\alpha \sqsubseteq \beta$ and $\gamma_1 \vee \dots \vee \gamma_n \leftrightarrow \alpha \wedge \beta$, where γ_i 's are the list of order-preserving merges of α and β . These axioms express the intuition that if a trace “happens”, then a subtrace of it also happens, and two traces happen if and only if one of their possible merges happen.

(III) Axioms for relationship between properties. We introduce the following set of formulas as non-logical axioms. These axioms represent some properties about nonces and cryptographic assumptions.

(1) Ordering:

$$\forall Q_1 Q_2 n m (n \sqsubseteq m \rightarrow \neg(Q_2 \text{ sends/receives/generates } m; Q_1 \text{ generates } n)).$$

Let $\overline{m_1 \sqsubseteq m_2 \sqsubseteq m_3}$ mean that $m_1 \sqsubseteq m_2 \sqsubseteq m_3$ and the only way m_1 occurs in m_3 is within m_2 . That is:

$$\overline{m_1 \sqsubseteq m_2 \sqsubseteq m_3} := \forall m (m_1 \sqsubseteq m \sqsubseteq m_3 \rightarrow m_2 \sqsubseteq m \vee (m \sqsubseteq m_2 \wedge \forall m_4 (m_2 \sqsubseteq m_4 \rightarrow \langle m, m_4 \rangle \not\sqsubseteq m_3 \wedge \langle m_4, m \rangle \not\sqsubseteq m_3)).$$

(2) Nonce verification 1: For each A, B constants of sort `name` and r^A constant of sort `coin`, we postulate

$$\begin{aligned} & \forall Q n_1 m_2 m_5 m_6 (A \text{ generates } n_1; A \text{ sends } m_2; Q \text{ receives } m_5 \\ & \wedge \overline{|n_1 \sqsubseteq \{m_6\}_B^{r^A} \sqsubseteq m_2|} \wedge n_1 \sqsubseteq m_5 \wedge \neg \overline{|n_1 \sqsubseteq \{m_6\}_B^{r^A} \sqsubseteq m_5|} \\ & \wedge \forall m_7 (A \text{ sends } m_7 \wedge n_1 \sqsubseteq m_7 \rightarrow \overline{|n_1 \sqsubseteq \{m_6\}_B^{r^A} \sqsubseteq m_7|}) \\ & \rightarrow \exists m_3 m_4 (A \text{ sends } m_2; B \text{ receives } m_3; B \text{ sends } m_4; Q \text{ receives } m_5 \wedge \{m_6\}_B^{r^A} \sqsubseteq m_3 \wedge n_1 \sqsubseteq m_4)) \end{aligned}$$

(3) Nonce verification 2: For each A, B, C of sort `name` (where A and C may coincide), r^A constant of sort `coin` _{A} and r^B constant of sort `coin` _{B} , we postulate

$$\begin{aligned} & \forall n_1 m_2 m_5 m_6 m_8 m_{10} (A \text{ generates } n_1; A \text{ sends } m_2; C \text{ receives } m_5 \\ & \wedge \overline{|n_1 \sqsubseteq \{m_6\}_B^{r^A} \sqsubseteq m_2|} \wedge n_1 \sqsubseteq m_5 \wedge \neg \overline{|n_1 \sqsubseteq \{m_6\}_B^{r^A} \sqsubseteq m_5|} \\ & \wedge \forall m_7 (A \text{ sends } m_7 \wedge n_1 \sqsubseteq m_7 \rightarrow \overline{|n_1 \sqsubseteq \{m_6\}_B^{r^A} \sqsubseteq m_7|}) \wedge B \text{ sends } m_4 \wedge \overline{|n_1 \sqsubseteq \{m_8\}_C^{r^B} \sqsubseteq m_4|} \\ & \wedge \forall m_9 (B \text{ sends } m_9 \wedge n_1 \sqsubseteq m_9 \rightarrow \overline{|n_1 \sqsubseteq \{m_8\}_C^{r^B} \sqsubseteq m_9|}) \wedge (\neg(C \text{ sends } m_{10} \wedge n_1 \sqsubseteq m_{10}) \vee A = C) \\ & \rightarrow \overline{|n_1 \sqsubseteq \{m_8\}_C^{r^B} \sqsubseteq m_5|} \end{aligned}$$

There are other possible axiomatizations, but the authors of [19] found this particularly useful (more exactly a somewhat less general version). The meaning of the Ordering axiom is clear. Nonce verification 1 and 2 are based on the idea of the authentication-tests [17]. Nonce verification I means that if A sent out a nonce n_1 encrypted with the public key of B that was not sent in any other way, and Q received this nonce in some other form, then the encrypted nonce had to go through B . The reason that we require A and B to be names and not arbitrary variables is that we do not want to require any principals in an arbitrary run to encrypt securely.

2.3 Query form and correctness properties

We introduce a general form of formulas, called *query form*, to represent our aimed correctness properties. In order to make the discussion simpler, we consider only the case of two party authentication protocols, however our query form can be easily extended so as to represent the correctness properties with respect to other types of protocols which include more than two principals.

Definition 1. (Query form) *Query form is a formula of the following form.*

$$Honest(\alpha^A) \wedge \beta^B \wedge Only(\beta^B) \rightarrow \gamma$$

We present the precise definition of $Only(\alpha^B)$ and of $Honest(\alpha^A)$ in the Appendix. $Only(\alpha^B)$ means that B performs only the actions of α^B , and nothing else, whereas $Honest(\alpha^A)$ represents “ A performs only a run of an initial segment of α^A which ends with a sending action or the last action of α^A ”. For example, from responder’s (namely, B ’s) view, the non-injective agreement of the protocol $\Pi = \{\alpha^A[B/Q_2, \mathbf{m}, \mathbf{s}], \beta^B[A/Q_1, \mathbf{m}, \mathbf{s}]\}$ can be described as the following formula: $Honest(\alpha^A[Q_2, \mathbf{m}, \mathbf{s}]) \wedge \beta^B[A/Q_1, \mathbf{N}/\mathbf{m}, \mathbf{r}/\mathbf{s}] \wedge Only(\beta^B[A/Q_1, \mathbf{N}/\mathbf{m}, \mathbf{r}/\mathbf{s}]) \rightarrow \alpha^A[B/Q_2, \mathbf{N}/\mathbf{m}, \mathbf{r}/\mathbf{s}]$

Remarks. In this paper we choose the set of formulas (III) of Section 2.2 as non-logical axioms, which have enough power to prove our aimed agreement properties for various protocols with public keys, such as the Needham-Schroeder-Lowe protocol, ISO/IEC 11770-3 Key Transport Mechanism 6 (cf. Protocol 4.16 in [9]), and so on. With additional syntax and non-logical axioms about shared keys or signature we can also prove correctness of the core part of Kerberos (cf. Protocol 3.25 in [9]) and the ISO 9798-3 protocol (used as an example in [12]). BPL however, does not provide flexible compositional treatment of proofs as Protocol Composition Logic does.

