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

Joe Halpern
Cornell University

Proving Correctness of Security Protocols

Security protocols are short, but notoriously difficult to prove correct.

Security flaws have been found in, for example,

- the 802.11 Wired Equivalent Privacy (WEP) protocol used to protect link-layer communications from eavesdropping
- standards and proposed standards for SSL (Secure Socket Layer)
- Kerberos
- the Needham-Schroeder public-key authentication protocol.

Some of these protocols are in wide use, so proving security is critical.

Two Approaches

Two (disjoint) approaches have been used for proving security:

1. The “logic” approach (qualitative—no numbers):
   - Ignore details of cryptography—assume crypto is unbreakable
   - Assume that the adversary controls the network
     - Can eavesdrop and inject messages into the system at will
   - Good news: can get axiomatic proofs of correctness.
2. The “crypto” approach (more quantitative):
   - Prove that a poly-time adversary has a negligible probability of causing damage
   - Probability is a function of a security parameter k (e.g., length of key used); “negligible” = $< 1/p(k)$, for all polynomials p

Bridging the Gap

Abadi and Rogaway [2000]: We want the best of both worlds:

- A logic that captures *quantitative* aspects of a protocol
- Machine checkable proofs and model checking

More relevant here: Datta et al. [2005]:

- Logic based on *Protocol Composition Logic* [Datta et al. 2007]
- Key new feature: an "implication" operator $\supset$ where $A \supset B$ means "the probability of B given A is high"
  - Example: `secret encrypted` $\supset$ `adversary does not decrypt the secret`
- $\supset$ has "unnatural" semantics, no axiomatization

Conditional Logic to the Rescue

In AI going back to the 1980s, there was great interest in *default logic*:

- $bird \rightarrow fly$ is interpreted as "birds typically / normally / by default fly"

Philosophers have been interested in *indicative conditionals*

- $a \rightarrow b$ could have a counterfactual interpretation: "if a were the case, then b"

Many interpretations have been given to $\rightarrow$.

- Using partial orders on worlds, possibility measures, ranking functions . . .
- They all lead to the same axiom system (at the propositional level)

The KLM Properties

The axiom system (called the *KLM* properties, after Kraus, Lehmann, and Magidor) has rules such as the following:

AND. From $\varphi \rightarrow \psi_1$ and $\varphi \rightarrow \psi_2$ infer $\varphi \rightarrow \psi_1 \wedge \psi_2$.

OR. From $\varphi_1 \rightarrow \psi$ and $\varphi_2 \rightarrow \psi$ infer $\varphi_1 \vee \varphi_2 \rightarrow \psi$.

CM. From $\varphi \rightarrow \psi_1$ and $\varphi \rightarrow \psi_2$ infer $\varphi \wedge \psi_2 \rightarrow \psi_1$ (cautious monotonicity).

Theorem: (For a number of diffferent semantic interpretations $\models$) if Σ is a finite set of $\rightarrow$ formulas, then $\Sigma \vdash_{KLM} \varphi \rightarrow \psi$ iff $\Sigma \models \varphi \rightarrow \psi$.

Epsilon Semantics

One interpretation of $\rightarrow$ is probabilistic:

- Roughly speaking, $\varphi \rightarrow \psi$ means that $\Pr(\psi \mid \varphi)$ is "high"
- But what does "high" mean?
  - The AND rule fails if $\Pr(\psi \mid \varphi) > 1 - \epsilon$ for any fixed ϵ
- [Goldszmidt, Morris, Pearl, '93] (also [Adams '75]) used not one probability measure, but a sequence:
  - $(\Pr_1, \Pr_2, \ldots) \models^\epsilon \varphi \rightarrow \psi$ if $\lim_{n\rightarrow\infty} \Pr_n(\psi \mid \varphi) = 1$.
  - The conditional probability is "high" if it approaches 1.
  - Historicallly, this was called "ϵ-semantics".

Theorem: The KLM properties are also sound and complete w.r.t. $\models^\epsilon$.

Back to Security

So what does all this have to do with security?

- It turns out that $\rightarrow$ using ϵ semantics is closely related to $\supset$ as defined by Datta et al.
- Moreover, using security, we can provide a concrete interpretation for the sequence $(\Pr_1, \Pr_2, \ldots)$:
  - $\Pr_k$ is the probability induced by security parameter k
  - For example, if cryptographic keys have length 100, use $\Pr_{100}$.

Super-polynomial convergence

Having a limit of 1 is not good enough for some security-related purposes:

- We want *super-polynomial* convergence: convergence faster than any inverse polynomial. Formally:

$$(\Pr_1, \Pr_2, \ldots) \models^{sp} \varphi \rightarrow \psi$$

$$\text{if} \quad \forall k \, \exists n_k \, \forall n \geq n_k \, \Pr_n(\psi \mid \varphi) \geq 1 - 1/n^k.$$

  - Eventually, the probability is greater than $1 - 1/n^k$
  - It suffices to consider polynomials n^k

- The KLM properties are also sound and complete w.r.t. $\models^{sp}$.

