
MODELING ADVERSARIES IN A LOGIC FOR SECURITY PROTOCOL ANALYSIS

JOSEPH Y. HALPERN ^a AND RICCARDO PUCELLA ^b

^a Cornell University, Ithaca, NY 14853
e-mail address: halpern@cs.cornell.edu

^b Northeastern University, Boston, MA 02115
e-mail address: riccardo@ccs.neu.edu

ABSTRACT. Logics for security protocol analysis require the formalization of an adversary model that specifies the capabilities of adversaries. A common model is the Dolev-Yao model, which considers only adversaries that can compose and replay messages, and decipher them with known keys. The Dolev-Yao model is a useful abstraction, but it suffers from some drawbacks: it cannot handle the adversary knowing protocol-specific information, and it cannot handle probabilistic notions, such as the adversary attempting to guess the keys. We show how we can analyze security protocols under different adversary models by using a logic with a notion of algorithmic knowledge. Roughly speaking, adversaries are assumed to use algorithms to compute their knowledge; adversary capabilities are captured by suitable restrictions on the algorithms used. We show how we can model the standard Dolev-Yao adversary in this setting, and how we can capture more general capabilities including protocol-specific knowledge and guesses.

1. INTRODUCTION

Many formal methods for the analysis of security protocols rely on specialized logics to rigorously prove properties of the protocols they study.¹ Those logics provide constructs for expressing the basic notions involved in security protocols, such as secrecy, recency, and message composition, as well as providing means (either implicitly or explicitly) for describing the evolution of the knowledge or belief of the principals as the protocol progresses. Every such logic aims at proving security in the presence of hostile adversaries. To analyze the effect of adversaries, a security logic specifies (again, either implicitly or explicitly) an *adversary model*, that is, a description of the capabilities of adversaries. Almost all existing logics are based on a Dolev-Yao adversary model (Dolev and Yao, 1983). Succinctly, a

1998 ACM Subject Classification: D.4.6, F.4.1.

Key words and phrases: Protocol analysis, security, attacker models, Dolev-Yao model, epistemic logic, algorithmic knowledge.

This work was done while the second author was at Cornell University.

¹Here, we take a very general view of logic, to encompass formal methods where the specification language is implicit, or where the properties to be checked are fixed, such as Casper (Lowe, 1998), Cryptyc (Gordon and Jeffrey, 2003), or the NRL Protocol Analyzer (Meadows, 1996).

Dolev-Yao adversary can compose messages, replay them, or decipher them if she knows the right keys, but cannot otherwise “crack” encrypted messages.

The Dolev-Yao adversary is a useful abstraction, in that it allows reasoning about protocols without worrying about the actual encryption scheme being used. It also has the advantage of being restricted enough that interesting theorems can be proved with respect to security. However, in many ways, the Dolev-Yao model is too restrictive. For example, it does not consider the information an adversary may infer from properties of messages and knowledge about the protocol that is being used. To give an extreme example, consider what we will call the *Duck-Duck-Goose* protocol: an agent has an n -bit key and, according to her protocol, sends the bits that make up its key one by one. Of course, after intercepting these messages, an adversary will know the key. However, there is no way for security logics based on a Dolev-Yao adversary to argue that, at this point, the adversary knows the key. Another limitation of the Dolev-Yao adversary is that it does not easily capture probabilistic arguments. After all, the adversary can always be lucky and just *guess* the appropriate key to use, irrespective of the strength of the encryption scheme.

The importance of being able to reason about adversaries with capabilities beyond those of a Dolev-Yao adversary is made clear when we look at the subtle interactions between the cryptographic protocol and the encryption scheme. It is known that various protocols that are secure with respect to a Dolev-Yao adversary can be broken when implemented using encryption schemes with specific properties (Moore, 1988), such as encryption systems with encryption cycles (Abadi and Rogaway, 2002) and ones that use exclusive-or (Ryan and Schneider, 1998). A more refined logic for reasoning about security protocols will have to be able to handle adversaries more general than the Dolev-Yao adversary.

Because they effectively build in the adversary model, many formal methods for analyzing protocols are not able to reason directly about the effect of running a protocol against adversaries with properties other than those built in. Some formal methods allow much flexibility in their description of adversaries—for example, Casper (Lowe, 1998), AVISPA (Viganò, 2005), ProVerif (Abadi and Blanchet, 2005)—but they are still bound to the restrictions of the models underlying their respective analysis methods, generally expressed in terms of an equational theory. The problem is even worse when it is not clear exactly what assumptions are implicitly being made about the adversary. One obvious assumption that needs to be made clear is whether the adversary is an insider in the system or an outsider. Lowe’s (1995) well-known man-in-the-middle attack against the Needham-Schroeder (1978) protocol highlights this issue. Until then, the Needham-Schroeder protocol had been analyzed under the assumption that the adversary had complete control of the network, and could intercept and inject arbitrary messages (up to the Dolev-Yao capabilities) into the protocol runs. However, the adversary was always assumed to be an outsider, not being able to directly interact with the protocol principals as himself. Lowe showed that if the adversary is allowed to be an *insider* of the system, that is, appear to the other principals as a bona fide protocol participant, then the Needham-Schroeder protocol does not guarantee the authentication properties it is meant to guarantee.

In this paper, we introduce a logic for reasoning about security protocols that allows us to model adversaries explicitly and naturally. The idea is to model the adversary in terms of what the adversary knows. This approach has some significant advantages. Logics of knowledge (Fagin et al., 1995) have been shown to provide powerful methods for reasoning about trace-based executions of protocols. They can be given semantics that is tied directly to protocol execution, thus avoiding problems of having to analyze an idealized form of the

protocol, as is required, for example, in BAN logic (Burrows et al., 1990). A straightforward application of logics of knowledge allows us to conclude that in the Duck-Duck-Goose protocol, the adversary knows the key. Logics of knowledge can also be extended with probabilities (Fagin and Halpern, 1994; Halpern and Tuttle, 1993) so as to be able to deal with probabilistic phenomena. Unfortunately, traditional logics of knowledge suffer from a well-known problem known as the *logical omniscience* problem: an agent knows all tautologies and all the logical consequences of her knowledge. The reasoning that allows an agent to infer properties of the protocol also allows an attacker to deduce properties that cannot be computed by realistic attackers in any reasonable amount of time.

To avoid the logical omniscience problem, we use the notion of *algorithmic knowledge* (Fagin et al., 1995, Chapter 10 and 11). Roughly speaking, we assume that agents (including adversaries) have “knowledge algorithms” that they use to compute what they know. The capabilities of the adversary are captured by its algorithm. Hence, Dolev-Yao capabilities can be provided by using a knowledge algorithm that can only compose messages or attempt to decipher using known keys. By changing the algorithm, we can extend the capabilities of the adversary so that it can attempt to crack the encryption scheme by factoring (in the case of RSA), using differential cryptanalysis (in the case of DES), or just by guessing keys, along the lines of a model due to Lowe (2002). Moreover, our framework can also handle the case of a principal sending the bits of its key, by providing the adversary’s algorithm with a way to check whether this is indeed what is happening. By explicitly using algorithms, we can therefore analyze the effect of bounding the resources of the adversary, and thus make progress toward bridging the gap between the analysis of cryptographic protocols and more computational accounts of cryptography. (See (Abadi and Rogaway, 2002) and the references therein for a discussion on work bridging this gap.) Note that we need both traditional knowledge and algorithmic knowledge in our analysis. Traditional knowledge is used to model a principal’s beliefs about what can happen in the protocol; algorithmic knowledge is used to model the adversary’s computational limitations (for example, the fact that it cannot factor).

The focus of this work is on developing a general and expressive framework for modeling and reasoning about security protocols, in which a wide class of adversaries can be represented naturally. In particular, we hope that the logic can provide a foundation for comparison and evaluation of formal methods that use different representation of adversaries. Therefore, we emphasize the expressiveness and representability aspects of the framework, rather than studying the kind of security properties that are useful in such a setting or developing techniques for proving that properties hold in the framework. These are all relevant questions that need to be pursued once the framework proves useful as a specification language.

The rest of the paper is organized as follows. In Section 2, we define our model for protocol analysis based on the well-understood multiagent system framework, and in Section 3 we present a logic for reasoning about implicit and explicit knowledge. In Section 4, we show how to model different adversaries from the literature. In Section 4.1, these adversaries are passive, in that they eavesdrop on the communication but do not attempt to interact with the principals of the system; in Section 4.2, we show how the framework can accommodate adversaries that actively interact with the principals by intercepting, forwarding, and replacing messages. We discuss related work in Section 5.

2. MODELING SECURITY PROTOCOLS

In this section, we review the multiagent system framework of Fagin et al. (1995, Chapters 4 and 5), and show it can be tailored to represent security protocols.

2.1. Multiagent Systems. The multiagent systems framework provides a model for knowledge that has the advantage of also providing a discipline for modeling executions of protocols. A multiagent system consists of n agents, each of which is in some local state at a given point in time. We assume that an agent’s local state encapsulates all the information to which the agent has access. In the security setting, the local state of an agent might include some initial information regarding keys, the messages she has sent and received, and perhaps the reading of a clock. In a poker game, a player’s local state might consist of the cards he currently holds, the bets made by other players, any other cards he has seen, and any information he may have about the strategies of the other players (for example, Bob may know that Alice likes to bluff, while Charlie tends to bet conservatively). The basic framework makes no assumptions about the precise nature of the local state.

We can then view the whole system as being in some global state, which is a tuple consisting of each agent’s local state, together with the state of the environment, where the environment consists of everything that is relevant to the system that is not contained in the state of the agents. Thus, a global state has the form (s_e, s_1, \dots, s_n) , where s_e is the state of the environment and s_i is agent i ’s state, for $i = 1, \dots, n$. The actual form of the agents’ local states and the environment’s state depends on the application.

A system is not a static entity. To capture its dynamic aspects, we define a run to be a function from time to global states. Intuitively, a run is a complete description of what happens over time in one possible execution of the system. A *point* is a pair (r, m) consisting of a run r and a time m . For simplicity, we take time to range over the natural numbers in the remainder of this discussion. At a point (r, m) , the system is in some global state $r(m)$. If $r(m) = (s_e, s_1, \dots, s_n)$, then we take $r_i(m)$ to be s_i , agent i ’s local state at the point (r, m) . We formally define a system \mathcal{R} to consist of a set of runs (or executions). It is relatively straightforward to model security protocols as systems. Note that the adversary in a security protocol can be modeled as just another agent. The adversary’s information at a point in a run can be modeled by his local state.

2.2. Specializing to Security. The multiagent systems framework is quite general. We have a particular application in mind, namely reasoning about security protocols, especially authentication protocols. We now specialize the framework in a way appropriate for reasoning about security protocols.