3 Computational Semantics

3.1 Computational Asymmetric Encryption Schemes

The fundamental objects of the computational world are strings, $\text{strings} = \{0, 1\}^*$, and families of probability distributions over strings. These families are indexed by a *security parameter* $\eta \in \mathbb{N}$ (which can be roughly understood as key-lengths). We need the notion of a negligible function: A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is said to be *negligible*, if for any $c > 0$, there is an $n_c \in \mathbb{N}$ such that $|f(\eta)| \leq \eta^{-c}$ whenever $\eta \geq n_c$.

Pairing is an injective *pairing function* $[\cdot, \cdot] : \text{strings} \times \text{strings} \rightarrow \text{strings}$ with a tag on each output string that shows it is a pair. We assume that changing a bit string in any of the argument to another bit string of the same length does not influence the length of the output of the pairing. An encryption scheme is a triple of algorithms $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ with key generation \mathcal{K} , encryption \mathcal{E} and decryption \mathcal{D} . Let plaintexts, ciphertexts, publickey and secretkey be nonempty subsets of strings. The set coins is some probability field that stands for coin-tossing, *i.e.*, randomness.

Definition 2 (Encryption Scheme). A computational asymmetric encryption scheme is a triple $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ where: $\mathcal{K} : \text{param} \times \text{coins} \rightarrow \text{publickey} \times \text{secretkey}$ is a key-generation algorithm with $\text{param} = \mathbb{N}$, $\mathcal{E} : \text{publickey} \times \text{plaintexts} \times \text{coins} \rightarrow \text{ciphertexts}$ is an encryption function, and $\mathcal{D} : \text{secretkey} \times \text{strings} \rightarrow \text{plaintexts}$ is such that for all (e, d) output of $\mathcal{K}(\eta, \cdot)$ and $c \in \text{coins}$ $\mathcal{D}(d, \mathcal{E}(e, m, c)) = m$ for all $m \in \text{plaintexts}$. All these algorithms are computable in polynomial time with respect to the security parameter.

We assume that the length of the output of the encryption depends only on the length of the plaintext. We also assume that a tag is attached to each output bit-string that shows it is a cyphertext. We further assume that the encryption scheme satisfies adaptive chosen ciphertext security (CCA-2) defined the following way:

Definition 3 (Adaptive Chosen Ciphertext Security). A computational public-key encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ provides indistinguishability under the adaptive chosen-ciphertext attack if for all PPT adversaries A and for all sufficiently large security parameter η : $|\Pr[(e, d) \leftarrow \mathcal{K}(1^\eta); b \leftarrow \{0, 1\}; m_0, m_1 \leftarrow A^{\mathcal{D}_1(\cdot)}(1^\eta, e); c \leftarrow \mathcal{E}(e, m_b); g \leftarrow A^{\mathcal{D}_2(\cdot)}(1^\eta, e, c) : b = g] - \frac{1}{2}|$ is a negligible function of η . The oracle $\mathcal{D}_1(x)$ returns $\mathcal{D}(d, x)$, and $\mathcal{D}_2(x)$ returns $\mathcal{D}(d, x)$ if $x \neq c$ and returns \perp otherwise. The adversary is assumed to keep state between the two invocations. It is required that m_0 and m_1 be of the same length. The probability includes all instances of randomness: key generation, the choices of the adversary, the choice of b , the encryption.

It has been shown in [8], that the above definition is equivalent with another that seems stricter at first, namely, when an n -tuple of encryption and decryption oracles are given, each with separate encryption and decryption keys, but using the same bit b to choose from the submitted plaintexts. The adversary is allowed to invoke the oracles in any order but it cannot submit a message that was received from an encryption oracle to the corresponding decryption oracle.

3.2 Stochastic Model for the Computational Execution of BPL

In the following, we discuss the mathematical objects that we use to represent a computational execution of a protocol. Our plan is to define a computational semantics, show that the syntactic axioms hold if the encryption scheme is CCA-2 secure, and, as a result, if the query-form (or anything else) is provable in the syntax, it must be true in any computational model.

The main improvements from computational semantics proposed by Datta et al. were explained in the introduction. Our approach avoids the necessity to tell on each trace whether something is satisfied or not because we want to avoid pure coincidences. Since computational behavior is probabilistic, it is often not natural to ask whether something is satisfied on a particular trace or not. The natural question is whether the traces produce the right kind of probability distribution for a certain variable.

First, since probabilities and complexity are involved, we need a probability space for each value of the security parameter. Since time plays an important role in the execution, what we need is the probability space for a *stochastic process*. For the presentation here, we limit ourselves to finite probability spaces as explaining the notion of measurability and stochastic processes is much simpler this way, but for anyone familiar with these notions in infinite spaces it is near to trivial to generalize the method to allowing infinite steps (but polynomial expected run-time). So, here we assume that for each security parameter, there is a maximum number of execution steps n^η . The following notions that we introduce are standard in probability theory.

We will denote the finite probability space for an execution of a protocol with security parameter η by Ω^η , subsets of which are called events. Let \mathcal{F}^η denote the set of all subsets of Ω^η (including the empty set). A subset containing only one element is called an *elementary event*.

The set Ω^n is meant to include all randomness of an execution of the protocol (and perhaps some additional information). A *probability measure* p^n assigns a probability to each subset such that it is additive with respect to disjoint unions of sets (so it is enough to assign a probability to each element of Ω^n , then the probability of any subset can be computed). When it is clear which probability space we are talking about, we will just use the notation \Pr .

