

Assurance and formal models

(in particular for security protocols)

Assurance

Specification and implementation

(review)

For any system:

- **Specification:** *What is it supposed to do?*
- **Implementation:** *How does it do it?*
- **Correctness:** *Does it really work?*

In security:

- **Specification:** *Policy*
- **Implementation:** *Mechanism*
- **Correctness:** *Assurance*

Assurance vs. security by obscurity

- In many systems, obscurity (not correctness) is a goal. E.g.,
 - spam filters,
 - censorship systems,
 - computer games,
 - military systems.
- Obscurity may *sometimes* help, at least *for a while, in combination* with other precautions.



The enemy knows the system.

[C. Shannon]

For example, motivated attackers typically can learn how cryptosystems work.

It is much easier to protect only the keys.

(See Kerckhoff's principle in cryptography.)

⇒ **The security of a system cannot depend on secrecy of specification or implementation.**

⇒ **Policies must be appropriate, mechanisms must actually be correct.**

Assurance

Some strategies and techniques:

- open design (maybe open source?),
- specifications and proofs,
- testing,
- processes,
- certification,
- economy of mechanism, and the trusted computing base (TCB).



The TCB

Trusted Computing Base: *the collection of hardware, software, and set-up information on which the security of the system depends.*

Also:

- The part of the system that has to be right.
- The part of the system that may appear to violate its security policy.

The TCB (cont.)

Ideally:

- The TCB should be precisely defined, small, and simple.
- The TCB should be specified, tested, and verified.

In practice:

- Often, lots of dubious code is put in the TCB.
- The TCB gets big and not trustworthy.

March 31, 2010 9:50 AM PDT

Vietnamese dissidents targeted by botnet attacks

by [Tom Krazit](#)

Malware that was disguised as a popular Vietnamese-language keyboard driver for Windows users was used to create a botnet, according to blog posts from [Google's Neel Mehta](#) and [McAfee Chief Technical Officer George Kurtz](#). That botnet was then used to target blogs rallying against a bauxite mining project in Vietnam, employing DDoS (Distributed Denial of Service) attacks to shut down those blogs, according to the posts.

Formal models and proofs
(in particular for security protocols)

What is different about security (1)

- Wish for some guarantees despite lucky, powerful, and persistent attackers.
 - Even if the attacker controls the network.
 - Even if a session key is compromised.
 - Even if an insider is dishonest.
 - ...



What is different about security (2)

- Attacks that exploit the limitations of models.
 - Binary-level exploits despite “secure” languages.
 - Power analysis on “secure” cryptography.
 - ...

