From Qualitative to Quantitative Proofs of Security Properties Using First-Order Conditional Logic: Abstract*

Joseph Y. Halpern
Dept. of Computer Science
Cornell University
Ithaca, NY 14853
halpern@cs.cornell.edu

Security protocols, such as key-exchange and key-management protocols, are short, but notoriously difficult to prove correct. Flaws have been found in numerous protocols, ranging from the the 802.11 Wired Equivalent Privacy (WEP) protocol used to protect link-layer communications from eavesdropping and other attacks [1] to standards and proposed standards for Secure Socket Layer [2], [3] to Kerberos [4]. Not surprisingly, a great deal of effort has been devoted to proving the correctness of such protocols. There are two largely disjoint approaches. The first essentially ignores the details of cryptography by assuming perfect cryptography (i.e., nothing encrypted can ever be decrypted without the encryption key) and an adversary that controls the network. By ignoring the cryptography, it is possible to give a more qualitative proof of correctness, using logics designed for reasoning about security protocols. Indeed, this approach has enabled axiomatic proofs of correctness and model checking of proofs (see, for example, [5], [6]). The second approach applies the tools of modern cryptography to proving correctness, using more quantitative arguments. Typically it is shown that, given some security parameter k (where k may be, for example, the length of the key used) an adversary whose running time is polynomial in k has a negligible probability of breaking the security, where "negligible" means "less than any inverse polynomial function of k" (see, for example, [7], [8]). There has been recent work on bridging the gap between these two approaches, with the goal of constructing a logic that can allow reasoning about quantitative aspects of security protocols while still being amenable to mechanization. This line of research started with the work of Abadi and Rogaway [9]. More recently, Datta et al. [10] showed that by giving a somewhat nonstandard semantics to their first-order *Protocol Composition Logic* [11], it was possible to reason about many features of the computational model. In this logic, an "implication" of the form $\phi \supset B$ is interpreted as, roughly speaking, the probability of B given ϕ is high. For example, a statement like `secret encrypted` $\supset$ `adversary does not decrypt the secret` says "with high probability, if the secret is encrypted, the adversary does not decrypt it". While the need for such statements should be clear, the probabilistic interpretation used is somewhat unnatural, and no axiomatization is provided by Datta et al. [10] for the $\supset$ operator (although some sound axioms are given that use it).

The interpretation of $\supset$ is quite reminiscent of one of the interpretations of $\rightarrow$ in conditional logic, where $\phi \rightarrow \psi$ can be interpreted as "typically, if ϕ then ψ" [12]. Indeed, one semantics given to $\rightarrow$, called *ϵ-semantics* [13], [14], is very close in spirit to that used in [10]; this is particularly true for the formulation of ϵ-semantics given by Goldszmidt, Morris, and Pearl [15]. In this formulation, a formula $\phi \rightarrow \psi$ is evaluated with respect to a sequence $(\Pr_1, \Pr_2, \ldots)$ of probability measures (*probability sequence*, for short): it is true if, roughly speaking, $\lim_{n\rightarrow\infty} \Pr_n(\psi \mid \phi) = 1$ (where $\Pr_k(\psi \mid \phi)$ is taken to be 1 if $\Pr_k(\phi) = 1$). This formulation is not quite strong enough for some security-related purposes, where the standard is *super-polynomial* convergence, that is, convergence faster than any inverse polynomial. To capture such convergence, we can take $\phi \rightarrow \psi$ to be true with respect to this probability sequence if, for all polynomials p, there exists n^* such that, for all $n \geq n^*$, $\Pr_n(\psi \mid \phi) \geq 1 - 1/p(n)$. (Note that this implies that $\lim_{n\rightarrow\infty} \Pr_n(\psi \mid \phi) = 1$.) In a companion paper [16], it is shown that reinterpreting $\rightarrow$ in this way gives an elegant, powerful variant of the logic considered in [10], which can be used to reason about security protocols of interest.