In order to describe what randomness was carried out until step $i \in \{0, 1, \dots, n^n\}$, we assign a subset $\mathcal{F}_i^n \subseteq \mathcal{F}^n$ to each i , such that \mathcal{F}_i^n is closed under union and intersection, and includes \emptyset and Ω^n , and $\mathcal{F}_i^n \subseteq \mathcal{F}_{i+1}^n$. The set $\{\mathcal{F}_i^n\}_{i=1}^{n^n}$ is called *filtration*. Since everything is finite, \mathcal{F}_i^n is *atomistic*, that is, each element of it can be obtained as a union of disjoint, minimal (with respect to inclusion) nonempty elements. The minimal nonempty elements are called *atoms*. We introduce the notation

$$\mathbf{Pr} = \{(\Omega^n, \{\mathcal{F}_i^n\}_{i=0}^{n^n}, p^n)\}_{\eta \in \text{param}}.$$

We included \mathcal{F}_0^n to allow some initial randomness such as key generation. A discrete *random variable* on Ω^n is a function on Ω^n taking some discrete value. Since \mathcal{F}_i^n contains the events determined until step i , a random variable g^n depends only on the randomness until i exactly if g is constant on the atoms of \mathcal{F}_i^n ; this is the same as saying that for any possible value c , the set $[g^n = c] := \{\omega \mid g^n(\omega) = c\}$ is an element of \mathcal{F}_i^n . In this case, we say that g^n is *measurable* with respect to \mathcal{F}_i^n . We will, however need a somewhat more complex dependence-notion. We will need to consider random variables that are determined by the randomness until step i_1 on certain random paths, but until step i_2 on other paths, and possibly something else on further paths. For this, we have to first consider a function $J^n : \Omega^n \rightarrow \{0, 1, \dots, n^n\}$ that tells us which time step to consider on each ω . This function should only depend on the past, so for each $i \in \{0, 1, \dots, n^n\}$, we require that the set $[J^n = i] \in \mathcal{F}_i^n$. We will call this function a *stopping time*. The events that have occurred until the stopping time J^n are contained in

$$\mathcal{F}_{J^n}^n := \{S \mid S \subseteq \Omega^n, \text{ and for all } i = 0, 1, \dots, n^n, S \cap [J^n = i] \in \mathcal{F}_i^n\}.$$

Then, a random variable f^n depends only on the events until the stopping time J^n iff for each c in its range, $[f^n = c] \in \mathcal{F}_{J^n}^n$. Furthermore, a random variable h^n on Ω^n is said to be independent of what happened until J^n iff for any $S \in \mathcal{F}_{J^n}^n$ and a c possible value of h^n , $\Pr([h^n = c] \cap S) = \Pr([h^n = c]) \Pr(S)$. Finally, it is easy to see that for each random variable f^n , there is a stopping time J_f^n such that f^n is measurable with respect to $\mathcal{F}_{J_f^n}^n$, and J_f^n is minimal in the sense that f^n is not measurable with respect to any other \mathcal{F}_J^n if there is an ω such that $J^n(\omega) < J_f^n(\omega)$.

Example 1. Suppose coins are tossed three times, one after the other. Then $\Omega = \{(a, b, c) \mid a, b, c = 0, 1\}$. Let $(1, \cdot, \cdot) := \{(1, b, c) \mid b, c = 0, 1\}$. $(0, \cdot, \cdot)$, etc. are defined analogously. At step $i = 1$, the outcome of the first coin-tossing becomes known. So, $\mathcal{F}_1 = \{\emptyset, (0, \cdot, \cdot), (1, \cdot, \cdot), \Omega\}$. At step $i = 2$, the outcome of the second coin becomes known too, therefore \mathcal{F}_2 , besides \emptyset and Ω , contains $(0, 0, \cdot)$, $(0, 1, \cdot)$, $(1, 0, \cdot)$ and $(1, 1, \cdot)$ as atoms, and all possible unions of these. \mathcal{F}_3 is all subsets. A function g that is measurable with respect to \mathcal{F}_1 , is constant on $(0, \cdot, \cdot)$ and on $(1, \cdot, \cdot)$, that is, g only depends on the outcome of the first coin tossing, but not the rest. Similarly, an f measurable on \mathcal{F}_2 , is constant on $(0, 0, \cdot)$, on $(0, 1, \cdot)$, on $(1, 0, \cdot)$ and on $(1, 1, \cdot)$. A stopping time is for example the J that equals the position of the first 1, or 3 if there is never 1: $J((a_1, a_2, a_3)) = i$ if $a_i = 1$ and $a_k = 0$ for $k < i$, and $J((a_1, a_2, a_3)) = 3$ if $a_k = 0$ for all $k = 1, 2, 3$. The atoms of \mathcal{F}_J are $(1, \cdot, \cdot)$, $(0, 1, \cdot)$, $\{(0, 0, 1)\}$ and $\{(0, 0, 0)\}$.

For each value of the security parameter, an execution of the protocol involves some principals. Each principal has a distinct name, a bit-string not longer than the upper bound n^n . Each principal generates an encryption-key, decryption-key pair at the initialization. Hence, if $\mathbf{Pr} =$

$\{(\Omega^\eta, \{\mathcal{F}_i^\eta\}_{i=0}^{n^\eta}, p^\eta)\}_{\eta \in \text{param}}$ is the stochastic space of the execution of the protocol, let \mathcal{P}^η be a set of (polynomially bounded number of) elements of the form $(A^\eta, (e_A^\eta, d_A^\eta))$ where $A^\eta \in \{0, 1\}^{n^\eta}$, and (e_A^η, d_A^η) is a pair of probability distributions on Ω^η measurable with respect to \mathcal{F}_0^η such that $\Pr[\omega : (e_A^\eta(\omega), d_A^\eta(\omega)) \notin \text{Range}(\mathcal{K}(\eta, \cdot))]$ is a negligible function of η . We assume that if $A = B$, then $(e_A^\eta, d_A^\eta) = (e_B^\eta, d_B^\eta)$. The set $\{\mathcal{P}^\eta\}_{\eta \in \text{param}}$ describes all the principals, corrupted and uncorrupted, that take part in the execution at a given security parameter, along with their public and secret keys. Let $\mathcal{P} = \{\mathcal{P}^\eta\}_{\eta \in \text{param}}$.

For nonces, we choose the following definition. Since CCA-2 security is length-revealing, we have to assume that nonces are always of some fixed length m^η for each security parameter η . Let \mathcal{N} be a set of elements of the form $\{N^\eta\}_{\eta \in \text{param}}$ where $N^\eta : \Omega^\eta \rightarrow \{0, 1\}^{m^\eta} \cup \{\perp\}$ (taking the value \perp means N^η has no bit-string value on that particular execution), such that over $\{0, 1\}^{m^\eta}$, N^η is uniformly distributed, and for any two $\{N_1^\eta\}_{\eta \in \text{param}}, \{N_2^\eta\}_{\eta \in \text{param}} \in \mathcal{N}$, N_1^η and N_2^η are independent (i.e. for any two $s_1, s_2 \in \{0, 1\}^{m^\eta}$, $\Pr[N_1^\eta = s_1 \wedge N_2^\eta = s_2] = \Pr[N_1^\eta = s_1] \Pr[N_2^\eta = s_2]$). This set describes the nonces that were generated with overwhelming probability during the execution of the protocol. The nonces have to be independent of each other, and have uniform distribution over the given length. The nonces also have to be independent of what happened earlier when they are being generated, but we will require this later.