First-Order Conditional Logic

Propositional logic cannot capture some security properties of interest.

- We need first-order quantification

Some notation:

- $\mathcal{L}^{fo}$: first-order logic (no $\rightarrow$'s)
- $\mathcal{L}_C$: first-order conditional logic
  - close off under $\rightarrow$, $\wedge$, $\vee$, and quantification

Theorem: [Friedman, Halpern, Koller, 2000] There is a sound and complete axiomatization of $\mathcal{L}_C$ w.r.t. $\models^{\epsilon}$.

What We Want

Old results were for ϵ-semantics; we need results for $\models^{sp}$

- Want the convergence to be faster than any inverse polynomial

Theorem: The same axioms hold for $\models$ and for $\models^{sp}$

So now we have a clean language for reasoning about probabilistic properties of security protocosl. But this language only makes qualitative statements (in the limit).

- We want to make more quantitative statements.

To do this, we need a detour . . .

Quantitative Analogues of KLM

Key observation: all the KLM rules have quantitative analogues!

$(\Pr_1, \Pr_2, \ldots) \models \varphi \rightarrow^r \psi$ if there exists some $n^* \geq 0$ such that for all $n \geq n^*$, $\Pr_n(\psi \mid \varphi) \geq 1 - r$).

Quantitative Rules:

AND^q. From $\varphi \rightarrow^{r_1} \psi_1$ and $\varphi \rightarrow^{r_2} \psi_2$ infer $\varphi \rightarrow^{r_3} \psi_1 \wedge \psi_2$,
where $r_3 = \min(r_1 + r_2, 1)$.

OR^q. From $\varphi_1 \rightarrow^{r_1} \psi$ and $\varphi_2 \rightarrow^{r_2} \psi$ infer $\varphi_1 \vee \varphi_2 \rightarrow^{r_3} \psi$,
where $r_3 = \min(\max(2r_1, 2r_2), 1)$.

CM^q. From $\varphi_1 \rightarrow^{r_1} \varphi_2$ and $\varphi_1 \rightarrow^{r_2} \psi$ infer $\varphi \wedge \varphi_2 \rightarrow^{r_3} \psi$,
where $r_3 = \max(r_1 + r_2, 1)$.

Another Language

Let $\mathcal{L}_C^0$ consist of all formulas of the form

$$\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi),$$

where φ and ψ are first order.

- No negated $\rightarrow$ formulas
- No nesting of $\rightarrow$ formulas (allowed in $\mathcal{L}_C$)

This language seems to suffice for most reasoning about security.

Main Theorem

Theorem: There exists a sound and complete axiomatization $\mathbf{P}^+$ extending the KLM properties for $\mathcal{L}_C^0$.

- There are some nonobvious rules but . . .

Why do we want an axiomatization of $\mathcal{L}_C^0$ when we already have an axiomatization of the richer language $\mathcal{L}_C$?

Key point: Each rule in the axiomatization of $\mathcal{L}_C^0$ has a quantitative analogue.

- This is not true for the axiomatization of $\mathcal{L}_C$.

The Payoff

Let $\mathbf{P}^q$ be the quantitative analogue of the rules in $\mathbf{P}^+$.

A *qualitative instantiation* of $\varphi \rightarrow \psi$ has the form $\varphi \rightarrow^r \psi$.

Theorem: If $\Delta \vdash_{\mathbf{P}^+} \varphi \rightarrow \psi$, then for all $r \in [0,1]$, in polynomial time in the length of the derivation

- we can find a quantitative instantiation Δ^q of Δ
- a derivation $\Delta^q \vdash_{\mathbf{P}^q} \varphi \rightarrow^r \psi$.

This justifies qualititative reasoning:

- You can construct a qualitative proof, without worrying about the numbers, and then convert it automatically to a quantitative proof

We get even more

We can get the conclusion with an arbitrary degree of confidence, by making sure the assumptions hold with high enough probability.

- Like ϵ-δ arguments in calculus
- Tell me the confidence that the conclusion should hold with (ϵ)
  - e.g., the desired degree of security

  and I'll tell you how strong you need to make the assumptions to get that conclusion (δ)
  - e.g., how long you should make the cryptographic keys
- There may be a number of quantitative instantiations that work
  - We can choose the instantiation that is easiest to implement
  - The constraints must just satisfy some linear inequalities

Future Work

Obvious next steps:

- Apply this logic to proving the correctness of security protocols!
  - Work currently ongoing with Datta, Mitchell, Roy, and Sen
  - Seems promising . . .
- We need to add explicit “dynamic-logic-like” aspects to the logic for reasoning about protocols.
  - There are some interactions between probability and these programming constructs
- Are there more efficient proof rules that also have quantitative analogues?

Stay tuned . . .