While it is already a pleasant surprise that conditional logic provides such a clean approach to reasoning about security, using conditional logic has two further significant advantages, which are the subject of this paper. The first is that, as I show here, the well-known complete axiomatization of conditional logic with respect to ϵ-semantics continues to be sound and complete with respect to the super-polynomial semantics for $\rightarrow$; thus, the axioms form a basis for automated proofs. The second is that the use of conditional logic allows for a

* This paper originally appeared in *AAAI-08 (Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence)*, 2008, pp. 454–459. The full paper is available at http://arxiv.org/abs/0804.2155. Work supported in part by NSF under under grants ITR-0325453, IIS-0534064, IIS-0812045, IIS-0911036, and CCF-1214844, and by AFOSR under grants FA9550-05-1-0055 and FA9550-08-1-0266, and by ARO grant W911NF-09-1-0281.

clean transition from qualitative to quantitative arguments. To explain these points, I need to briefly recall some well-known results from the literature.

As is well known, the *KLM properties* [12] provide a sound and complete axiomatization for reasoning about $\rightarrow$ formulas with respect to ϵ-semantics [17]. More precisely, if Δ is a collection of formulas of the form $\phi' \rightarrow \psi'$, then Δ *(ϵ-)entails* $\phi \rightarrow \psi$ (that is, for every probability sequence $\mathcal{P}$, if every formula in Δ is true in $\mathcal{P}$ according to ϵ semantics, then so is $\phi \rightarrow \psi$), then $\phi \rightarrow \psi$ is provable from Δ using the KLM properties. This result applies only when Δ is a collection of $\rightarrow$ formulas. Δ cannot include negations or disjunctions of $\rightarrow$ formulas. *Conditional logic* extends the KLM framework by allowing Boolean combinations of $\rightarrow$ statements. A sound and complete axiomatization of propositional conditional logic with semantics given by what are called preferential structures was given by Burgess [18]; Friedman and Halpern [19] proved it was also sound and complete for ϵ-semantics.

Propositional conditional logic does not suffice for reasoning about security. The logic of [10] is first-order; quantification is needed to capture important properties of security protocols. A sound and complete axiomatization for the language of first-order conditional logic, denoted $\mathcal{L}_C$, with respect to ϵ-semantics is given by Friedman, Halpern, and Koller [20]. The first major result of this paper shows a conditional logic formula ϕ is satisfiable in some model M with respect to ϵ-semantics iff it is satisfiable in some model M' with respect to the super-polynomial semantics. It follows that all the completeness results for ϵ-semantics apply without change to the super-polynomial semantics.

I then consider the language $\mathcal{L}_C^0$ which essentially consists of universal $\rightarrow$ formulas, that is, formulas of the form $\forall x_1 \ldots \forall x_n(\phi{\rightarrow}\psi)$, where ϕ and ψ are first-order formulas. As in the KLM framework, there are no nested $\rightarrow$ formulas or negated $\rightarrow$ formulas. The second major result of this paper is to provide a sound and complete axiomatization that extends the KLM properties for reasoning abut when a collection of formulas in $\mathcal{L}_C^0$ entails a formula in $\mathcal{L}_C^0$.

It might seem strange to be interested in an axiomatization for $\mathcal{L}_C^0$ when there is already a sound and complete axiomatization for the full language $\mathcal{L}_C$. However, $\mathcal{L}_C^0$ has some significant advantages. In reasoning about concrete security, asymptotic complexity results do not suffice; more detailed information about security guarantees is needed. For example, we may want to prove that an SSL server that supports 1,000,000 sessions using 1024 bit keys has a probability of 0.999999 of providing the desired service without being compromised. I show how to convert a qualitative proof of security in the language $\mathcal{L}_C^0$, which provides only asymptotic guarantees, to a quantitative proof. Moreover, the conversion shows exactly how strong the assumptions have to be in order to get the desired 0.999999 level of security. Such a conversion is not possible with $\mathcal{L}_C$.

This conversion justifies reasoning at the qualitative level. A qualitative proof can be constructed without worrying about the details of the numbers, and then automatically converted to a quantitative proof for the desired level of security.