Let \mathcal{R} be a set of elements of the form $R = \{R^\eta\}_{\eta \in \text{param}}$ where $R^\eta : \Omega^\eta \rightarrow \text{coins} \cup \{\perp\}$. Let \mathcal{R}_g be the subset of \mathcal{R} which are properly randomized, that is, for which the values in coins have the distribution required for the encryption scheme (on the condition that the value is not \perp). That is, they have good distribution (hence the g).

Messages: Let the set of messages be \mathcal{M} elements of the form $M = \{M^\eta\}_{\eta \in \text{param}}$, where $M^\eta : \Omega^\eta \rightarrow \{0, 1\}^{n^\eta} \cup \{\perp\}$. For any two messages, M_1, M_2 , we will denote that $M_1 \approx M_2$ iff $p^\eta[\omega : M_1(\omega) \neq M_2(\omega)]$ is a negligible function of η . This way, \approx is an equivalence notion on the set of messages. Let $D_M := \mathcal{M}/\approx$, let $D_N := \mathcal{N}/\approx \subset D_M$, and let

$$D_P := \{A \in \mathcal{M} : (A^\eta, (e_A^\eta, d_A^\eta)) \in \mathcal{P} \text{ for some } (e_A^\eta, d_A^\eta)\}/\approx \subset D_M$$

We have to define what we mean by a computational pairing and encryption. For any $X, X_1, X_2 \in D_M$, we write that $X =_C \langle X_1, X_2 \rangle$, if for some (hence for all) $M_1 = \{M_1^\eta\}_{\eta \in \text{param}} \in X_1$ and $M_2 = \{M_2^\eta\}_{\eta \in \text{param}} \in X_2$, the ensemble of random variables $\{\omega \mapsto [M_1^\eta(\omega), M_2^\eta(\omega)]\}_{\eta \in \text{param}}$ is an element of X . Further, if $A \in \mathcal{P}$, and $R \in \mathcal{R}$, then we will write that $X =_C \{X_1\}_A^R$ if for any (hence for all) $M_1 = \{M_1^\eta\}_{\eta \in \text{param}} \in X_1$, the ensemble of random variables $\{\omega \mapsto \mathcal{E}(e_A^\eta(\omega), M_1^\eta(\omega), R(\omega))\}_{\eta \in \text{param}}$ is an element of X . If the value of any of the input distributions is \perp then we take the output to be \perp as well. This way, we can consider an element of the free term algebra $T(D_M)$ over D_M as an element of D_M . Let $\sqsubseteq_{T(D_M)}$ denote the subterm relation on $T(D_M)$. This generates a subterm relation \sqsubseteq_C on D_M by defining $X_1 \sqsubseteq_C X_2$ to be true iff there is an element $X \in T(D_M)$ such that $X_1 \sqsubseteq_{T(D_M)} X$ and $X_2 =_C X$.

For any set of subsets $\mathcal{D}^\eta \in \mathcal{F}^\eta$, $\mathcal{D} = \{\mathcal{D}^\eta\}_{\eta \in \eta}$ with non-negligible $p^\eta(\mathcal{D}^\eta)$, we say that for $X_1, X_2 \in D_M$, $X_1 = X_2$ on \mathcal{D} if there are $M_1 = \{M_1^\eta\}_{\eta \in \text{param}} \in X_1$ and $M_2 = \{M_2^\eta\}_{\eta \in \text{param}} \in X_2$ with $M_1^\eta(\omega) = M_2^\eta(\omega)$ for all $\omega \in \mathcal{D}^\eta$. We say that $X_1 \sqsubseteq_C X_2$ on \mathcal{D} iff there is an element $X \in T(D_M)$ such that $X_1 \sqsubseteq_{T(D_M)} X$ and $X_2 =_C X$ on \mathcal{D} .

Execution trace: The execution trace is defined as $Tr = \{Tr^\eta\}_{\eta \in \text{param}}$ where $Tr^\eta : \omega \mapsto Tr^\eta(\omega)$ with either

$$Tr^\eta(\omega) = P_1^\eta(\omega) \text{ acts}_1^\eta(\omega) s_1^\eta(\omega); \dots; P_{n^\eta(\omega)}^\eta(\omega) \text{ acts}_{n^\eta(\omega)}^\eta(\omega) s_{n^\eta(\omega)}^\eta(\omega)$$

where for each η security parameter, $\omega \in \Omega^\eta$, $n^\eta(\omega)$ is a natural number less than n^η , $P_i^\eta(\omega) \in D_P$, $\text{acts}_i^\eta(\omega)$ is one of *generates*, *sends*, *receives* and $s_i^\eta(\omega) \in \{0, 1\}^*$; or $Tr^\eta(\omega) = \perp$

with $n^\eta(\omega) = 0$ meaning that no generate, send or receive action happened. For each η , ω , and $i \in \{1, \dots, n^\eta\}$, let

$$Tr_i^\eta(\omega) = \begin{cases} P_i^\eta(\omega) \text{ acts}_i^\eta(\omega) s_i^\eta(\omega) & \text{if } i \in \{1, \dots, n^\eta(\omega)\} \\ \perp & \text{otherwise} \end{cases}$$

We also require that Tr_i^η be measurable with respect to \mathcal{F}_i^η for all i . Moreover, we require that any of Tr is PPT computable from the earlier ones.

3.3 Computational Semantics

We now explain how to give computational semantics to the syntax, and what it means that a formula of the syntax is true in the semantics. For a given security parameter, an *execution* is played by a number of participants.

Assumptions. In a particular execution, we assume that the principals corresponding to names in the syntax (that is, they correspond to elements in $\mathcal{C}_{\text{name}}$) are *regular* (non-corrupted). We assume that these participants generate their keys and encrypt correctly (that is, the keys are properly distributed, and also r^A is properly randomized) with a CCA-2 encryption scheme, and never use their private keys in any computation except for decryption. For other participants (possibly corrupted), we do not assume this. (Encrypting correctly is essential to able to prove the nonce verification axioms.) We further assume that pairing of any two messages differs from any nonce and from any principal name on sets of non-negligible probability. The network is completely controlled by an adversary. The sent and received bit strings are recorded in a trace in the order they happen. Freshly generated bit-strings produced by the regular participants are also recorded. The combined algorithms of the participants and the adversary are assumed to be probabilistic polynomial time.