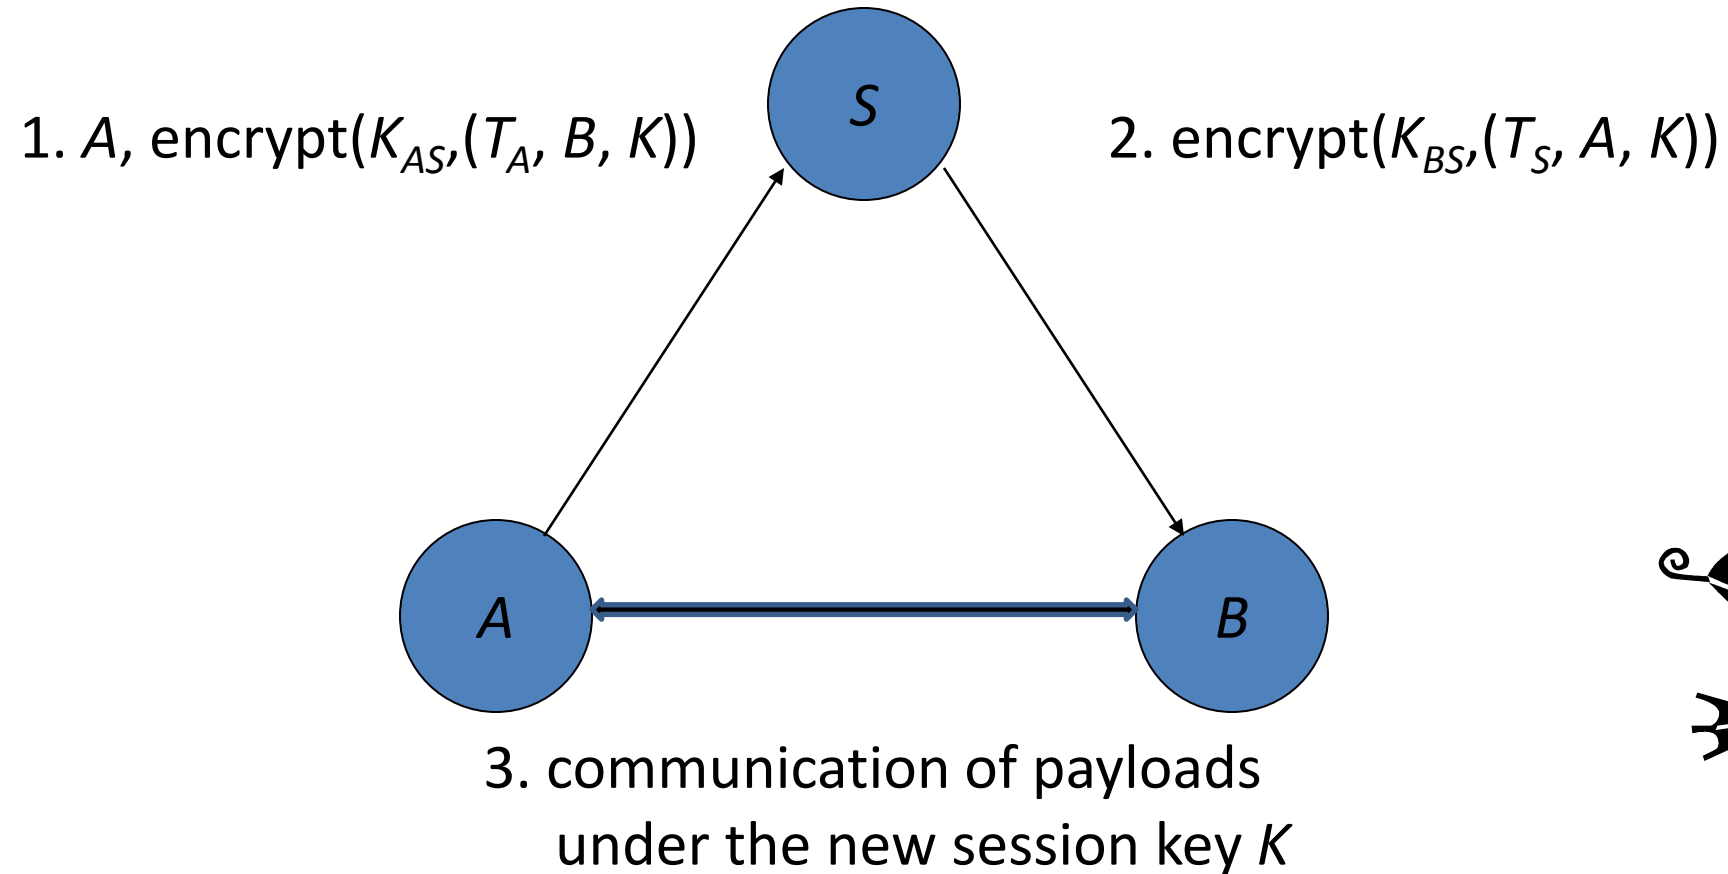


- Doing without full functional correctness:

*Message authenticity and secrecy,
not message correctness.*

These characteristics impact models and proofs.

The WMF protocol (reminder)



T_A, T_S are timestamps.

Here encrypt is symmetric encryption. It may include authentication.