Since the vast majority of security protocols studied in the literature are message-based, a natural class of multiagent systems to consider is that of *message passing systems* (Fagin et al., 1995). Let \mathcal{M} be a fixed set of messages. A *history* for agent i (over \mathcal{M}) is a sequence of elements of the form $\text{send}(j, m)$ and $\text{recv}(m)$, where $m \in \mathcal{M}$. We think of $\text{send}(j, m)$ as representing the event “message m is sent to j ” and $\text{recv}(m)$ as representing the event “message m is received.” (We also allow events corresponding to internal actions; for the purpose of this paper, the internal actions we care about concern adversaries eavesdropping or intercepting messages—we return to those in Section 4.) Intuitively, i ’s history at (r, m) consists of i ’s initial state, which we take to be the empty sequence, followed by the sequence describing i ’s actions up to time m . If i performs no actions in round m , then

her history at (r, m) is the same as her history at $(r, m - 1)$. In such a message-passing system, we speak of $\text{send}(j, m)$ and $\text{recv}(m)$ as *events*. For an agent i , let $r_i(m)$ be agent i 's history in (r, m) . We say that an event e occurs in i 's history in round $m + 1$ of run r if e is in (the sequence) $r_i(m + 1)$ but not in $r_i(m)$.

In a message-passing system, the agent's local state at any point is her history. Of course, if h is the history of agent i at the point (r, m) , then we want it to be the case that h describes what happened in r up to time m from i 's point of view. To do this, we need to impose some consistency conditions on global states. In particular, we want to ensure that message histories do not shrink over time, and that every message received in round m corresponds to a message that was sent at some earlier round.

Given a set \mathcal{M} of messages, we define a *message-passing system* (over \mathcal{M}) to be a system satisfying the following constraints at all points (r, m) for each agent i :

MP1. $r_i(m)$ is a history over \mathcal{M} .

MP2. For every event $\text{recv}(m)$ in $r_i(m)$ there exists a corresponding event $\text{send}(i, m)$ in $r_j(m)$.

MP3. $r_i(0)$ is the empty sequence and $r_i(m + 1)$ is either identical to $r_i(m)$ or the result of appending one event to $r_i(m)$.

MP1 says that an agent's local state is her history, MP2 guarantees that every message received at round m corresponds to one that was sent earlier, and MP3 guarantees that histories do not shrink. We note that there is no guarantee that messages are not delivered twice, or that they are delivered at all.

A *security system* is a message passing system where the message space has a structure suitable for the interpretation of security protocols. Therefore, a security system assumes a set \mathcal{P} of plaintexts, as well as a set \mathcal{K} of keys. An encryption scheme \mathcal{C} over \mathcal{P} and \mathcal{K} is the closure \mathcal{M} of \mathcal{P} and \mathcal{K} under a key inverse operation $\text{inv} : \mathcal{K} \rightarrow \mathcal{K}$, a concatenation operation $\text{conc} : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$, decomposition operators $\text{first} : \mathcal{M} \rightarrow \mathcal{M}$ and $\text{second} : \mathcal{M} \rightarrow \mathcal{M}$, an encryption operation $\text{encr} : \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{M}$, and a decryption operation $\text{decr} : \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{M}$, subject to the constraints:

$$\begin{aligned} \text{first}(\text{conc}(m_1, m_2)) &= m_1 \\ \text{second}(\text{conc}(m_1, m_2)) &= m_2 \\ \text{decr}(\text{encr}(m, k), \text{inv}(k)) &= m. \end{aligned}$$

In other words, decrypting an encrypted message with the inverse of the key used to encrypt the message yields the original message. (For simplicity, we restrict ourselves to nonprobabilistic encryption schemes in this paper.) We often write k^{-1} for $\text{inv}(k)$, $m_1 \cdot m_2$ for $\text{conc}(m_1, m_2)$, and $\{m\}_k$ for $\text{encr}(m, k)$. There is no difficulty in adding more operations to the encryption schemes, for instance, to model hashes, signatures, the ability to take the exclusive-or of two terms, or the ability to compose two keys together to create a new key. We make no assumption in the general case as to the properties of encryption. Thus, for instance, most concrete encryption schemes allow collisions, that is, $\{m_1\}_{k_1} = \{m_2\}_{k_2}$ without $m_1 = m_2$ and $k_1 = k_2$. (In contrast, most security protocol analyses assume that there are no properties of encryption schemes beyond those specified above; this is part of the Dolev-Yao adversary model, which we examine in more detail in Section 4.1.1.)

Define \preceq on \mathcal{M} as the smallest relation satisfying the following constraints:

- (1) $m \preceq m$
- (2) if $m \preceq m_1$, then $m \preceq m_1 \cdot m_2$

- (3) if $m \preceq m_2$, then $m \preceq m_1 \cdot m_2$
(4) if $m \preceq m_1$, then $m \preceq \{\!\{m_1\}\!\}_k$.

Intuitively, $m_1 \preceq m_2$ if m_1 *could* be used in the construction of m_2 , or, equivalently, if m_1 can be considered part of message m_2 , under any of the possible ways of decomposing m_2 . For example, if $m = \{\!\{m_1\}\!\}_k = \{\!\{m_2\}\!\}_k$, then both $m_1 \preceq m$ and $m_2 \preceq m$. If we want to establish that $m_1 \preceq m_2$ for a given m_1 and m_2 , then we have to recursively look at all the possible ways in which m_2 can be constructed to see if m_1 can possibly be used to construct m_2 . Clearly, if encryption does not result in collisions, there is a single way in which m_2 can be taken apart. If the cryptosystem under consideration supports other operations (e.g., an exclusive-or operation, or operations that to combine keys to create new keys), additional constraints must be added to the \preceq relation to account for those operations.

To analyze a particular security protocol, we first derive the multiagent system corresponding to the protocol, using the approach of Fagin et al. (1995, Chapter 5). Intuitively, this multiagent system contains a run for every possible execution of the protocol, for instance, for every possible key used by the principals, subject to the restrictions above (such as MP1–MP3).

Formally, a protocol for agent i is a function from her local state to the set of actions that she can perform at that state. For ease of exposition, the only actions we consider here are those of sending messages (although we could easily incorporate other actions, such as choosing keys, or tossing coins to randomize protocols). A *joint protocol* $P = (P_e, P_1, \dots, P_n)$, consisting of a protocol for each of the agents (including a protocol for the environment), associates with each global state a set of possible *joint actions* (i.e., tuples of actions) in the obvious way. Joint actions transform global states. To capture their effect, we associate with every joint action \mathbf{a} a function $\tau(\mathbf{a})$ from global states to global states. This function captures, for instance, the fact that a message sent by an agent will be received by another agent, and so on. Given a context consisting of a set of initial global states, an interpretation τ for the joint actions, and a protocol P_e for the environment, we can generate a system corresponding to the joint protocol P in a straightforward way. Intuitively, the system consists of all the runs r that could have been generated by the joint protocol P , that is, for all m , $r(m+1)$ is the result of applying $\tau(\mathbf{a})$ to $r(m)$, where \mathbf{a} is a joint action that could have been performed according to the joint protocol P to $r(m)$.²

This way of generating systems from protocols is quite general. For instance, multiple protocol sessions can be modeled by using an action specifically for starting a new session of the protocol, in combination with protocol action being tagged by the session to which the action applies. More generally, most existing approaches to describing systems can be modeled using these kinds of protocols—it is a simple matter, for example, to take a process calculus expression and derive a protocol in our framework that generates the same system as the original process calculus expression. (Of course, process calculi come equipped with reasoning techniques that cannot be easily captured in our framework, but we are concerned with the ability to represent systems here.)

² It is also possible to represent a protocol in other ways, such as in terms of *strand spaces* (Thayer et al., 1999). Whichever representation is used, it should be possible to get a system corresponding to the protocol. For example, Halpern and Pucella (2003) show how to get a system from a strand space representation. For the purposes of this paper, the precise mechanism used to derive the multiagent system is not central, although it is an important issue for the development of formal tools for analyzing protocols.

3. A LOGIC FOR SECURITY PROPERTIES

The aim is to be able to reason about properties of security systems as defined in the last section, including properties involving the knowledge of agents in the system. To formalize this type of reasoning, we first need a language. The logic of algorithmic knowledge (Fagin et al., 1995, Chapters 10 and 11) provides such a framework. It extends the classical logic of knowledge by adding algorithmic knowledge operators.

The syntax of the logic $\mathcal{L}_n^{\text{KX}}$ for algorithmic knowledge is straightforward. Starting with a set Φ_0 of primitive propositions, which we can think of as describing basic facts about the system, such as “the key is k ” or “agent A sent the message m to B ”, formulas of $\mathcal{L}_n^{\text{KX}}(\Phi_0)$ are formed by closing off under negation, conjunction, and the modal operators K_1, \dots, K_n and X_1, \dots, X_n .

The formula $K_i\varphi$ is read as “agent i (implicitly) knows the fact φ ”, while $X_i\varphi$ is read as “agent i explicitly knows fact φ ”. In fact, we will read $X_i\varphi$ as “agent i can compute fact φ ”. This reading will be made precise when we discuss the semantics of the logic. As usual, we take $\varphi \vee \psi$ to be an abbreviation for $\neg(\neg\varphi \wedge \neg\psi)$ and $\varphi \Rightarrow \psi$ to be an abbreviation for $\neg\varphi \vee \psi$.

The standard models for this logic are based on the idea of possible worlds and Kripke structures (Kripke, 1963). Formally, a Kripke structure M is a tuple $(S, \pi, \mathcal{K}_1, \dots, \mathcal{K}_n)$, where S is a set of states or possible worlds, π is an interpretation which associates with each state in S a truth assignment to the primitive propositions (i.e., $\pi(s)(p) \in \{\mathbf{true}, \mathbf{false}\}$ for each state $s \in S$ and each primitive proposition p), and \mathcal{K}_i is an equivalence relation on S (recall that an equivalence relation is a binary relation which is reflexive, symmetric, and transitive). \mathcal{K}_i is agent i 's possibility relation. Intuitively, $(s, t) \in \mathcal{K}_i$ if agent i cannot distinguish state s from state t (so that if s is the actual state of the world, agent i would consider t a possible state of the world).

A system can be viewed as a Kripke structure, once we add a function π telling us how to assign truth values to the primitive propositions. An *interpreted system* \mathcal{I} consists of a pair (\mathcal{R}, π) , where \mathcal{R} is a system and π is an interpretation for the propositions in Φ that assigns truth values to the primitive propositions at the global states. Thus, for every $p \in \Phi$ and global state s that arises in \mathcal{R} , we have $\pi(s)(p) \in \{\mathbf{true}, \mathbf{false}\}$. Of course, π also induces an interpretation over the points of \mathcal{R} ; simply take $\pi(r, m)$ to be $\pi(r(m))$. We refer to the points of the system \mathcal{R} as points of the interpreted system \mathcal{I} .