Such a situation, with the definitions of the previous section, produces a *computational trace structure associated to the execution* of the form

$$\mathfrak{M} = (II, [\cdot, \cdot], \mathbf{Pr}, \mathcal{P}, \mathcal{N}, \mathcal{R}_g, Tr, \Phi_{\mathcal{C}}, \mathcal{D}),$$

where $\Phi_{\mathcal{C}}$ is a one-to-one function on $\mathcal{C}_{\text{name}} \cup \mathcal{C}_{\text{nonce}} \cup \mathcal{C}_{\text{coin}}$ such that **(i)** $\Phi_{\mathcal{C}}(A) \in D_P$ for any $A \in \mathcal{C}_{\text{name}}$ such that $(e_{\Phi_{\mathcal{C}}(A)}^\eta, d_{\Phi_{\mathcal{C}}(A)}^\eta)$ is measurable with respect to \mathcal{F}_0 and has the correct key distribution, and for different constants are independent of each other; **(ii)** $\Phi_{\mathcal{C}}(N) \in D_N$ for any $N \in \mathcal{C}_{\text{nonce}}$; **(iii)** $\Phi_{\mathcal{C}}(r) \in \mathcal{R}_g$ for any $r \in \mathcal{C}_{\text{coin}}$; and $\mathcal{D} = \{\mathcal{D}^\eta\}_{\eta \in \eta}$, $\mathcal{D}^\eta \in \mathcal{F}^\eta$ a sequence of subsets where we focus our attention with $p^\eta(\mathcal{D}^\eta)$ non-negligible. An *extension* of $\Phi_{\mathcal{C}}$ to evaluation of free variables is a function Φ that is the same on constants as $\Phi_{\mathcal{C}}$, and for variables Q , n , m , s^A , s of sort name, nonce, message, coin_A and coin respectively, $\Phi(Q) \in D_P$, $\Phi(n) \in D_N$, $\Phi(m) \in D_M$, $\Phi(s^A) \in \mathcal{R}_g$ and $\Phi(s) \in \mathcal{R}$ hold. Then, for any t term, $\Phi(t) \in D_M$ is defined on terms as **(i)** $\Phi(\langle t_1, t_2 \rangle) = \langle \Phi(t_1), \Phi(t_2) \rangle$; **(ii)** $\Phi(\{t\}_P^r) = \{\Phi(t)\}_{\Phi(P)}^{\Phi(r)}$; where, as we mentioned earlier, elements of $T(D_M)$ are considered as elements of D_M .

We say that an ensemble of random variables $M = \{M^\eta\}_{\eta \in \text{param}}$ such that M^η is defined on \mathcal{D}^η is a *realization* of the term t through Φ on \mathcal{D} , which we denote $M \lll_{\Phi, \mathcal{D}} t$, if there is an $M_1 \in \Phi(t)$ with $M_1^\eta(\omega) = M^\eta(\omega) \neq \perp$ for all $\omega \in \mathcal{D}^\eta$; and if also $t = \{t'\}_P^{\rho^A}$, then we further require that there is an $M' \in \Phi(t')$ such that $M' \lll_{\Phi, \mathcal{D}} t'$ and $\Phi(\rho^A)^\eta$ on $\{0, 1\}^{m^\eta}$ is independent of $\mathcal{F}_{M'}^\eta$ on the condition that $\Phi(\rho^A)^\eta \neq \perp$ (where for $J_{M'}^\eta$ see the paragraph before Example 1).

In the following, we give the interpretation of BPL. Note, that the interpretation of conjunction, disjunction, negation and conclusion are defined in the most standard manner. We first define when a formula φ is *satisfied* by Φ (remember, $\mathcal{D} = \{\mathcal{D}^\eta\}_{\eta \in \eta}$ with $\mathcal{D}^\eta \in \mathcal{F}^\eta$ from \mathfrak{M}):

- For any terms t_1, t_2 , $\varphi \equiv t_1 = t_2$ is satisfied by Φ , iff $\Phi(t_1) = \Phi(t_2)$, and $\varphi \equiv t_1 \sqsubseteq t_2$ is satisfied by Φ on \mathcal{D} iff $\Phi(t_1) \sqsubseteq_C \Phi(t_2)$ on \mathcal{D} .
- For any term u and $acts = sends/receives$, $\varphi \equiv P acts u$ is satisfied by Φ iff there are stopping times J^η such that apart from sets of negligible probability, $Tr_{J^\eta(\omega)}^\eta(\omega)$ is of the form $A^\eta acts M^\eta(\omega)$ for $\omega \in \mathcal{D}^\eta$ where $M := \{M^\eta\}_{\eta \in \text{param}} \lll_{\Phi, \mathcal{D}} u$ and $A := \{A^\eta\}_{\eta \in \text{param}} \lll_{\Phi, \mathcal{D}} P$. We will denote this as $Tr_J \lll_{\Phi, \mathcal{D}} P acts u$.
- If $acts = generates$ then the u above is a nonce ν , and so $M := \{M^\eta\}_{\eta \in \text{param}} \lll_{\Phi, \mathcal{D}} u$ means there is an $N \in \Phi(\nu)$ such that $M^\eta|_{\mathcal{D}^\eta} = N^\eta|_{\mathcal{D}^\eta}$ in this case, and we further require that N^η be *independent* up to negligible probability of $\mathcal{F}_{J_{-1}^\eta}$ for all η on the condition that $[N \neq \perp]$. (More precisely, there is an $N' \approx N$ such that N'^η is independent of $\mathcal{F}_{J_{-1}^\eta}$ on $[N' \neq \perp]$.)
- $\varphi \equiv \beta_1, \dots, \beta_n$ sequence of actions is satisfied by Φ if each of β_k ($k = 1, \dots, n$) is satisfied by Φ , and if J_k is the stopping time belonging to β_k , then we require that $J_k < J_l$ on \mathcal{D} whenever $k < l$ (that is, for each $\eta \in \text{param}$ and $\omega \in \mathcal{D}^\eta$, $J_k^\eta(\omega) < J_l^\eta(\omega)$).
- For any formulas $\varphi, \varphi_1, \varphi_2$, $\neg\varphi$ is satisfied by Φ iff φ is not satisfied by Φ ; $\varphi_1 \vee \varphi_2$ is satisfied by Φ iff φ_1 is satisfied by Φ or φ_2 is satisfied by Φ ; $\varphi_1 \wedge \varphi_2$ is satisfied by Φ iff φ_1 is satisfied by Φ and φ_2 is also satisfied by Φ . $\varphi_1 \rightarrow \varphi_2$ is satisfied by Φ iff $\neg\varphi_1 \vee \varphi_2$ is satisfied by Φ .
- If φ is a formula, m' a bound variable, and φ' is obtained from φ by substituting m' for every occurrence in φ of some free variable m of the same sort as m' , then $\forall m' \varphi'$ (or $\exists m' \varphi'$, resp.) is satisfied by Φ iff φ is satisfied by each (or some, resp.) Φ' extension of Φ_C when Φ' only differs from Φ on m .