What the messages actually mean

- 1) K is a good key for A and B around time T_A
- 2) A says that K is a good key for A and B around time T_S

Understanding the meaning of messages is central to designing and analyzing protocols.

Even imprecise, informal meanings can be extremely helpful.

A first analysis in a logic of authentication (late 1980s)

- Replace messages with formal representations of their meanings.
- Set out assumptions:
 - S believes (K_{AS} is a good key for A and S)*
 - S believes fresh(T_A)*
 - B believes (A controls (K is a good key for A and B))*
 - ...
- Reason with a few general rules.
- Conclude:
 - A believes (K is a good key for A and B)*
 - B believes (K is a good key for A and B)*

Comments on a logic of authentication

- Served for finding many subtleties and errors.
 - Explained protocols.
 - Highlighted assumptions and conclusions.
 - Used by many people, including protocol designers.
- Lacked a clear link with operational models of protocols or clear cryptographic justification (but see *PCL*).
 - Required more creativity as one moved away from the classic key-exchange protocols.


Some other approaches

(not a complete or orthogonal list)

- Informal but rigorous frameworks based on probabilities and complexity theory.
- Theorem proving, e.g., with Coq or Isabelle.
- Finite-state model checking, e.g., with FDR.
- Type systems and other static analyses for programming languages (and process calculi).

Four current research directions

1. Models, proof techniques, and tools (e.g., type systems).
2. Analysis of particular protocols.
3. Analysis of actual implementations.
4. Relating and combining symbolic and computational approaches.



Some observations

- Most security protocols have ambiguities, subtleties, and flaws.
 - Many of these have to do with cryptography.
 - Many of these don't have to do with the details of cryptography.
- ⇒ *For design, implementation, and analysis, abstract views of cryptography are practical.*

Other low-hanging fruit



An example of a mundane ambiguity:


The simplest fix is to require that a SSL implementation receive a change cipher spec message before accepting a finished message. (Indeed, there is some language in the specification which could be interpreted to mandate this restriction, although it is not entirely clear.) [. . .] at least one implementation has fallen for this pitfall.

(Wagner and Schneier)

The WMF protocol, more abstractly

1. new channel c
for A and B ,
on a long-term channel c_{AS}
between A and S

2. the channel c
on a long-term channel c_{BS}
between B and S



3. communication of payloads
on new private channel c
(e.g., A sends M to B on c)

Towards a language for protocols

- The pi calculus is a general, simple language for concurrent processes that communicate by sending messages on named channels.
- It includes an operator ν (“new”) for generating fresh channels.

$$(\nu c)(\bar{c}\langle M \rangle \mid c(x)\dots)$$

“(new c)(send M on c and, in parallel, receive x on c ...)”

- Here two processes run in parallel.
- One sends M to the other on a fresh channel c .
- x is a bound variable.



Syntax

M, N	$::=$	terms (i.e., data)
		x variable
		n name
P, Q	$::=$	processes (i.e., programs)
		nil nil process (may be omitted)
		$\overline{M}\langle N \rangle.P$ sending
		$M(x).P$ receiving
		$(\nu n)P$ restriction (“new”)
		$P \mid Q$ parallelism
		$!P$ replication

An abstract version of the protocol

$$A(M) = (\nu c) \overline{c}_{AS} \langle c \rangle . \bar{c} \langle M \rangle$$

“With new c , send c on c_{AS} , send M on c .”

An abstract version of the protocol

$$A(M) = (\nu c) \overline{c}_{AS} \langle c \rangle . \bar{c} \langle M \rangle$$

$$S = c_{AS}(x) . \overline{c}_{SB} \langle x \rangle$$

“Receive x on c_{AS} , forward it on c_{SB} .”

An abstract version of the protocol

$$A(M) = (\nu c)\overline{c}_{AS}\langle c\rangle.\overline{c}\langle M\rangle$$

$$S = c_{AS}(x).\overline{c}_{SB}\langle x\rangle$$

$$B = c_{SB}(x).x(y).nil$$

“Receive x on c_{SB} , receive y on x .”