The interpreted system $\mathcal{I} = (\mathcal{R}, \pi)$ can be made into a Kripke structure by taking the possible worlds to be the points of \mathcal{R} , and by defining \mathcal{K}_i so that $((r, m), (r', m')) \in \mathcal{K}_i$ if $r_i(m) = r'_i(m')$. Clearly \mathcal{K}_i is an equivalence relation on points. Intuitively, agent i considers a point (r', m') possible at a point (r, m) if i has the same local state at both points. Thus, the agents' knowledge is completely determined by their local states.

To account for X_i , we provide each agent with a knowledge algorithm that he uses to compute his knowledge. We will refer to $X_i\varphi$ as *algorithmic knowledge*. An *interpreted algorithmic knowledge system* has the form $(\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$, where (\mathcal{R}, π) is an interpreted system and \mathbf{A}_i is the knowledge algorithm of agent i . In local state ℓ , the agent computes whether he knows φ by applying the knowledge algorithm \mathbf{A} to input (φ, ℓ) . The output is either “Yes”, in which case the agent knows φ to be true, “No”, in which case the agent does *not* know φ to be true, or “?”, which intuitively says that the algorithm has insufficient resources to compute the answer. It is the last clause that allows us to deal with resource-bounded reasoners.

We define what it means for a formula φ to be true (or satisfied) at a point (r, m) in an interpreted system \mathcal{I} , written $(\mathcal{I}, r, m) \models \varphi$, inductively as follows:

$$\begin{aligned} (\mathcal{I}, r, m) &\models p \text{ if } \pi(r, m)(p) = \mathbf{true} \\ (\mathcal{I}, r, m) &\models \neg\varphi \text{ if } (\mathcal{I}, r, m) \not\models \varphi \\ (\mathcal{I}, r, m) &\models \varphi \wedge \psi \text{ if } (\mathcal{I}, r, m) \models \varphi \text{ and } (\mathcal{I}, r, m) \models \psi \\ (\mathcal{I}, r, m) &\models K_i\varphi \text{ if } (\mathcal{I}, r', m') \models \varphi \text{ for all } (r', m') \text{ such that } r_i(m) = r'_i(m') \\ (\mathcal{I}, r, m) &\models X_i\varphi \text{ if } \mathbf{A}_i(\varphi, r_i(m)) = \text{“Yes”}. \end{aligned}$$

The first clause shows how we use the π to define the semantics of the primitive propositions. The next two clauses, which define the semantics of \neg and \wedge , are the standard clauses from propositional logic. The fourth clause is designed to capture the intuition that agent i knows φ exactly if φ is true in all the worlds that i thinks are possible. The last clause captures the fact that explicit knowledge is determined using the knowledge algorithm of the agent.

As we pointed out, we think of K_i as representing *implicit knowledge*, facts that the agent implicitly knows, given its information, while X_i represents *explicit knowledge*, facts whose truth the agent can compute explicitly. As is well known, implicit knowledge suffers from the logical omniscience problem; agents implicitly know all valid formulas and agents implicitly know all the logical consequences of their knowledge (that is, $(K_\varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i\psi$ is valid). Explicit knowledge does not have that problem. Note that, as defined, there is no necessary connection between $X_i\varphi$ and $K_i\varphi$. An algorithm could very well claim that agent i knows φ (i.e., output “Yes”) whenever it chooses to, including at points where $K_i\varphi$ does not hold. Although algorithms that make mistakes are common, we are often interested in knowledge algorithms that are correct. A knowledge algorithm is *sound* for agent i in the system \mathcal{I} if for all points (r, m) of \mathcal{I} and formulas φ , $\mathbf{A}(\varphi, r_i(m)) = \text{“Yes”}$ implies $(\mathcal{I}, r, m) \models K_i\varphi$, and $\mathbf{A}(\varphi, r_i(m)) = \text{“No”}$ implies $(\mathcal{I}, r, m) \models \neg K_i\varphi$. Thus, a knowledge algorithm is sound if its answers are always correct.³

To reason about security protocols, we use the following set Φ_0^S of primitive propositions:

- $\text{send}_i(j, m)$: agent i sent message m to agent j ;
- $\text{recv}_i(m)$: agent i received message m ;
- $\text{has}_i(m)$: agent i has message m .

Intuitively, $\text{send}_i(j, m)$ is true when agent i has sent message m at some point, intended for agent j , and $\text{recv}_i(m)$ is true when agent i has received message m at some point. Agent i *has* a submessage m_1 at a point (r, m) , written $\text{has}_i(m_1)$, if there exists a message $m_2 \in \mathcal{M}$ such that $\text{recv}(m_2)$ is in $r_i(m)$, the local state of agent i , and $m_1 \preceq m_2$. Note that the has_i predicate is not constrained by encryption. If $\text{has}_i(\{\!|m|\!\}_k)$ holds, then so does $\text{has}_i(m)$, whether or not agent i knows the decryption key k^{-1} . Intuitively, the has_i predicate is true of messages that the agent considers (possible) part of the messages she has received, as captured by the \preceq relation.

An *interpreted algorithmic knowledge security system* is simply an interpreted algorithmic knowledge system $\mathcal{I} = (\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$, where \mathcal{R} is a security system, the set Φ_0 of primitive propositions includes Φ_0^S , and π is an *acceptable* interpretation, that is, it gives the following fixed interpretation to the primitive propositions in Φ_0^S :

- $\pi(r, m)(\text{send}_i(j, m)) = \mathbf{true}$ if and only if $\text{send}(j, m) \in r_i(m)$
- $\pi(r, m)(\text{recv}_i(m)) = \mathbf{true}$ if and only if $\text{recv}(m) \in r_i(m)$

³Note that for the purpose of this paper, there is no need for us to distinguish between “No” and “?” — none of our algorithms ever return “No”. We do keep the distinction, however, as well as the definition of soundness in cases the algorithm returns “No”, for consistency with existing uses of algorithmic knowledge.

- $\pi(r, m)(\text{has}_i(m)) = \mathbf{true}$ if and only if there exists m' such that $m \preceq m'$ and $\text{recv}(m') \in r_i(m)$.

This language can easily express the type of confidentiality (or secrecy) properties that we focus on here. Intuitively, we want to guarantee that throughout a protocol interaction, the adversary does not know a particular message. Confidentiality properties are stated naturally in terms of knowledge, for example, “agent 1 knows that the key k is a key known only to agent 2 and herself”. Confidentiality properties are well studied, and central to most of the approaches to reasoning about security protocols.⁴ Higher-level security properties, such as authentication properties, can often be established via confidentiality properties. See (Syverson and Cervesato, 2001) for more details.

To illustrate some of the issues involved, consider an authentication protocol such as the Needham-Schroeder-Lowe protocol (Lowe, 1995). A simplified version of the protocol is characterized by the following message exchange between two agents A and B :

$$\begin{aligned} A &\rightarrow B : \{n_A, A\}_{k_B} \\ B &\rightarrow A : \{n_A, n_B, B\}_{k_A} \\ A &\rightarrow B : \{n_B\}_{k_B}. \end{aligned}$$

An authentication property of this protocol can be expressed informally as follows: under suitable assumptions on the keys known to the adversary and the fact that B is running his part of the protocol, A knows that n_A and n_B are kept confidential between her and B .⁵ From this, she knows that she is interacting with B , because she has received a message containing n_A , which only B could have produced. Similarly, A also knows that when B receives her message, B will know that he is interacting with A , because only A knows the nonce n_B which is part of the last message. Similar reasoning can be applied to B . This argument relies on the confidentiality of the nonces n_a and n_b . It is tempting to capture this fact by stating that no agent i other than A and B knows $\text{has}_i(n_A)$ or $\text{has}_i(n_B)$. (We must write this as $K_i\text{has}_i(n_A)$ or $K_i\text{has}_i(n_B)$ because we cannot express directly the idea of knowing a message—an agent can only know if they have received a message, or if a message is a possible component of a message they have received.)

Unfortunately, because the implicit knowledge operator suffers from logical omniscience, such a statement does not capture the intent of confidentiality. At every point where an adversary a intercepts a message $\{n_A, n_B, B\}_{k_A}$, $K_a\text{has}_a(n_A)$ is true (since $n_A \preceq \{n_A, n_B, B\}_{k_A}$), and hence the adversary knows that he has seen the nonce n_A , irrespective of whether he knows the decryption key corresponding to k_A). This shows that the standard interpretation of knowledge expressed via the implicit knowledge operator does not capture important aspects of reasoning about security. The adversary having the implicit knowledge that n_A is part of the message does not suffice, in general, for the adversary to *explicitly* know that n_A is part of the message. Intuitively, the adversary may not have the capabilities to realize he has seen n_A .

A more reasonable interpretation of confidentiality in this setting is $\neg X_a\text{has}_a(n_A)$: the adversary does not explicitly know (i.e., cannot compute or derive) whether he has seen the nonce n_A . Most logics of security, instead of relying on a notion of knowledge, introduce

⁴A general definition of secrecy in terms of knowledge is presented by Halpern and O’Neill (2002) in the context of information flow, a setting that does not take into account cryptography.

⁵It may be more reasonable to talk about *belief* rather than *knowledge* that n_A and n_B are kept confidential. For simplicity, we talk about knowledge in this paper. Since most representations of belief suffer from logical omniscience, what we say applies to belief as well as knowledge.

special primitives to capture the fact that the adversary can see a message m encrypted with key k only if he has access to the key k . Doing this hardwires the capabilities of the adversary into the semantics of the logic. Changing these capabilities requires changing the semantics. In our case, we simply need to supply the appropriate knowledge algorithm to the adversary, capturing his capabilities. In the following section, we examine in more detail the kind of knowledge algorithms that correspond to interesting capabilities.

4. MODELING ADVERSARIES

As we showed in the last two sections, interpreted algorithmic knowledge security systems can be used to provide a foundation for representing security protocols, and support a logic for writing properties based on knowledge, both traditional (implicit) and algorithmic (explicit). For the purposes of analyzing security protocols, we use traditional knowledge to model a principal’s beliefs about what can happen in the protocol, while we use algorithmic knowledge to model the adversary’s capabilities, possibly resource-bounded. To interpret algorithmic knowledge, we rely on a knowledge algorithm for each agent in the system. We use the adversary’s knowledge algorithm to capture the adversary’s ability to draw conclusions from what he has seen. In this section, we show how we can capture different capabilities for the adversary in a natural way in this framework. We first show how to capture the standard model of adversary due to Dolev and Yao. We then show how to account for the adversary in the Duck-Duck-Goose protocol, and the adversary considered by Lowe (2002) that can perform self-validating guesses.