A formula φ is *true* in the structure \mathfrak{M} , iff φ is satisfied by every Φ extension of Φ_C .

If in a structure, the Basic Protocol Logic axioms are true (in which case the structure is called *model*), then by standard arguments of first order logic, it follows that everything provable in the syntax is true in the model. In particular, if the query form is provable in the syntax, then it must be true in any model. We now turn our attention to whether the axioms are satisfied by a structure.

Truth of the Term axioms.

- (a) These axioms are true since if terms are equal in the free algebra $\bar{\mathcal{A}}$, then their interpretations are also equal, no matter how Φ is extended to variables. Further, if $t \sqsubseteq t'$ holds in the free algebra, then the way we receive t' from t by pairing and encryptions carries over to the computational world, no matter how Φ is evaluated on variables.
- (b) These axioms hold as computational equality is also symmetric, reflexive and transitive. Further, subterm relation is also transitive for the interpretations, and equality implies computation subterm relation by definition of computational subterm.
- (c) If the interpretations of $\{t_1\}_Q^s$ and $\{t_2\}_Q^{s'}$ are computationally equal up to negligible probability, then the interpretations of t_1 and t_2 must also be equal up to negligible probability as $\Phi(t_1) = \mathcal{D}(d_{\Phi(Q)}, \Phi(\{t_1\}_Q^s))$ and $\Phi(t_2) = \mathcal{D}(d_{\Phi(Q)}, \Phi(\{t_2\}_Q^{s'}))$ and the right-hand sides are equal up to negligible probability.
- (d) First observe, that it cannot happen that both $\{t_1\}_A^{s_C} \sqsubseteq t_2$ and $\{t_2\}_B^{s_D} \sqsubseteq t_1$ are true in $\bar{\mathcal{A}}$. So assume that $\{t_1\}_A^{s_C} \not\sqsubseteq t_2$ in $\bar{\mathcal{A}}$. Since A and C are honest participants, the interpretation of the encryption $\{t_1\}_A^{s_C}$ is CCA-2 secure. Therefore, $\Phi(\{t_1\}_A^{s_C}) = \Phi(\{t_2\}_B^{s_D})$ is only possible if $A = B$ and $C = D$ (and $\Phi(t_1) = \Phi(t_2)$), otherwise CCA-2 security would be broken as $\Phi(\{t_2\}_B^{s_D})$ would reproduce the encryption $\Phi(\{t_1\}_A^{s_C})$ without using it because $\{t_1\}_A^{s_C} \not\sqsubseteq t_2$ in $\bar{\mathcal{A}}$, and t_2 is built of constants.
- (e) Soundness of this axiom follows as we supposed that computational pairing is one-to-one.
- (f) These follow as we assumed that the outputs of computational encryption and pairing are tagged.

- (g) Follows from tagging.
- (h) Soundness of the first formula follows as if $t \sqsubseteq \langle t_1, t_2 \rangle$ is satisfied, then either the interpretations of the two sides are equal (up to negligible probability) and hence $t = \langle t_1, t_2 \rangle$ is satisfied, or the interpretation of $\langle t_1, t_2 \rangle$ can be received from the interpretation of t via encryptions and pairing, of which the last has to be pairing because the tags have to match; then by soundness of (e), it follows that the paired items must in fact be interpretations of t_1 and t_2 , which implies that either of the interpretations of t_1 or of t_2 was received from the interpretation of t via pairing and encryptions, which means that either $t \sqsubseteq t_1$ or $t \sqsubseteq t_2$ is satisfied, and that proves the soundness of this formula. As for the second formula, if $t_1 \sqsubseteq \{t_2\}_Q^s$ is satisfied, then either the interpretations of the two sides are equal, or the interpretation of $\{t_2\}_Q^s$ can be received from the interpretation of t_1 via encryptions and pairing, of which the last has to be encryption because the tags have to match, and so soundness follows.
- (i) Also follows from tagging.

Truth of the Ordering axiom. Suppose that there is an extension Φ and a domain \mathcal{D} such that the formula $Q_2 \text{ sends } m; Q_1 \text{ generates } n$ is satisfied on \mathcal{D} with non-negligible probability as well as the formula $n \sqsubseteq m$. Then, there are stopping times J_1^η, J_2^η such that $\text{Tr}_{J_1} \lll_{\Phi, \mathcal{D}} Q \text{ sends } m$, and $\text{Tr}_{J_2} \lll_{\Phi, \mathcal{D}} Q \text{ generates } n$, and $J_1 < J_2$. Then, $\text{Tr}_{J_1} \lll_{\Phi, \mathcal{D}} Q \text{ sends } m$ implies that there is $M \lll_{\Phi, \mathcal{D}} m$ such that M^η is measurable with respect to $\mathcal{F}_{J_1}^\eta$ and since $n \sqsubseteq m$ is satisfied, some $N \in \Phi(n)$ can be obtained as a series of decryptions and breaking up pairs from M . Since there is no new randomness used there, N^η only depends on the randomness until J_1 , so N^η is measurable with respect to $\mathcal{F}_{J_1}^\eta$. But, $\text{Tr}_{J_2} \lll_{\Phi, \mathcal{D}} Q \text{ generates } n$ implies that $\Phi(n)$ has an element N'^η measurable with respect to $\mathcal{F}_{J_2}^\eta$ and independent of $\mathcal{F}_{J_2-1}^\eta$ on $[N'^\eta \neq \perp]$, and hence independent of $\mathcal{F}_{J_1}^\eta$ and of N^η on $[N'^\eta \neq \perp]$. So, N and N' only differ up to negligible probability, but N^η and N'^η are independent for all η , which is possible only if $\Pr[N'^\eta \neq \perp]$ is negligible. That means \mathcal{D} has negligible probability, a contradiction.