An abstract version of the protocol

$$A(M) = (\nu c)\overline{c}_{AS}\langle c\rangle.\overline{c}\langle M\rangle$$

$$S = c_{AS}(x).\overline{c}_{SB}\langle x\rangle$$

$$B = c_{SB}(x).x(y).nil$$

$$P(M) = (\nu c_{AS})(\nu c_{SB})(S \mid A(M) \mid B)$$

“With new c_{AS} and c_{SB} ,
run S , $A(M)$, and B in parallel.”

Adding concurrent sessions

$$A(M) = (\nu c)\overline{c}_{AS}\langle c \rangle.\overline{c}\langle M \rangle$$

$$S = !c_{AS}(x).\overline{c}_{SB}\langle x \rangle$$

$$B = !c_{SB}(x).x(y).nil$$

$$P(M_1, \dots, M_n) = (\nu c_{AS})(\nu c_{SB})(S \mid A(M_1) \mid \dots \mid A(M_n) \mid B)$$

Secrecy as equivalence

- Secrecy properties can be phrased as equivalences between processes.
 - For example, $P(M)$ and $P(N)$ are equivalent, for all M and N .

Secrecy as equivalence

- Secrecy properties can be phrased as equivalences between processes.
 - For example, $P(M)$ and $P(N)$ are equivalent, for all M and N .

Here, many notions of “equivalence” will do.

Secrecy as equivalence

- Secrecy properties can be phrased as equivalences between processes.
 - For example, $P(M)$ and $P(N)$ are equivalent, for all M and N .
- Other security properties can also be presented and proved formally, as equivalences or as properties of executions.


Here, many notions of “equivalence” will do.

Extending the pi calculus

- In the pure pi calculus, we can easily represent systems like



- But it is much harder (or impossible) to represent the use of cryptography, as in:



The applied pi calculus

We add function symbols, as in:

$(\nu k)(\dots \text{encrypt}(k, M) \dots \text{decrypt}(k, x) \dots)$

- Here the operator ν generates a key.
- Encryption and decryption are function symbols, with equations.

$$\text{decrypt}(\text{encrypt}(x, y), x) = y$$

Expressiveness

- Representing protocols such as WMF is just “a matter of programming”.
- For example, we may write:

$$(\nu k).\bar{c}\langle \text{encrypt}(k, 0) \rangle$$

Here, a process reveals a term that uses a fresh name k without revealing k itself.

- This does not arise in the pure pi calculus.
- It is a source of expressiveness and complications.

Syntax for terms

M, N	$::=$	terms
		x variable
		n name
		$f(M_1, \dots, M_k)$ function application

where f is a function symbol of arity k
(and optionally also with conditions on types)

Syntax for processes

$P, Q ::=$	processes
nil	nil process
$\overline{M}\langle N \rangle.P$	sending
$M(x).P$	receiving
$(\nu n)P$	restriction
$P \mid Q$	parallelism
$!P$	replication
<i>if $M = N$ then P else Q</i>	<i>conditional</i>

Equality

Equality is defined by an equational theory (basically, a set of equations). For example:

- For pairs:

$$\text{fst}(\text{pair}(x, y)) = x$$

$$\text{snd}(\text{pair}(x, y)) = y$$

Equality

Equality is defined by an equational theory (basically, a set of equations). For example:

- For pairs:

$$\text{fst}(\text{pair}(x, y)) = x$$

$$\text{snd}(\text{pair}(x, y)) = y$$

- For symmetric encryption:

$$\text{decrypt}(\text{encrypt}(x, y), x) = y$$

Equality (cont.)

- For asymmetric encryption:

$$\text{adecrypt}(\text{sk}(x), \text{aencrypt}(\text{pk}(x), y)) = y$$

Equality (cont.)