References

[1] N. Borisov, I. Goldberg, and D. Wagner, "Intercepting mobile communications: the insecurity of 802.11," in *Proc. 7th Annual International Conference on Mobile Computing and Networking*, 2001, pp. 180–189.

[2] D. Wagner and B. Schneier, "Analysis of the SSL 3.0 protocol," in *Proc. 2nd USENIX Workshop on Electronic Commerce*, 1996.

[3] J. C. Mitchell, V. Shmatikov, and U. Stern, "Finite-state analysis of SSL 3.0," in *Proc. Seventh USENIX Security Symposium*, 1998, pp. 201–216.

[4] G. Bella and L. C. Paulson, "Kerberos version IV: Inductive analysis of the secrecy goals," in *Proc. 5th European Symposium on Research in Computer Security*, ser. LNCS, Volume 1485, J.-J. Quisquater, Ed. Springer-Verlag, 1998, pp. 361–375.

[5] J. Mitchell, M. Mitchell, and U. Stern, "Automated analysis of cryptographic protocols using Murφ," in *Proc. 1997 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 1997, pp. 141–151.

[6] L. C. Paulson, *Isabelle, A Generic Theorem Prover*, ser. Lecture Notes in Computer Science. Springer-Verlag, 1994, vol. 828.

[7] M. Bellare, R. Canetti, and H. Krawczyk, "A modular approach to the design and analysis of authentication and key exchange protocols," in *Proc. 30th Annual Symposium on the Theory of Computing*, 1998.

[8] O. Goldreich, *Foundations of Cryptography, Vol. 1*. Cambridge University Press, 2001.

[9] M. Abadi and P. Rogaway, "Reconciling two views of cryptography (the computational soundness of formal encryption)," in *Proc. IFIP International Conference on Theoretical Computer Science (TCS'00)*, ser. Lecture Notes in Computer Science, vol. 1872. Springer-Verlag, 2000, pp. 3–22.

[10] A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, and M. Turuani, "Probabilistic polynomial-time semantics for a protocol security logic," in *32nd International Colloquium on Automata, Languages, and Programming (ICALP)*, 2005, pp. 16–29.

[11] A. Datta, A. Derek, J. C. Mitchell, and A. Roy, "Protocol composition logic (PCL)," *Electronic Notes Theoretical Computer Science*, vol. 172, pp. 311–358, 2007.

[12] S. Kraus, D. Lehmann, and M. Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics," *Artificial Intelligence*, vol. 44, pp. 167–207, 1990.

[13] E. Adams, *The Logic of Conditionals*. Dordrecht, Netherlands: Reidel, 1975.

[14] M. Goldszmidt and J. Pearl, "Rank-based systems: a simple approach to belief revision, belief update and reasoning about evidence and actions," in *Principles of Knowledge Representation and Reasoning: Proc. Third International Conference (KR '92)*, 1992, pp. 661–672.

[15] M. Goldszmidt, P. Morris, and J. Pearl, "A maximum entropy approach to nonmonotonic reasoning," *IEEE Transactions of Pattern Analysis and Machine Intelligence*, vol. 15, no. 3, pp. 220–232, 1993.

[16] A. Datta, J. Y. Halpern, J. C. Mitchell, R. Pucella, and A. Roy, "Reasoning about conditional probability and concrete security in protocol proofs," 2009, unpublished manuscript.

[17] H. Geffner, "High probabilities, model preference and default arguments," *Mind and Machines*, vol. 2, pp. 51–70, 1992.

[18] J. Burgess, "Quick completeness proofs for some logics of conditionals," *Notre Dame Journal of Formal Logic*, vol. 22, pp. 76–84, 1981.

[19] N. Friedman and J. Y. Halpern, "Plausibility measures and default reasoning," *Journal of the ACM*, vol. 48, no. 4, pp. 648–685, 2001.

[20] N. Friedman, J. Y. Halpern, and D. Koller, "First-order conditional logic for default reasoning revisited," *ACM Trans. on Computational Logic*, vol. 1, no. 2, pp. 175–207, 2000.