Truth of the Nonce verification axioms. In order to show that the axioms are satisfied, we use the assumption that regular participants encrypt with a CCA-2 secure encryption scheme. Suppose there is a Φ and non-negligible \mathcal{D} such that the premise of the axioms are satisfied by Φ on \mathcal{D} , but the conclusion is not. Then, if the conclusion is not satisfied, that means that with non-negligible probability, $\{m_6\}_B^{r_A}$ does not go through B . The premise however says that n_1 shows up in m_5 later, which can be recovered from there up to negligible probability via a series of de-coupling and decryption such that $\{m_6\}_B^{r_A}$ does not have to be decrypted. We have to show that a PPT algorithm can be constructed that breaks CCA-2 security. The algorithm that breaks CCA-2 security is simply the protocol execution itself with the following modifications: **(1)** The decryption oracle (that the algorithm may access according to the definition of CCA-2 security) does the job for all decryptions with the private key d_B . **(2)** The algorithm generates two samples of n_1 when the protocol execution samples n_1 . **(3)** When the protocol execution is to produce $\{m_6\}_B^{r_A}$, compute two samples of the realization of m_6 using the two samples of n_1 and using the same samples for the other parts of m_6 . **(4)** Submit to the encryption oracle of the CCA-2 game the pair of samples of m_6 , and use the ciphers that it outputs whenever $\{m_6\}_B^{r_A}$ occurs again. **(5)** If the sample for $\{m_6\}_B^{r_A}$ goes through B , terminate. If not, continue until the Q receives the sample for m_5 . **(6)** Recover the sample for n_1 via de-coupling and decryption using the decryption oracle if necessary. The bit string hence obtained is the one that was in the plaintext encrypted by the oracle, so the bit value b of the game can be determined. If the conclusion of the axiom is not satisfied, then this algorithm has non-negligible probability of winning the CCA-2 game as \mathcal{D} is non-negligible.

In order to show the validity of the second nonce-verification axiom, we have to use the modified version of CCA-2 (equivalent to the original) when there are two encryption - decryption pairs

of oracles, each corresponding to independently generated encryption key - decryption key pairs. The algorithm then is the following: **(1)** The decryption oracles (that the algorithm may access according to the modified definition of CCA-2 security) do the job for all decryptions with the private keys d_B and d_C . **(2)** The algorithm generates two samples of n_1 when the protocol execution samples n_1 . **(3)** When the protocol execution is to produce $\{m_6\}_B^{r^A}$, compute two samples of the realization of m_6 using the two samples of n_1 and using the same samples for the other parts of m_6 . **(4)** Submit to the first encryption oracle of the CCA-2 game the pairs of samples of m_6 . **(5)** Skip the step when B decrypts $\{m_6\}_B^{r^A}$. **(6)** When $\{m_8\}_C^{r^B}$ is constructed, compute two samples of m_8 just as in the case of m_6 . Stop if the samples have different length, otherwise submit the results to the second encryption oracle. **(7)** Continue until C receives the sample for m_5 . **(8)** Recover the sample for n_1 via de-coupling and decryption using the decryption oracle if necessary. The bit string hence obtained is the one that was in the plaintext encrypted by the oracles, so the bit value b of the game can be determined. This is again PPT algorithm given that the protocol execution was PPT, so it breaks CCA-2 security. Therefore, the Nonce- verification axioms hold.

Soundness Since the axioms are true in the structure \mathfrak{M} , by a standard argument of first order logic, the following theorem is true:

Theorem 1. *With our assumptions on the execution of the protocol, if the associated computational trace structure is $\mathfrak{M} = (\Pi, [\cdot, \cdot], \mathbf{Pr}, \mathcal{P}, \mathcal{N}, \mathcal{R}_g, Tr, \Phi_C, \mathcal{D})$, then, if a formula (the query form in particular) is provable in the syntax with first-order predicate logic and axioms (I), (II), (III), then it is true in \mathfrak{M} .*

Proof. We have showed that the term axioms and non-logical axioms of BPL are true in the model. It is routine to check that all the logical axioms and logical inference rules of first order logic are also true in the model, because we followed the usual first-order logical operations of composed formulas in the interpretation. Hence the theorem holds.

We would like to reflect again on the four points in the introduction: 1. Removing the bound n^n from the length of executions is a trivial step (change the finite sequence of the filtration to an infinite one, and the definition of measurability to the standard one for infinite spaces) in our framework, only the presentation of the definition of measurability is more involved in this case, that is why we chose to stick to the bound. 2. We did not introduce modal formulas here in the syntax, and it is our work in progress to extend our approach to PCL. As we keep track of the actual probability distributions and correlations, it should be no problem to define the semantics of modal formulas so that these axioms hold as long as the interpretations (distributions, not bit-strings) of b and $\text{Send}(\tilde{X}, t)$ are different up to negligible probability. 3. As we use filtrations, according to which random variables have to be measurable, dependence on the future is taken care of by measurability. 4. We required that the distribution of keys are measurable with respect to \mathcal{F}_0^n , and generated nonces are independent of the past, so the anomaly mentioned in the introduction here cannot happen as N and K must have independent interpretations. The reader may be worried that we don't require that the generated R has to be dependent of N as R is generated by the adversary or a corrupted participant. It is true that we could introduce another filtration that indicates the knowledge of the adversary up to a certain time, which may be needed in a more complex syntax (for example if we allow corrupted participants to generate their keys sometime in the middle), however, in BPL this is not necessary as this does not result in undesired coincidences and the proofs work even without this tool.

4 Conclusions

We have given a computational semantics to Basic Protocol Logic that uses stochastic structures, and showed a soundness theorem. In order to show that the axioms of BPL were true in the semantics, we had to modify BPL as the original axioms were not all computationally sound. We showed our method on BPL as it is simple enough for a first, concise presentation. We argued why this semantics looks more promising than the one by Datta et al. Next, we would like to apply our methods to the much more complex formal syntax of Protocol Composition Logic. A formal completeness theorem for BPL has also been provided in [19], we would like to investigate completeness in the computational case as well.