- For asymmetric encryption:

$$\text{adecrypt}(\text{sk}(x), \text{aencrypt}(\text{pk}(x), y)) = y$$

and optionally with other equations, e.g.,

$$\text{pdecrypt}(x, \text{pencrypt}(y, z)) = \text{pencrypt}(y, \text{pdecrypt}(x, z))$$

Equality (cont.)

- For asymmetric encryption:

$$\text{adecrypt}(\text{sk}(x), \text{aencrypt}(\text{pk}(x), y)) = y$$

and optionally with other equations, e.g.,

$$\text{pdecrypt}(x, \text{pencrypt}(y, z)) = \text{pencrypt}(y, \text{pdecrypt}(x, z))$$

- For probabilistic encryption:

$$\text{adecrypt}(\text{sk}(x), \text{aencrypt}(\text{pk}(x), y, z)) = y$$

Other examples

- MACs
- Digital signatures
- One-way hash functions
- XOR
- Exponentiation as used in Diffie-Hellman
- Errors
- ...

Semantics: reduction

$$\bar{c}\langle M \rangle.P \mid c(x).Q \rightarrow P \mid Q[M/x]$$

where $Q[M/x]$ is the result of replacing x with M in Q

$$\text{if } M = M \text{ then } P \text{ else } Q \rightarrow P$$

$$\text{if } M = N \text{ then } P \text{ else } Q \rightarrow Q$$

$$\text{if } M \neq N$$

(In addition, some other trivial rules allow rearranging processes, e.g., by commutativity and associativity of parallel composition.)

Semantics: structural equivalence

$$P \mid nil \equiv P$$

$$P \mid Q \equiv Q \mid P$$

$$P \mid (Q \mid R) \equiv (P \mid Q) \mid R$$

$$!P \equiv P \mid !P$$

$$(\nu m)(\nu n)P \equiv (\nu n)(\nu m)P$$

$$(\nu n)nil \equiv nil$$

$$(\nu n)(P \mid Q) \equiv P \mid (\nu n)Q \quad \text{if } n \notin fn(P)$$

$$P[M/x] \equiv P[N/x] \quad \text{if } M = N$$

Equivalence

- Two processes P and Q are *testing equivalent* if no context R can distinguish them.
 - For a given channel n , $P \mid R$ may output on n if and only if $Q \mid R$ may output on n .
 - The context R may represent an attacker.

Equivalence

- Two processes P and Q are *testing equivalent* if no context R can distinguish them.
 - For a given channel n , $P \mid R$ may output on n if and only if $Q \mid R$ may output on n .
 - The context R may represent an attacker.
- This equivalence is coarse enough that, for example, it relates the two processes:

$(\nu k).\bar{c}\langle \text{encrypt}(k, 0) \rangle$ $(\nu k).\bar{c}\langle \text{encrypt}(k, 1) \rangle$

ProVerif

- ProVerif is a mature tool for formal proofs.
- It has been applied to many protocols.
- Its input language is the applied pi calculus.
- Internally, ProVerif compiles systems into logical formulas (Horn clauses)
 - for expressing protocol actions, e.g., “if a message x is received, then x is forwarded”,
 - for expressing knowledge, e.g., “if the attacker has a ciphertext and the key then it has the plaintext”.

A small example (first informally)

Message 1. $A \rightarrow B : \text{penc}((n, pK_A), pK_B)$

Message 2. $B \rightarrow A : \text{penc}((n, k), pK_A)$

Message 3. $A \rightarrow B : \text{senc}(s, k)$

n is a fresh nonce, k is a fresh shared key.