We start by considering passive (or eavesdropping) adversaries, which simply record every message exchanged by the principals; in Section 4.2, we consider active adversaries. For simplicity, we assume a single adversary per system; our results extend to the general case immediately, but the notation becomes cumbersome.

Adversaries generally have the ability to eavesdrop on all communications, and in the case of active adversaries, to furthermore intercept messages and forward them at will. To model both the eavesdropping of a message and its interception, we assume an internal action $\text{intercept}(m)$ meant to capture the fact that an adversary has intercepted message m . Because messages may or may not be delivered to their final destination, this models both eavesdropping—in which case the message does in fact get received by its intended recipient—and actual interception—in which case the message does not get received by its intended recipient. Because an intercepted message is available to an agent in its local state, we adapt the notion of an acceptable interpretation from Section 3 to allow intercept messages to be used for determining which messages an agent has:

- $\pi(r, m)(\text{has}_i(m)) = \mathbf{true}$ if and only if there exists m' such that $m \preceq m'$ and either $\text{rcv}(m') \in r_i(m)$ or $\text{intercept}(m') \in r_i(m)$.

4.1. Passive Adversaries. Passive adversaries can be modeled formally as follows. An *interpreted algorithmic knowledge security system with passive adversary a* ($a \in \{1, \dots, n\}$) is an interpreted algorithmic knowledge security system $\mathcal{I} = (\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$ satisfying the following constraints at all points (r, m) :

- P1. $r_a(m)$ consists only of $\text{intercept}(m)$ events.
- P2. For all j and events $\text{send}(j, m)$ in $r_j(m)$, there exists an event $\text{intercept}(m)$ in $r_a(m)$.

P3. For every $\text{intercept}(\mathbf{m})$ in $r_a(m)$, there is a corresponding $\text{send}(j, \mathbf{m})$ in $r_i(m)$ for some i .

P4. For all i and every $\text{send}(j, \mathbf{m})$ in $r_i(m)$, $j \neq a$.

P1 captures the passivity of the adversary—he can only intercept messages, not send any; P2 says that every message sent by a principal is basically copied to the adversary’s local state, while P3 says that only messages sent by principals appear in the adversary’s local state; P4 ensures that a passive adversary is not an agent to which other agents may intentionally send messages to (i.e., the adversary is an outsider in the system). We next consider various knowledge algorithms for the adversary.

4.1.1. *The Dolev-Yao Adversary.* Consider the standard Dolev-Yao adversary (Dolev and Yao, 1983). This model is a combination of assumptions on the encryption scheme used and the capabilities of the adversaries. Specifically, the encryption scheme is seen as the free algebra generated by \mathcal{P} and \mathcal{K} over operations \cdot and $\{\!\!\}\}_k$. Perhaps the easiest way to formalize this is to view the set \mathcal{M} as the set of expressions generated by the grammar

$$\mathbf{m} ::= \mathbf{p} \mid \mathbf{k} \mid \{\!\!\}\mathbf{m}\}_k \mid \mathbf{m} \cdot \mathbf{m}$$

(with $\mathbf{p} \in \mathcal{P}$ and $\mathbf{k} \in \mathcal{K}$). We assume that there are no collisions; messages always have a unique decomposition. The only way that $\{\!\!\}\mathbf{m}\}_k = \{\!\!\}\mathbf{m}'\}_{k'}$ is if $\mathbf{m} = \mathbf{m}'$ and $k = k'$. We also make the standard assumption that concatenation and encryption have enough redundancy to recognize that a term is in fact a concatenation $\mathbf{m}_1 \cdot \mathbf{m}_2$ or an encryption $\{\!\!\}\mathbf{m}\}_k$.

The classical Dolev-Yao model can be formalized by a relation $H \vdash_{DY} \mathbf{m}$ between a set H of messages and a message \mathbf{m} . (Our formalization is equivalent to many other formalizations of Dolev-Yao in the literature, and is similar in spirit to that of Paulson (1998).) Intuitively, $H \vdash_{DY} \mathbf{m}$ means that an adversary can “extract” message \mathbf{m} from a set of received messages and keys H , using the allowable operations. The derivation is defined using the following inference rules:

$$\frac{\mathbf{m} \in H}{H \vdash_{DY} \mathbf{m}} \quad \frac{H \vdash_{DY} \{\!\!\}\mathbf{m}\}_k \quad H \vdash_{DY} k^{-1}}{H \vdash_{DY} \mathbf{m}} \quad \frac{H \vdash_{DY} \mathbf{m}_1 \cdot \mathbf{m}_2}{H \vdash_{DY} \mathbf{m}_1} \quad \frac{H \vdash_{DY} \mathbf{m}_1 \cdot \mathbf{m}_2}{H \vdash_{DY} \mathbf{m}_2}.$$

This presentation of the Dolev-Yao capabilities is restricted, because it only allows an adversary to deconstruct messages, and not to construct them. In many situations with passive adversaries—that is, adversaries that cannot inject new messages into the system—this is not a significant restriction. It is easy to extend the relation to allow for the construction of new messages using concatenation and encryption. (Of course, extending the relation in such a way would require a corresponding modification to the knowledge algorithm below.)

In our framework, to capture the capabilities of a Dolev-Yao adversary, we specify how the adversary can explicitly know that she *has* a message, by defining a knowledge algorithm \mathbf{A}_i^{DY} for adversary i . Recall that a knowledge algorithm for agent i takes as input a formula and agent i ’s local state (which we are assuming contains the messages received by i). The most interesting case in the definition of \mathbf{A}_i^{DY} is when the formula is $\text{has}_i(\mathbf{m})$. To compute $\mathbf{A}_i^{DY}(\text{has}_i(\mathbf{m}), \ell)$, the algorithm simply checks, for every message \mathbf{m}' received by the adversary, whether \mathbf{m} is a submessage of \mathbf{m}' , according to the keys that are known to the adversary. We assume that the adversary’s initial state consists of the set of keys initially known by the adversary. This will typically contain, in a public-key cryptography setting, the public keys of all the agents. We use $\text{initkeys}(\ell)$ to denote the set of initial keys known by agent i in local state ℓ . (Recall that a local state for agent i is the sequence of events pertaining to

```

partof( $m, m', K$ ) : if  $m = m'$  then
    return true
    if  $m'$  is  $\{\!\{m_1\}\!\}_k$  and  $k^{-1} \in K$  then
    return partof( $m, m_1, K$ )
    if  $m'$  is  $m_1 \cdot m_2$  then
    return partof( $m, m_1, K$ )  $\vee$  submsg( $m, m_2, K$ )
    return false

getkeys( $m, K$ ) : if  $m \in \mathcal{K}$  then
    return  $\{m\}$ 
    if  $m'$  is  $\{\!\{m_1\}\!\}_k$  and  $k^{-1} \in K$  then
    return getkeys( $m_1, K$ )
    if  $m'$  is  $m_1 \cdot m_2$  then
    return getkeys( $m_1, K$ )  $\cup$  getkeys( $m_2, K$ )
    return  $\{\}$ 

keysof( $\ell$ ) :  $K \leftarrow \textit{initkeys}(\ell)$ 
    loop until no change in  $K$ 
     $K \leftarrow \bigcup_{\textit{rcv}(m) \in \ell} \textit{getkeys}(m, K)$ 
    return  $K$ 

```

Figure 1: Dolev-Yao knowledge algorithm auxiliary functions

agent i , including any initial information in the run, in this case, the keys initially known.) The function *partof*, which can take apart messages created by concatenation, or decrypt messages as long as the adversary knows the decryption key, is used to check whether m is a submessage of m' . $\mathbf{A}_i^{\text{DY}}(\textit{has}_i(m), \ell)$ is defined as follows:

```

if  $m \in \textit{initkeys}(\ell)$  then return “Yes”
 $K \leftarrow \textit{keysof}(\ell)$ 
for each  $\textit{rcv}(m')$  and  $\textit{intercept}(m')$  in  $\ell$ 
    if partof( $m, m', K$ ) then
    return “Yes”
return “?”.

```

The auxiliary functions used by the algorithm are given in Figure 1. In particular, the function *partof* captures the \preceq relation.

In the Dolev-Yao model, an adversary cannot explicitly compute anything interesting about what other messages agents have. Hence, for other primitives, including $\textit{has}_j(m)$ for $j \neq i$, \mathbf{A}_i^{DY} returns “?”. For formulas of the form $K_j\varphi$ and $X_j\varphi$, \mathbf{A}_i^{DY} also returns “?”. For Boolean combinations of formulas, \mathbf{A}_i^{DY} returns the corresponding Boolean combination (where the negation of “?” is “?”, the conjunction of “No” and “?” is “No”, and the conjunction of “Yes” and “?” is “?”) of the answer for each $\textit{has}_i(m)$ query.

The following result shows that an adversary using \mathbf{A}_i^{DY} recognizes (i.e., returns “Yes” to) $\textit{has}_i(m)$ in state ℓ if and only if m is one of the messages that can be derived (according

to \vdash_{DY}) from the messages received in that state together with the keys initially known. Moreover, if a $\text{has}_i(m)$ formula is derived at the point (r, m) , then $\text{has}_i(m)$ is actually true at (r, m) (so that \mathbf{A}_i^{DY} is sound).

Proposition 4.1. *Let $\mathcal{I} = (\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$ be an interpreted algorithmic knowledge security system where $\mathbf{A}_i = \mathbf{A}_i^{\text{DY}}$. Then*

$$(\mathcal{I}, r, m) \models X_i(\text{has}_i(m)) \text{ if and only if } \{\mathbf{m}' : \text{recv}(\mathbf{m}') \in r_i(m)\} \cup \{\mathbf{m}' : \text{intercept}(\mathbf{m}') \in r_i(m)\} \cup \text{initkeys}(\ell) \vdash_{DY} \mathbf{m}.$$

Moreover, if $(\mathcal{I}, r, m) \models X_i(\text{has}_i(m))$ then $(\mathcal{I}, r, m) \models \text{has}_i(m)$.