References

1. Website of www.stanford.edu/~danupam/logic-derivation.html.
2. M. Abadi, M. Baudet, and B. Warinschi. Guessing attacks and the computational soundness of static equivalence. In *Proceedings of FoSSaCS'06*, Lecture Notes in Computer Science. Springer-Verlag, March-April 2006.
3. M. Abadi and P. Rogaway. Reconciling two views of cryptography. *Journal of Cryptology*, 15(2):103–127, January 2002.
4. P. Adão, G. Bana, J. Herzog, and A. Scedrov. Soundness of formal encryption in the presence of key-cycles. In *Proceedings of ESORICS'05*, volume 3679 of *Lecture Notes in Computer Science*, pages 374–396, Milan, Italy, September 12–14 2005. Springer.
5. P. Adão, G. Bana, and A. Scedrov. Computational and information-theoretic soundness and completeness of formal encryption. In *Proceedings of CSFW'05*, pages 170–184, Aix-en-Provence, France, June 20–22 2005. IEEE Computer Society Press.
6. M. Backes, B. Pfizmann, and M. Waidner. A composable cryptographic library with nested operations. In *Proceedings of CCS'03*, pages 220–230, Washington D.C., USA, October 27–30 2003. ACM Press.
7. M. Baudet, V. Cortier, and S. Kremer. Computationally sound implementations of equational theories against passive adversaries. In *Proceedings of ICALP'05*, volume 3580 of *Lecture Notes in Computer Science*, pages 652–663. Springer-Verlag, July 2005.
8. M. Bellare, A. Boldyreva, and S. Micali. Public-key encryption in a multi-user setting. In *Proceedings of EUROCRYPT'00*, pages 258–274. Springer-Verlag, 2000.
9. C. Boyd and Anish Mathuria. *Protocols for Authentication and Key Establishment*. Springer, 2003.
10. Ran Canetti and Jonathan Herzog. Universally composable symbolic analysis of mutual authentication and key exchange protocols. In *Proceedings, Theory of Cryptography Conference (TCC)*, March 2006.
11. I. Cervesato, C. Meadows, and D. Pavlovic. An encapsulated authentication logic for reasoning about key distribution protocols. In *Proceedings of CSFW'05*, pages 48–61, 2005.
12. A. Datta, A. Derek, J. C. Mitchell, and D. Pavlovic. A derivation system and compositional logic for security protocols. *Journal of Computer Security (Special Issue of Selected Papers from CSFW-16)*, 13:423–482, 2005.
13. A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, and M. Turuani. Probabilistic polynomial-time semantics for a protocol security logic. In *Proceedings of ICALP'05*, volume 3580 of *Lecture Notes in Computer Science*, pages 16–29, Lisbon, Portugal, July 11–15 2005. Springer.
14. D. Dolev and A. C. Yao. On the security of public-key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, March 1983. Preliminary version presented at FOCS'81.
15. N.A. Durgin, J.C. Mitchell, and D. Pavlovic. A compositional logic for proving security properties of protocols. *Journal of Computer Security*, 11:677–721, 2003.
16. S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and Systems Sciences*, 28(2):270–299, April 1984. Preliminary version presented at STOC'82.
17. J. D. Guttman and F. J. Thayer Fábrega. Authentication tests. In *IEEE Symposium on Security and Privacy*, pages 96–109, 2002.

18. J. D. Guttman, F. J. Thayer, and L. D. Zuck. The faithfulness of abstract protocol analysis: Message authentication. In *Proceedings of CCS'01*, pages 186–195. ACM Press, November 05–08 2001.
19. K. Hasebe and M. Okada. Completeness and counter-example generations of a basic protocol logic. In *Proceedings of RULE'05*, volume 147(1), pages 73–92. Elsevier Science, 2005. Available also at <http://abelard.flet.keio.ac.jp/person/hasebe/downloads/rule05.pdf> or <http://dx.doi.org/10.1016/j.entcs.2005.06.038>.
20. P. Laud. Symmetric encryption in automatic analyses for confidentiality against active adversaries. In *Proceedings of S&P'04*, pages 71–85. IEEE Computer Society Press, May 9–12 2004.
21. D. Micciancio and B. Warinschi. Completeness theorems for the Abadi-Rogaway logic of encrypted expressions. *Journal of Computer Security*, 12(1):99–130, 2004.

A Definition of *Only* and *Honest*

Our aimed correctness properties are described in a special form of formulas, called *query form*. Let $\alpha^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}]$ be a role A $acts_{s_1} t_1; A acts_{s_2} t_2; \dots; A acts_{s_k} t_k$ where each $acts_i$ ($1 \leq i \leq k$) is one of *sends*, *receives* and *generates*, t_i is a term built from messages in $\mathbf{m} = \{m_1, \dots, m_h\}$ (some of which may be nonces) from coins in $\mathbf{s} = \{s_1, \dots, s_i\}$ and from names A and $\mathbf{Q} = \{Q_1, \dots, Q_l\}$. Let $\alpha_{\leq i}^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}]$ denote an initial segment of $\alpha^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}]$ ending with $A acts_i t_i$ (for $1 \leq i \leq k$), i.e., $\alpha_{\leq i}^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}] \equiv A acts_{s_1} t_1; \dots A acts_i t_i$. Let $\alpha_{\leq 0}^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}] \equiv A = A$.

The query form includes a formalization of *principal's honesty* $Honest(\alpha^A)$, which is defined as follows, the intuitive meaning being that A follows the role α^A and does nothing else, but it may not complete it:

$$Honest(\alpha^A)[\mathbf{Q}, \mathbf{m}, \mathbf{s}] \stackrel{\text{def}}{\equiv} \exists \mathbf{Qms} \bigvee_{i \in \{0\} \cup \{j \mid acts_j = \text{sends}\} \cup \{k\}} \alpha_{\leq i}^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}] \wedge Only(\alpha_{\leq i}^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}])$$

$Only(\alpha^A)$ denotes the following formula, whose intuitive meaning is “ A performs only α^A ”.

$$Only(\alpha^A) \equiv \forall n (A \text{ generates } n_1 \rightarrow n \in Generates(\alpha^A)) \wedge \forall m_1 (A \text{ sends } m_1 \rightarrow m_1 \in Sends(\alpha^A)) \\ \wedge \forall m_2 (A \text{ receives } m_2 \rightarrow m_2 \in Receives(\alpha^A))$$

Here, $Sends(\alpha^A)$ denotes the set $\{t_j \mid A \text{ sends } t_j \subseteq \alpha^A\}$, and $(Receives(\alpha^A), Generates(\alpha^A))$ are defined similarly. Set theoretical notation as $m \in Sends(\alpha^A)$ (as well as $m \in Receives(\alpha^A)$ and $m \in Generates(\alpha^A)$) is an abbreviation of a disjunctive form: for example, if $Sends(\alpha^A) = \{t'_1, \dots, t'_j\}$, then $m \in Sends(\alpha^A)$ denotes the formula $(m = t'_1) \vee (m = t'_2) \vee \dots \vee (m = t'_j)$. (As a special case, if $Sends(\alpha^A)$ is empty then $m \in Sends(\alpha^A)$ denotes $A \neq A$, that is, impossible.)

Intuitively, each disjunct $\alpha_{\leq i}^A \wedge Only(\alpha_{\leq i}^A)$ in $Honest(\alpha^A)$ represents a historical record of P 's actions at each step of his run: the sequence of actions $\alpha_{\leq i}^A$ represents A 's performance until this step, and $Only(\alpha_{\leq i}^A)$ represents that A performs only $\alpha_{\leq i}^A$. $Only(\alpha_{\leq 0}^A)$ means that nothing was performed. Thus, $Honest(\alpha^A[\mathbf{Q}, \mathbf{m}, \mathbf{s}])$ represents “ A performs only a run of an initial segment of α^A which ends with a sending action or the last action of α^A , and uses the data items \mathbf{Q}, \mathbf{m} and \mathbf{s} for each run”.