A small example (a simple version, formally)

$$P \triangleq (\nu sK_A)(\nu sK_B)$$

let $pK_A = \text{pk}(sK_A)$ *in* *let* $pK_B = \text{pk}(sK_B)$ *in*
 $\bar{e}\langle pK_A \rangle . \bar{e}\langle pK_B \rangle . (A \mid B)$

$$A \triangleq (\nu n)(\bar{e}\langle \text{penc}((n, pK_A), pK_B) \rangle \mid$$

e(z).*let* $(x, y) = \text{pdec}(z, sK_A)$ *in*
if $x = n$ *then* $\bar{e}\langle \text{senc}(s, y) \rangle$

$$B \triangleq e(z).let (x, y) = \text{pdec}(z, sK_B) \text{ in}$$

$(\nu k)(\bar{e}\langle \text{penc}((x, k), y) \rangle \mid e(z').let s' = \text{sdec}(z', k) \text{ in } 0)$

A small example

(in ProVerif, a little optimized)

$attacker(x) \wedge attacker(y) \rightarrow attacker(\text{penc}(x, y))$
 $attacker(x) \rightarrow attacker(\text{pk}(x))$
 $attacker(\text{penc}(m, \text{pk}(k))) \wedge attacker(k) \rightarrow attacker(m)$
 $attacker(x) \wedge attacker(y) \rightarrow attacker(\text{senc}(x, y))$
 $attacker(\text{senc}(m, k)) \wedge attacker(k) \rightarrow attacker(m)$
 $attacker(x) \wedge attacker(y) \rightarrow \text{mess}(x, y)$
 $\text{mess}(x, y) \wedge attacker(x) \rightarrow attacker(y)$

$attacker(e[])$

$attacker(\text{pk}(sK_A[]))$
 $attacker(\text{pk}(sK_B[]))$

$attacker(\text{penc}((n[], \text{pk}(sK_A[])), \text{pk}(sK_B[])))$
 $attacker(\text{penc}((n[], x), \text{pk}(sK_A[]))) \rightarrow attacker(\text{senc}(s[], x))$

$attacker(\text{penc}((x, y), \text{pk}(sK_B[])))$
 $\rightarrow attacker(\text{penc}((x, k[\text{penc}((x, y), \text{pk}(sK_B[]))]), y))$

Can we derive $attacker(s[])$?