Proof. Let $K = \text{keysof}(r_i(m))$. We must show that $\mathbf{A}_i^{\text{DY}}(\text{has}_i(m), r_i(m)) = \text{“Yes”}$ if and only if $K \cup \{\mathbf{m}' : \text{recv}(\mathbf{m}') \in r_i(m)\} \cup \{\mathbf{m}' : \text{intercept}(\mathbf{m}') \in r_i(m)\} \vdash_{DY} \mathbf{m}$. It is immediate from the description of \mathbf{A}_i^{DY} and \vdash_{DY} that this is true if $\mathbf{m} \in \text{initkeys}(r_i(m))$. If $\mathbf{m} \notin \text{initkeys}(r_i(m))$, then $\mathbf{A}_i^{\text{DY}}(\text{has}_i(m), r_i(m)) = \text{“Yes”}$ if and only if $\text{partof}(\mathbf{m}, \mathbf{m}', K) = \text{true}$ for some \mathbf{m}' such that $\text{recv}(\mathbf{m}') \in r_i(m)$ or $\text{intercept}(\mathbf{m}') \in r_i(m)$. Next observe that $\text{partof}(\mathbf{m}, \mathbf{m}', K) = \text{true}$ if and only if $K \cup \{\mathbf{m}'\} \vdash_{DY} \mathbf{m}$: the “if” direction follows by a simple induction on the length of the derivation; the “only if” direction follows by a straightforward induction on the structure of \mathbf{m} . Finally, observe that if M is a set of messages, then $K \cup M \vdash_{DY} \mathbf{m}$ if and only if $K \cup \{\mathbf{m}'\} \vdash_{DY} \mathbf{m}$ for some $\mathbf{m}' \in M$. The “if” direction is trivial. The “only if” direction follows by induction on the number of times the rule “from $\mathbf{m}' \in H$ infer $H \vdash_{DY} \mathbf{m}''$ ” is used to derive some $\mathbf{m}' \in M$. If it is never used, then it is easy to see that $K \vdash_{DY} \mathbf{m}'$. If it is used more than once, and the last occurrence is used to derive \mathbf{m}' , then it is easy to see that $K \cup \{\mathbf{m}'\} \vdash_{DY} \mathbf{m}'$ (the derivation just starts from the last use of this rule). The desired result is now immediate. \square

In particular, if we have an interpreted algorithmic knowledge security system with a passive adversary a such that $\mathbf{A}_a = \mathbf{A}_a^{\text{DY}}$, then Proposition 4.1 captures the knowledge of a passive Dolev-Yao adversary.

4.1.2. The Duck-Duck-Goose Adversary. The key advantage of our framework is that we can easily change the capabilities of the adversary beyond those prescribed by the Dolev-Yao model. For example, we can capture the fact that if the adversary knows the protocol, she can derive more information than she could otherwise. For instance, in the Duck-Duck-Goose example, assume that the adversary maintains in her local state a list of all the bits received corresponding to the key of the principal. We can write the algorithm so that if the adversary’s local state contains all the bits of the key of the principal, then the adversary can decode messages that have been encrypted with that key. Specifically, assume that key k is being sent in the Duck-Duck-Goose example. Then for an adversary i , $\text{has}_i(k)$ will be false until all the bits of the key have been received. This translates immediately into the following algorithm $\mathbf{A}_i^{\text{DDG}}$:

if all the bits recorded in ℓ form k then
return “Yes” else return “?”.

$\mathbf{A}_i^{\text{DDG}}$ handles other formulas in the same way as \mathbf{A}_i^{DY} .

Of course, nothing keeps us from combining algorithms, so that we can imagine an adversary intercepting both messages and key bits, and using an algorithm \mathbf{A}_i that is a

combination of the Dolev-Yao algorithm and the Duck-Duck-Goose algorithm; $A_i(\varphi, \ell)$ is defined as follows:

if $A_i^{\text{DY}}(\varphi, \ell) = \text{“Yes”}$ then
 return “Yes”
 else return $A_i^{\text{DDG}}(\varphi, \ell)$.

This assumes that the adversary knows the protocol, and hence knows when the key bits are being sent. The algorithm above captures this protocol-specific knowledge.

To see why this adversary is not completely trivial, note that the obvious way of trying to capture this kind of adversary in a Dolev-Yao model, where we allow operations for combining keys to form new keys, fails in that the resulting adversary is too powerful. If we assume that keys are formed out of bits, and we allow a bit-concatenation operation that lets the adversary create keys out of bits, then a Dolev-Yao adversary with such extended operations, as soon as they have received bits 1 and 0, would be able to construct any possible key, and therefore such an adversary would be able to decrypt any encrypted message.

4.1.3. *The Lowe Adversary.* For a more realistic example of an adversary model that goes beyond Dolev-Yao, consider the following adversary model introduced by Lowe (2002) to analyze protocols subject to guessing attacks. The intuition is that some protocols provide for a way to “validate” the guesses of an adversary. For a simple example of this, here is a simple challenge-based authentication protocol:

$A \rightarrow S : A$
 $S \rightarrow A : n_s$
 $A \rightarrow S : \{n_s\}_{p_a}$.

Intuitively, A tells the server S that she wants to authenticate herself. S replies with a challenge n_s . A sends back to S the challenge encrypted with her password p_a . Presumably, S knows the password, and can verify that she gets $\{n_s\}_{p_a}$. Unfortunately, an adversary can overhear both n_s and $\{n_s\}_{p_a}$, and can “guess” a value g for p_a and verify his guess by checking if $\{n_s\}_g = \{n_s\}_{p_a}$. The key feature of this kind of attack is that the guessing (and the validation) can be performed offline, based only on the intercepted messages.

To account for this capability of adversaries is actually fairly complicated. We present a slight variation of Lowe’s description, mostly to make it notationally consistent with the rest of the section; we refer the reader to Lowe (2002) for a discussion of the design choices.

Lowe’s model relies on a basic one-step reduction function, $S \triangleright_l m$, saying that the messages in S can be used to derive the message m . Its definition is reminiscent of \vdash_{DY} , except that it represents a single step of derivation. Note that the derivation relation \triangleright_l is “tagged” by the kind of derivation performed (l).

$$\begin{aligned} \{m, k\} &\triangleright_{\text{enc}} \{m\}_k \\ \{\{m\}_k, k^{-1}\} &\triangleright_{\text{dec}} m \\ \{m_1 \cdot m_2\} &\triangleright_{\text{fst}} m_1 \\ \{m_1 \cdot m_2\} &\triangleright_{\text{snd}} m_2. \end{aligned}$$

Lowe also includes a reduction to derive $m_1 \cdot m_2$ from m_1 and m_2 . We do not add this reduction to simplify the presentation. It is straightforward to extend our approach to deal with it.

Given a set H of messages, and a sequence t of one-step reductions, we define inductively the set $[H]_t$ of messages obtained from the one-step reductions given in t :

$$[H]_{\langle \rangle} = H$$

$$[H]_{\langle S \triangleright_l \mathbf{m} \rangle \cdot t} = \begin{cases} [H \cup \{\mathbf{m}\}]_t & \text{if } S \subseteq H \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Here, $\langle \rangle$ denotes the empty trace, and $t_1 \cdot t_2$ denotes trace concatenation. A trace t is said to be *monotone* if, intuitively, it does not perform any one-step reduction that “undoes” a previous one-step reduction. For example, the reduction $\{\mathbf{m}, \mathbf{k}\} \triangleright \{\mathbf{m}\}_{\mathbf{k}}$ undoes the reduction $\{\{\mathbf{m}\}_{\mathbf{k}}, \mathbf{k}^{-1}\} \triangleright \mathbf{m}$. (See Lowe (2002) for more details on undoing reductions.)

We say that a set H of messages *validates* a guess \mathbf{m} if H contains enough information to verify that \mathbf{m} is indeed a good guess. Intuitively, this happens if a value v (called a validator) can be derived from the messages in $H \cup \{\mathbf{m}\}$ in a way that uses the guess \mathbf{m} , and either that (a) validator v can be derived in a different way from $H \cup \{\mathbf{m}\}$, (b) the validator v is already in $H \cup \{\mathbf{m}\}$, or (c) the validator v is a key whose inverse is derivable from $H \cup \{\mathbf{m}\}$. For example, in the protocol exchange at the beginning of this section, the adversary sees the messages $H = \{n_s, \{\mathbf{n}_s\}_{p_a}\}$, and we can check that H validates the guess $\mathbf{m} = p_a$: clearly, $\{n_s, \mathbf{m}\} \triangleright_{\text{enc}} \{\mathbf{n}_s\}_{p_a}$, and $\{\mathbf{n}_s\}_{p_a} \in H \cup \{\mathbf{m}\}$. In this case, the validator $\{\mathbf{n}_s\}_{p_a}$ is already present in $H \cup \{\mathbf{m}\}$. For other examples of validation, we again refer to Lowe (2002).

We can now define the relation $H \vdash_L \mathbf{m}$ that says that \mathbf{m} can be derived from H by a Lowe adversary. Intuitively, $H \vdash_L \mathbf{m}$ if \mathbf{m} can be derived by Dolev-Yao reductions, or \mathbf{m} can be guessed and validated by the adversary, and hence susceptible to an attack. Formally, $H \vdash_L \mathbf{m}$ if and only if $H \vdash_{DY} \mathbf{m}$ or there exists a monotone trace t , a set S , and a “validator” v such that

- (1) $[H \cup \{\mathbf{m}\}]_t$ is defined;
- (2) $S \triangleright_l v$ is in t ;
- (3) there is no trace t' such that $S \subseteq [H]_{t'}$; and
- (4) either:
 - (a) there exists $(S', l') \neq (S, l)$ with $S' \triangleright_{l'} v$ in t
 - (b) $v \in H \cup \{\mathbf{m}\}$ or
 - (c) $v \in \mathcal{K}$ and $v^{-1} \in [H \cup \{\mathbf{m}\}]_t$.

It is not hard to verify that this formalization captures the intuition about validation given earlier. Specifically, condition (1) says that the trace t is well-formed, condition (2) says that the validator v is derived from $H \cup \{\mathbf{m}\}$, condition (3) says that deriving the validator v depends on the guess \mathbf{m} , and condition (4) specifies when a validator v validates a guess \mathbf{m} , as given earlier.