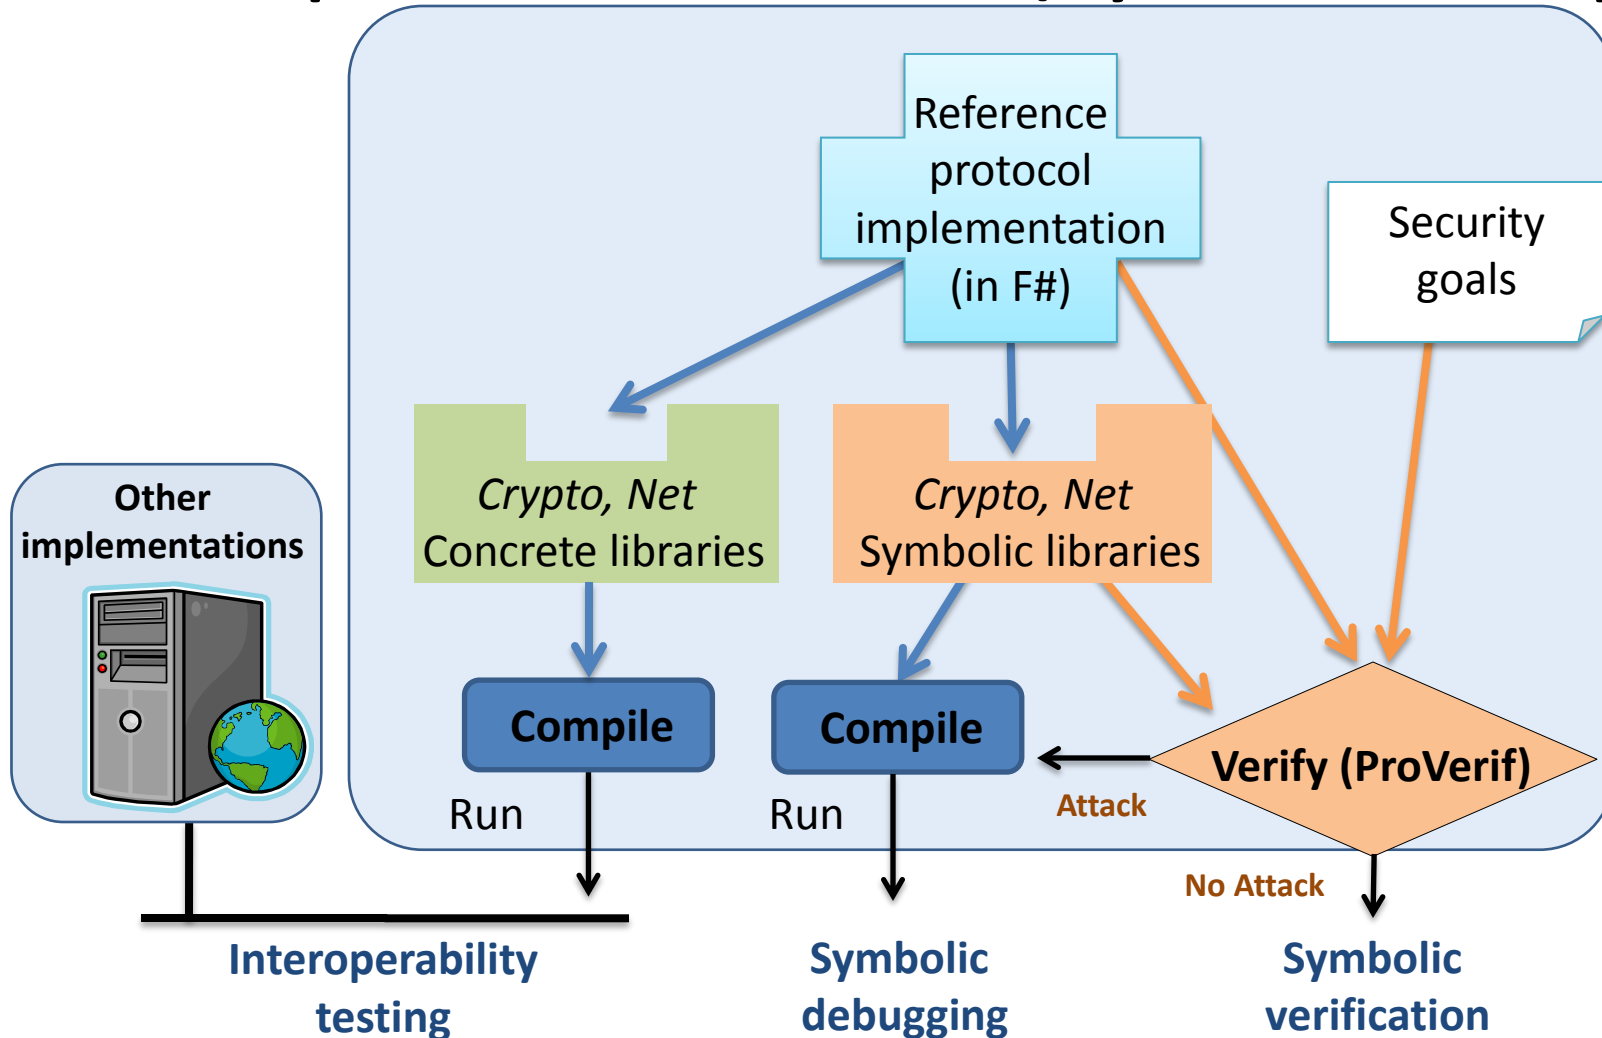
ProVerif (cont.)

- ProVerif can prove various kinds of properties:
 - correspondence properties, of the form “in any behavior where event e happens event e' must also happen”,
 - secrecy properties of the form “the attacker does not obtain s in any behavior”,
 - certain important equivalences, e.g., equivalences that express secrecy properties.
- Proofs are automatic and reasonably fast (milliseconds to minutes, typically).