We would now like to define a knowledge algorithm \mathbf{A}_i^L to capture the capabilities of the Lowe adversary. Again, the only case of real interest is what \mathbf{A}_i^L does on input $\text{has}_i(\mathbf{m})$. $\mathbf{A}_i^L(\text{has}_i(\mathbf{m}), \ell)$ is defined as follows:

```

if  $\mathbf{A}_i^{\text{DY}}(\text{has}_i(\mathbf{m}), \ell) = \text{“Yes”}$  then
  return “Yes”
if  $\text{guess}(\mathbf{m}, \ell)$  then
  return “Yes”
return “?”.
```

```

guess( $m, \ell$ ) :  $H \leftarrow \text{reduce}(\{m' : \text{recv}(m') \in \ell \text{ or } \text{intercept}(m') \in \ell\} \cup \text{initkeys}(\ell)) \cup \{m\}$ 
                 $\text{reds} \leftarrow \{\}$ 
                loop until  $\text{reductions}(H) - \text{reds}$  is empty
                ( $S, l, v$ )  $\leftarrow$  pick an element of  $\text{reductions}(H) - \text{reds}$ 
                if  $\exists(S', l', v) \in \text{reds}$  s.t.  $S' \neq S$  and  $l' \neq l$  then
                return “Yes”
                if  $v \in H$  then
                return “Yes”
                if  $v \in \mathcal{K}$  and  $v^{-1} \in H$  then
                return “Yes”
                 $\text{reds} \leftarrow \text{reds} \cup \{(S, l, v)\}$ 
                 $H \leftarrow H \cup \{v\}$ 
                return “No”

reduce( $H$ ) : loop until no change in  $H$ 
                 $r \leftarrow \text{reductions}(H)$ 
                for each ( $S, l, v$ ) in  $r$ 
                 $H \leftarrow H \cup \{v\}$ 
                return  $H$ 

reductions( $H$ ) :  $\text{reds} \leftarrow \{\}$ 
                for each  $m_1 \cdot m_2$  in  $H$ 
                 $\text{reds} \leftarrow \{(\{m\}, \text{fst}, m_1), (\{m\}, \text{snd}, m_2)\}$ 
                for each  $m_1, m_2$  in  $H$ 
                if  $m_2 \in \mathcal{K}$  and  $\text{sub}(\{\!|m_1|\!\}_{m_2}, H)$  then
                 $\text{reds} \leftarrow \{(\{m_1, m_2\}, \text{enc}, \{\!|m_1|\!\}_{m_2})\}$ 
                if  $m_1$  is  $\{\!|m'\!\}_k$  and  $m_2$  is  $k^{-1}$  then
                 $\text{reds} \leftarrow \{(\{m_1, m_2\}, \text{dec}, m')\}$ 
                return  $\text{reds}$ 

sub( $m, H$ ) : if  $H = \{m\}$  then
                return true
                if  $H = \{m_1 \cdot m_2\}$  then
                return  $\text{sub}(m, \{m_1\}) \vee \text{sub}(m, \{m_2\})$ 
                if  $H = \{\!|m'\!\}_k$  then
                return  $\text{sub}(m, \{m'\})$ 
                if  $|H| > 1$  and  $H = \{m'\} \cup H'$  then
                return  $\text{sub}(m, \{m'\}) \vee \text{sub}(m, H')$ 
                return false

```

Figure 2: Lowe knowledge algorithm auxiliary functions

The auxiliary functions used by the algorithm are given in Figure 2. (We have not concerned ourselves with matters of efficiency in the description of \mathbf{A}_i^L ; again, see Lowe (2002) for a discussion of implementation issues.)

As before, we can check the correctness and soundness of the algorithm:

Proposition 4.2. *Let $\mathcal{I} = (\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$ be an interpreted algorithmic knowledge security system where $\mathbf{A}_i = \mathbf{A}_i^L$. Then*

$$(\mathcal{I}, r, m) \models X_i(\text{has}_i(m)) \text{ if and only if } \{\mathbf{m}' : \text{recv}(\mathbf{m}') \in r_i(m)\} \cup \{\mathbf{m}' : \text{intercept}(\mathbf{m}') \in r_i(m)\} \cup \text{initkeys}(\ell) \vdash_L m.$$

Moreover, if $(\mathcal{I}, r, m) \models X_i(\text{has}_i(m))$ then $(\mathcal{I}, r, m) \models \text{has}_i(m)$.

Proof. Let $K = \text{keysof}(r_i(m))$. The proof is similar in spirit to that of Proposition 4.1, using the fact that if $m \notin \text{initkeys}(r_i(m))$ and $K \cup \{\mathbf{m}' \mid \text{recv}(\mathbf{m}') \in r_i(m)\} \cup \{\mathbf{m}' \mid \text{intercept}(\mathbf{m}') \in r_i(m)\} \not\vdash_{DY} m$, then $\text{guess}(m, r_i(m)) = \text{“Yes”}$ if and only if $K \cup \{\mathbf{m}' \mid \text{recv}(\mathbf{m}') \in r_i(m)\} \cup \{\mathbf{m}' \mid \text{intercept}(\mathbf{m}') \in r_i(m)\} \vdash_L m$. The proof of this fact is essentially given by Lowe (2002), the algorithm \mathbf{A}_i^L being a direct translation of the CSP process implementing the Lowe adversary. Again, soundness with respect to $\text{has}_i(m)$ follows easily. \square

4.2. Active Adversaries. Up to now we have considered passive adversaries, which can intercept messages exchanged by protocol participants, but cannot actively participate in the protocol. Passive adversaries are often appropriate when the concern is confidentiality of messages. However, there are many attacks on security protocols that do not necessarily involve a breach of confidentiality. For instance, some authentication properties are concerned with ensuring that no adversary can pass himself off as another principal. This presumes that the adversary is able to interact with other principals. Even when it comes to confidentiality, there are clearly attacks that an active adversary can make that cannot be made by a passive adversary.

To analyze active adversaries, we need to consider what messages they can send. This, in turn depends on their capabilities, which we already have captured using knowledge algorithms. Formally, at a local state ℓ , an adversary using knowledge algorithm \mathbf{A}_i can construct the messages in the set $C(\ell)$, defined to be the closure under *conc* and *encr* of the set $\{m \mid \mathbf{A}_i(\text{has}_i(m), \ell) = \text{“Yes”}\}$ of messages that adversary i has. For more complex cryptosystems, the construction and destruction of messages to determine which can in fact be created by the adversary will be more complex—Paulson (1998), for instance, defines a general approach based on two operations, analysis and synthesis, interleaved to generate all possible messages that an adversary can construct. This technique can be readily adapted to our framework.

Once we consider active adversaries, we must consider whether they are insiders or outsiders. Intuitively, an insider is an adversary that other agents know about, and can initiate interactions with. (Insider adversaries are sometimes called *corrupt principals* or *dishonest principals*.) As we mentioned in the introduction, the difference between insiders and outsiders was highlighted by Lowe’s (1995) man-in-the-middle attack of the Needham-Schroeder protocol.

An interpreted algorithmic knowledge security system with active (insider) adversary a ($a \in \{1, \dots, n\}$) is an interpreted algorithmic knowledge security system $\mathcal{I} = (\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$ satisfying the following constraints at all points (r, m) .

- A1. For every $\text{intercept}(m) \in r_a(m)$, there is a corresponding $\text{send}(j, m)$ in $r_i(m)$ for some i, j .
- A2. For every $\text{send}(j, m) \in r_a(m)$, we have $m \in C(r_a(m))$.

A1 says that every message sent by the agents can be intercepted by the adversary and end up in the adversary’s local state, and every intercepted message in the adversary’s local state is a message that has been sent by an agent. A2 says that every message sent by the adversary must have been constructed out of the messages in his local state according to his capabilities. (Note that the adversary can forge the “send” field of the messages.)

To accommodate outsider adversaries, it suffices to add the restriction that no message is sent directly to the adversary. Formally, an *interpreted algorithmic knowledge security system with active (outsider) adversary* a ($a \in \{1, \dots, n\}$) is an interpreted algorithmic knowledge security system $\mathcal{I} = (\mathcal{R}, \pi, \mathbf{A}_1, \dots, \mathbf{A}_n)$ with an active insider adversary a such that for all points (r, m) and for all agents i , the following additional constraint is satisfied.

A3. For every $\text{send}(j, \mathbf{m}) \in r_i(m)$, $j \neq a$.

5. RELATED WORK

The issues we raise in this paper are certainly not new, and have been addressed, up to a point, in the literature. In this section, we review this literature, and discuss where we stand with respect to other approaches that have attempted to tackle some of the same problems.

As we mentioned in the introduction, the Dolev-Yao adversary is the most widespread adversary in the literature. Part of its attraction is its tractability, making it possible to develop formal systems to automatically check for safety with respect to such adversaries (Millen et al., 1987; Mitchell et al., 1997; Paulson, 1998; Lowe, 1998; Meadows, 1996). The idea of moving beyond the Dolev-Yao adversary is not new. As we pointed out in Section 4.1.3, Lowe (2002) developed an adversary that can encode some amount of off-line guessing; we showed in Section 4.1.3 that we could capture such an adversary in our framework. More recent techniques for detecting off-line guessing attacks (Corin et al., 2005; Baudet, 2005) can also be similarly modeled. Other approaches have the possibility of extending the adversary model. For instance, the framework of Paulson (1998), Clarke, Jha and Morrero (1998), and Lowe (1998) describe the adversary via a set of derivation rules, which could be modified by adding new derivation rules. We could certainly capture these adversaries by appropriately modifying our \mathbf{A}_i^{DY} knowledge algorithm. (Pucella (2006) studies the properties of algorithmic knowledge given by derivation rules in more depth.) However, these other approaches do not seem to have the flexibility of our approach in terms of capturing adversaries. Not all adversaries can be conveniently described in terms of derivation rules.

There are other approaches that weaken the Dolev-Yao adversary assumptions by either taking concrete encryption schemes into account, or at least adding new algebraic identities to the algebra of messages. Bieber (1990) does not assume that the encryption scheme is a free algebra, following an idea due to Merritt and Wolper (1985). Even et al. (1985) analyze ping-pong protocols under RSA, taking the actual encryption scheme into account. The applied π -calculus of Abadi and Fournet (2001) permits the definition of an equational theory over the messages exchanged between processes, weakening some of the encryption scheme assumptions when the applied π -calculus is used to analyze security protocols. Since the encryption scheme used in our framework is a simple parameter to the logic, there is no difficulty in modifying our logic to reason about a particular encryption scheme, and hence we can capture these approaches in our framework. However, again, it seems that

our approach is more flexible than these other approaches; not all adversaries can be defined simply by starting with a Dolev-Yao adversary and adding identities.

Recent tools for formal analysis of security protocols in the symbolic setting have moved to a general way of describing adversaries based on specifying the cryptosystem (including the adversary’s abilities) using an equational theory. This can be used to model the Dolev-Yao adversary, but also can move beyond. AVISPA (Viganò, 2005) and ProVerif (Abadi and Blanchet, 2005) are representative of that class of analysis tools. It is clear that relying on equational theories leads to the possibility of modeling very refined adversaries, but there are restrictions. For instance, the equational theories used in the analysis must be decidable, in the sense that there must exist an algorithm for determining whether two terms in the cryptosystem are equal with respect to the equational theory. Much of the recent research has focused on identifying large decidable classes of equational theories based on some identifiable characteristic structure of those theories (Abadi and Cortier, 2004, 2005; Chevalier and Rusinowitch, 2006). Clearly, we can model decidable equational theories in our setting by simply implementing the algorithm witnessing the decidability, and in that sense we benefit from the promising work on the subject. But we can also naturally support approximation algorithms for undecidable theories in a completely transparent way, and therefore we are not restricted in the same way. On the other hand, we do not seek to support automated tools, but rather to provide an expressive framework for modeling and specifying security protocols.

Another class of formal approaches to security protocol analysis has recently been proving popular, the class of approaches based on *computational cryptography* (Goldreich, 2001). These approaches take seriously the view that messages are strings of bits, and that adversaries are efficient Turing machines, generally, randomized polynomial-time algorithms. A protocol satisfies a security property in this setting if it satisfies it in the presence of an arbitrary adversary taken from a given class of algorithms. Thus, this corresponds to modeling a security protocol in the presence not of a single adversary, but rather a family of adversaries. Early work by Datta et al. (2005) and more recently the development of automated tools such as CryptoVerif (Blanchet, 2008) has shown the applicability of the approach to protocol analysis.⁶ Our framework is amenable to supporting multiple adversaries via a simple extension: rather than having a single global knowledge algorithm, we can make the knowledge algorithm part of the local state of the adversary—this is in fact the original and most general presentation of algorithmic knowledge (Halpern et al., 1994). With the knowledge algorithm now part of the local state, we can model protocol execution in the presence of a class of adversaries, each represented by its knowledge algorithm. The system generated by protocol P in the presence of a class of adversaries ADV_s is the union of the runs of P executed under each adversary $A \in ADV_s$. The initial states of the system are the initial states of P under each adversary. Intuitively, each run of the system corresponds to a possible execution of protocol P under some nondeterministically chosen adversary in ADV_s . In the case of computational cryptography models, the knowledge algorithms can simply represent all possible polynomial-time algorithms. Such an approach to modeling systems under different adversaries is along the lines of the model developed by Halpern and Tuttle (1993).

⁶An orthogonal line of research investigates the relationship between computational cryptography and the extent to which it can be soundly approximated by more symbolic approaches (Backes et al., 2003; Micciancio and Warinschi, 2004). We do not address this question here.

On a related note, the work of Abadi and Rogaway (2002), building on previous work by Bellare and Rogaway (1993), compare the results obtained by a Dolev-Yao adversary with those obtained by a more computational view of cryptography. They show that, under various conditions, the former is sound with respect to the latter, that is, terms that are assumed indistinguishable in the Dolev-Yao model remain indistinguishable under a concrete encryption scheme. It would be interesting to see the extent to which their analysis can be recast in our setting, which, as we argued, can capture both the Dolev-Yao adversary and more concrete adversaries.

The use of a security-protocol logic based on knowledge or belief is not new. Several formal logics for analysis of security protocols that involve knowledge and belief have been introduced, going back to BAN logic (Burrows et al., 1990), such as (Bieber, 1990; Gong et al., 1990; Syverson, 1990; Abadi and Tuttle, 1991; Stubblebine and Wright, 1996; Wedel and Kessler, 1996; Accorsi et al., 2001). The main problem with some of those approaches is that semantics of the logic (to the extent that one is provided) is typically not tied to protocol executions or attacks. As a result, protocols are analyzed in an idealized form, and this idealization is itself error-prone and difficult to formalize (Mao, 1995).⁷ While some of these approaches have a well-defined semantics and do not rely on idealization (e.g., (Bieber, 1990; Accorsi et al., 2001)), they are still restricted to (a version of) the Dolev-Yao adversary. In contrast, our framework goes beyond Dolev-Yao, as we have seen, and our semantics is directly tied to protocol execution. Other approaches have notions of knowledge that can be interpreted as a form of algorithmic knowledge ((Durgin et al., 2003), for instance), but the interpretation of knowledge is fixed in the semantics of the logic. One limitation that our logic shares with other logics for security protocol analysis based on multiagent systems is that we can only reason about a fixed finite number of agents participating in the protocol. This is in contrast to approaches such as process calculi that can implicitly deal with an arbitrary number of agents.

The problem of logical omniscience in logics of knowledge is well known, and the literature describes numerous approaches to try to circumvent it. (See (Fagin et al., 1995, Chapter 10 and 11) for an overview.) In the context of security, this takes the form of using different semantics for knowledge, either by introducing *hiding* operators that hide part of the local state for the purpose of indistinguishability or by using notions such as *awareness* (Fagin and Halpern, 1988) to capture an intruder's inability to decrypt (Accorsi et al., 2001).⁸ We now describe these two approaches in more detail.

The hiding approach is used in many knowledge-based frameworks as a way to define an essentially standard semantics for knowledge not subject to logical omniscience, at least as far as cryptography is concerned. Abadi and Tuttle (1991), for instance, map all messages that the agent cannot decrypt to a fixed symbol \square ; the semantics of knowledge is modified so that s and s' are indistinguishable to agent i when the local state of agent i in s and s' is the same after applying the mapping described above. Syverson and van Oorschot (1994) use a variant: rather than mapping all messages that an agent cannot decrypt to the same

⁷While more recent logical approaches (e.g., (Clarke et al., 1998; Durgin et al., 2003)) do not suffer from an idealization phase and are more tied to protocol execution, they also do not attempt to capture knowledge and belief in any general way.

⁸A notion of algorithmic knowledge was defined by Moses (1988) and used by Halpern, Moses and Tuttle (1988) to analyze zero-knowledge protocols. Although related to algorithmic knowledge as defined here, Moses' approach does not use an explicit algorithm. Rather, it checks whether there exists an algorithm of a certain class (for example, a polynomial-time algorithm) that could compute such knowledge.

symbol \square , they use a distinct symbol \square_x for each distinct term x of the free algebra modeling encrypted messages, and take states containing these symbols to be indistinguishable if they are the same up to permutation of the set of symbols \square_x . Thus, an adversary may still do comparisons of encrypted messages without attempting to decrypt them. Hutter and Schairer (2004) use this approach in their definition of information flow in the presence of symbolic cryptography, and Garcia et al. (2005) use it in their definition of anonymity in the presence of symbolic cryptography.⁹ This approach deals with logical omniscience for encrypted messages: when the adversary receives a message m encrypted with a key that he does not know, the adversary does not know that he has m if there exists another state where he has received a different message m' encrypted with a key he does not know. However, the adversary can still perform arbitrary computations with the data that he does know. Therefore, this approach does not directly capture *computational limitations*, something algorithmic knowledge takes into account.

Awareness is a more syntactical approach. Roughly speaking, the semantics for awareness can specify for every point a set of formulas of which an agent is aware. For instance, an agent may be aware of a formula without being aware of its subformulas. A general problem with awareness is determining the set of formulas of which an agent is aware at any point. One interpretation of algorithmic knowledge is that it characterizes what formulas an agent is aware of: those for which the algorithm says “Yes”. In that sense, we subsume approaches based on awareness by providing them with an intuition. We should note that not every use of awareness in the security protocol analysis literature is motivated by the desire to model more general adversaries. Accorsi et al. (2001), for instance, describe a logic for reasoning about beliefs of agents participating in a protocol, much in the way that BAN logic is used to reason about beliefs of agents participating in a protocol. To deal with the logical omniscience problem, Accorsi et al. use awareness to restrict the set of facts that an agent can believe. Thus, an agent may be aware of which agent sent a message if she shares a secret with the sender of the message, and not be aware of that fact otherwise. This makes the thrust of their work different from ours.

6. CONCLUSION

We have presented a framework for security analysis using algorithmic knowledge. The knowledge algorithm can be tailored to account for both the capabilities of the adversary and the specifics of the protocol under consideration. Of course, it is always possible to take a security logic and extend it in an ad hoc way to reason about adversary with different capabilities. Our approach has a number of advantages over ad hoc approaches. In particular, it is a general framework (we simply need to change the algorithm used by the adversary to change its capabilities, or add adversaries with different capabilities), and it permits reasoning about protocol-specific issues (for example, it can capture situations such as an agent sending the bits of her key).

Another advantage of our approach is that it naturally extends to the probabilistic setting. For instance, we can easily handle probabilistic protocols by considering multiagent systems with a probability distribution on the runs (see (Halpern and Tuttle, 1993)). We can also deal with knowledge algorithms that are probabilistic, although there are some

⁹A variant of this approach is developed by Cohen and Dam (2005) to deal with logical omniscience in a first-order interpretation of BAN logic. Rather than using a symbol \square to model that a message is encrypted with an unknown key, they identify messages in different states encrypted using an unknown key.

additional subtleties that arise, since the semantics for X_i given here assumes that the knowledge algorithm is deterministic. In a companion paper (Halpern and Pucella, 2005), we extend our approach to deal with probabilistic algorithmic knowledge, which lets us reason about a Dolev-Yao adversary that attempts to guess keys subject to a distribution. We hope to use this approach to capture probabilistic adversaries of the kind studied by Lincoln et al. (1998).

The goal of this paper was to introduce a general framework for handling different adversary models in a natural way, not specifically to devise new attacks or adversary capabilities. In fact, finding new attacks or defining new adversaries are difficult tasks orthogonal to the problem of using the framework. One potential application for the framework would be to formalize and compare within the same knowledge-based framework new attacks that are introduced by the community. We gave a concrete example of this with the “guess-and-confirm” attacks of Lowe (2002).

It is fair to ask at this point what we can gain by using this framework. For one thing, we believe that the ability of the framework to describe the capabilities of the adversary will make it possible to specify the properties of security protocols more precisely. In particular, we can now express security properties directly in terms of knowledge, which has the advantage of matching fairly well the informal specification of many properties, and we can give a clear algorithmic semantics to knowledge based on the capabilities of the adversary. More complex security properties can be expressed by extending the logic with additional operators such as temporal operators, extensions that are all well understood (Fagin et al., 1995). Of course, not all security properties can be conveniently expressed with our logic. Properties such as computational indistinguishability (Goldwasser and Micali, 1984), which is not a property of single executions but of sets of executions, or observational equivalence (Milner, 1980), which requires not only the protocol to be analyzed but also an idealized version of the protocol that is obviously correct, cannot be expressed directly in the logic. It would be of interest to study the sort of extensions required to capture those properties, and others.

Of course, it may be the case that to prove correctness of a security protocol with respect to certain types of adversaries (for example, polynomial-time bounded adversaries) we will need to appeal to techniques developed in the cryptography community. However, we believe that it may well be possible to extend current model-checking techniques to handle more restricted adversaries (for example, Dolev-Yao extended with random guessing). This is a topic that deserves further investigation. In any case, having a logic where we can specify the abilities of adversaries is a necessary prerequisite to using model-checking techniques.

Acknowledgments. A preliminary version of this paper appeared in the *Proceedings of the Workshop on Formal Aspects of Security*, LNCS 2629, pp. 115-132, 2002.

This research was inspired by discussions between the first author, Pat Lincoln, and John Mitchell, on a wonderful hike in the Dolomites. We also thank Sabina Petride for useful comments. Authors supported in part by NSF under grant CTC-0208535, by ONR under grants N00014-00-1-03-41 and N00014-01-10-511, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grant N00014-01-1-0795, and by AFOSR under grant F49620-02-1-0101.

REFERENCES

- M. Abadi and B. Blanchet. Analyzing security protocols with secrecy types and logic programs. *Journal of the ACM*, 52(1):102–146, 2005.
- M. Abadi and V. Cortier. Deciding knowledge in security protocols under equational theories. In *Proc. 31st Colloquium on Automata, Languages, and Programming (ICALP'04)*, volume 3142 of *Lecture Notes in Computer Science*, 2004.
- M. Abadi and V. Cortier. Deciding knowledge in security protocols under (many more) equational theories. In *Proc. 18th IEEE Computer Security Foundations Workshop (CSFW'05)*, pages 62–76. IEEE Computer Society Press, 2005.
- M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *Proc. 28th Annual ACM Symposium on Principles of Programming Languages (POPL'01)*, pages 104–115, 2001.
- M. Abadi and P. Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
- M. Abadi and M. R. Tuttle. A semantics for a logic of authentication. In *Proc. 10th ACM Symposium on Principles of Distributed Computing (PODC'91)*, pages 201–216, 1991.
- R. Accorsi, D. Basin, and L. Viganò. Towards an awareness-based semantics for security protocol analysis. In Jean Goubault-Larrecq, editor, *Proc. Workshop on Logical Aspects of Cryptographic Protocol Verification*, volume 55.1 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science Publishers, 2001.
- M. Backes, B. Pfitzmann, and M. Waidner. A composable cryptographic library with nested operations. In *Proc. 10th ACM Conference on Computer and Communications Security (CCS'03)*, pages 220–230. ACM Press, 2003.
- M. Baudet. Deciding security of protocols against off-line guessing attacks. In *Proc. 12th ACM Conference on Computer and Communications Security (CCS'05)*, pages 16–25. ACM Press, 2005.
- M. Bellare and P. Rogaway. Entity authentication and key distribution. In *Proc. 13th Annual International Cryptology Conference (CRYPTO'93)*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. Springer-Verlag, 1993.
- P. Bieber. A logic of communication in hostile environment. In *Proc. 3rd IEEE Computer Security Foundations Workshop (CSFW'90)*, pages 14–22. IEEE Computer Society Press, 1990.
- B. Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, 2008.
- M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, 1990.
- Y. Chevalier and M. Rusinowitch. Hierarchical combination of intruder theories. In *Proc. 17th International Conference on Rewriting Techniques and Applications (RTA'06)*, 2006.
- E.M. Clarke, S. Jha, and W. Marrero. Using state space exploration and a natural deduction style message derivation engine to verify security protocols. In *Proc. IFIP Working Conference on Programming Concepts and Methods (PROCOMET)*, 1998.
- M. Cohen and M. Dam. Logical omniscience in the semantics of BAN logic. In *Proc. Workshop on Foundations of Computer Security (FCS'05)*, 2005.
- R. Corin, J. Doumen, and S. Etalle. Analysing password protocol security against off-line dictionary attacks. In *Proc. 2nd International Workshop on Security Issues with Petri*

- Nets and other Computational Models (WISP'04)*, volume 121 of *Electronic Notes in Theoretical Computer Science*, pages 47–63. Elsevier Science Publishers, 2005.
- A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, and M. Turuani. Probabilistic polynomial-time semantics for a protocol security logic. In *Proc. 32nd Colloquium on Automata, Languages, and Programming (ICALP'05)*, pages 16–29, 2005.
- D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
- N. A. Durgin, J. C. Mitchell, and D. Pavlovic. A compositional logic for proving security properties of protocols. *Journal of Computer Security*, 11(4):677–722, 2003.
- S. Even, O. Goldreich, and A. Shamir. On the security of ping-pong protocols when implemented using the RSA. In *Proc. Conference on Advances in Cryptology (CRYPTO'85)*, volume 218 of *Lecture Notes in Computer Science*, pages 58–72. Springer-Verlag, 1985.
- R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.
- R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *Journal of the ACM*, 41(2):340–367, 1994.
- R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- F. D. Garcia, I. Hasuo, W. Pieters, and P. van Rossum. Provable anonymity. In *Proc. 3rd ACM Workshop on Formal Methods in Security Engineering (FMSE 2005)*, pages 63–72. ACM Press, 2005.
- O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2001.
- S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and Systems Sciences*, 28(2):270–299, 1984.
- L. Gong, R. Needham, and R. Yahalom. Reasoning about belief in cryptographic protocols. In *Proc. 1990 IEEE Symposium on Security and Privacy*, pages 234–248. IEEE Computer Society Press, 1990.
- A. D. Gordon and A. Jeffrey. Authenticity by typing for security protocols. *Journal of Computer Security*, 11(4):451–520, 2003.
- J. Y. Halpern and K. O'Neill. Secrecy in multiagent systems. In *Proc. 15th IEEE Computer Security Foundations Workshop (CSFW'02)*, pages 32–46. IEEE Computer Society Press, 2002.
- J. Y. Halpern and R. Pucella. On the relationship between strand spaces and multi-agent systems. *ACM Transactions on Information and System Security*, 6(1):43–70, 2003.
- J. Y. Halpern and R. Pucella. Probabilistic algorithmic knowledge. *Logical Methods in Computer Science*, 1(3:1), 2005.
- J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *Journal of the ACM*, 40(4):917–962, 1993.
- J. Y. Halpern, Y. Moses, and M. R. Tuttle. A knowledge-based analysis of zero knowledge. In *Proc. 20th Annual ACM Symposium on the Theory of Computing (STOC'88)*, pages 132–147, 1988.
- J. Y. Halpern, Y. Moses, and M. Y. Vardi. Algorithmic knowledge. In *Proc. 5th Conference on Theoretical Aspects of Reasoning about Knowledge (TARK'94)*, pages 255–266. Morgan Kaufmann, 1994.
- D. Hutter and A. Schairer. Possibilistic information flow control in the presence of encrypted communication. In *Proc. 9th European Symposium on Research in Computer*

- Security (ESORICS'04)*, volume 3193 of *Lecture Notes in Computer Science*, pages 209–224. Springer-Verlag, 2004.
- S. Kripke. A semantical analysis of modal logic I: normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.
- P. Lincoln, J. C. Mitchell, M. Mitchell, and A. Scedrov. A probabilistic poly-time framework for protocol analysis. In *Proc. 5th ACM Conference on Computer and Communications Security (CCS'98)*, pages 112–121, 1998.
- G. Lowe. Analysing protocols subject to guessing attacks. In *Proc. Workshop on Issues in the Theory of Security (WITS'02)*, 2002.
- G. Lowe. An attack on the Needham-Schroeder public-key authentication protocol. *Information Processing Letters*, 56:131–133, 1995.
- G. Lowe. Casper: A compiler for the analysis of security protocols. *Journal of Computer Security*, 6:53–84, 1998.
- W. Mao. An augmentation of BAN-like logics. In *Proc. 8th IEEE Computer Security Foundations Workshop (CSFW'95)*, pages 44–56. IEEE Computer Society Press, 1995.
- C. Meadows. The NRL protocol analyzer: An overview. *Journal of Logic Programming*, 26(2):113–131, 1996.
- M. Merritt and P. Wolper. States of knowledge in cryptographic protocols. Unpublished manuscript, 1985.
- D. Micciancio and B. Warinschi. Soundness of formal encryption in the presence of active adversaries. In *Proc. Theory of Cryptography Conference (TCC'04)*, volume 2951 of *Lecture Notes in Computer Science*, pages 133–151. Springer-Verlag, 2004.
- J. K. Millen, S. C. Clark, and S. B. Freedman. The Interrogator: Protocol security analysis. *IEEE Transactions on Software Engineering*, 13(2):274–288, 1987.
- R. Milner. *A Calculus of Communicating Systems*. Number 92 in *Lecture Notes in Computer Science*. Springer-Verlag, 1980.
- J. Mitchell, M. Mitchell, and U. Stern. Automated analysis of cryptographic protocols using $\text{Mur}\varphi$. In *Proc. 1997 IEEE Symposium on Security and Privacy*, pages 141–151. IEEE Computer Society Press, 1997.
- J. H. Moore. Protocol failures in cryptosystems. *Proceedings of the IEEE*, 76(5):594–602, 1988.
- Y. Moses. Resource-bounded knowledge. In *Proc. 2nd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK'88)*, pages 261–276. Morgan Kaufmann, 1988.
- R. M. Needham and M. D. Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 21(12):993–999, 1978.
- L. C. Paulson. The inductive approach to verifying cryptographic protocols. *Journal of Computer Security*, 6(1/2):85–128, 1998.
- R. Pucella. Deductive algorithmic knowledge. *Journal of Logic and Computation*, 16(2):287–309, 2006.
- P. Y. A. Ryan and S. A. Schneider. An attack on a recursive authentication protocol: A cautionary tale. *Information Processing Letters*, 65(1):7–10, 1998.
- S. Stubblebine and R. Wright. An authentication logic supporting synchronization, revocation, and recency. In *Proc. 3rd ACM Conference on Computer and Communications Security (CCS'96)*. ACM Press, 1996.
- P. Syverson. A logic for the analysis of cryptographic protocols. NRL Report 9305, Naval Research Laboratory, 1990.

- P. Syverson and I. Cervesato. The logic of authentication protocols. In *Proc. 1st International School on Foundations of Security Analysis and Design (FOSAD'00)*, volume 2171 of *Lecture Notes in Computer Science*, pages 63–137, 2001.
- P. F. Syverson and P. C. van Oorschot. On unifying some cryptographic protocol logics. In *Proc. 1994 IEEE Symposium on Security and Privacy*, pages 14–28. IEEE Computer Society Press, 1994.
- F. J. Thayer, J. C. Herzog, and J. D. Guttman. Strand spaces: Proving security protocols correct. *Journal of Computer Security*, 7(2/3):191–230, 1999.
- L. Viganò. Automated security protocol analysis with the AVISPA tool. In *Proc. 21th Conf. Mathematical Foundations of Programming Semantics (MFPS'05)*, volume 155 of *Electronic Notes in Theoretical Computer Science*, pages 61–86. Elsevier Science Publishers, 2005.
- G. Wedel and V. Kessler. Formal semantics for authentication logics. In *Proc. 4th European Symposium on Research in Computer Security (ESORICS'96)*, volume 1146 of *Lecture Notes in Computer Science*, pages 219–241. Springer-Verlag, 1996.