Analyzing the WMF protocol in ProVerif

- We do some trivial syntactic transformation to fit into the actual ProVerif syntax.
- ProVerif establishes secrecy and authenticity properties automatically.
- E.g., $WMF(payload) \sim WMF(payload')$ where \sim is observational equivalence.
- *It just works!*

Analyzing reference implementations (symbolically)



A TLS reference implementation

- A subset of TLS in F# (10 kLOC), supports:
 - SSL3.0, TLS1.0, TLS1.1 with session resumption,
 - ciphersuites using DES, AES, RC4, SHA1, MD5,
 - server-only authentication, RSA mode only,
 - no compression, fragmentation, or alerts.
- Tested in a few basic scenarios, including:
 - an HTTPS client (with IIS and Apache servers),
 - an HTTPS server (with IE and Firefox clients).

Sample code

HTTPS Client Implementation

```
let client_request url =  
  let tcp = Net.connect url in  
  let connectionId, sessionId =  
    Handshake.connect tcp url in  
  let request = httpRequest url in  
  Record.send connectionId request;  
  let response =  
    Record.recv connectionId in  
  response
```

Open TCP connection

Run Handshake protocol

Build HTTP request

Send it over Record protocol


Get response over Record protocol

Symbolic verification: some results

- Formal analysis is still non-trivial.
- It requires some tweaks in protocol code.
- Still, it is feasible and increasingly practical.

Verified TLS code	Time	Memory
Handshake (authenticity queries)	16 sec	60 MB
Handshake (secrecy queries)	10 sec	80 MB
Handshake + Resumption (resumption authenticity queries)	4 min	460 MB
Handshake + Resumption + Record (record authenticity queries)	6 min	700 MB
Handshake + Resumption + Record	2 hours	1.7 GB

Meanwhile, in the computational world...



An encryption scheme consists of functions $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ on bitstrings for key generation, encryption, and decryption.

It is *IND-CPA* and *INT-CTXT* if:

- **Secrecy:** For all PPTIME adversaries \mathcal{C} ,

$$\Pr \left[k \xleftarrow{R} \mathcal{K}(\eta) : \mathcal{C}^{\mathcal{E}_k(\cdot)}(\eta) = 1 \right] - \Pr \left[k \xleftarrow{R} \mathcal{K}(\eta) : \mathcal{C}^{\mathcal{E}_k(0^{|\cdot|})}(\eta) = 1 \right]$$

is negligible (“very small”) as a function of the security parameter η .

- **Integrity:**

A soundness theorem

[Comon-Lundh & Cortier]

- If $P \sim_s P'$ in the formal model, then P and P' are computationally indistinguishable.
- Assumptions:
 - The encryption scheme is IND-CPA and INT-CTXT.
 - Only the key-generation algorithm creates keys.
 - There are no encryption cycles.
 - There is an ordering $<$ on private keys such that, if k appears in the plaintext of a ciphertext encrypted under k' , then $k < k'$.
 - We have proved this manually for WMF.
 - It is possible to compute a symbolic representation of any bitstring.
 - This "parsing assumption" may not be needed but eases proofs.

Applying the theorem



- A priori, \sim (used with ProVerif) is different from \sim_s (treated by the theorem).
- They can be reconciled by
 - *in the ProVerif model*: adding functions that reveal the length or structure of plaintexts (not so easy),
or
 - *in the computational model*: requiring encryption to be length-concealing.
- So $WMF(\text{payload})$ and $WMF(\text{payload}')$ are computationally indistinguishable, either way.

CryptoVerif

- CryptoVerif is a tool for computational proofs via sequences of games.
- Its input language is basically a process calculus for describing games.
- Its scope overlaps that of ProVerif.
- It is more recent and still a little less mature than ProVerif.

Analyzing the WMF protocol in CryptoVerif (1)

- We also model the protocol in CryptoVerif.
E.g., for A,

***!**^N $c_2(xA : host, xB : host)$; **if** $xA = a \vee xA = b$ **then**
let $KAs = (\text{if } xA = a \text{ then } Kas \text{ else } Kbs)$ **in**
new $rKab : keyseed$; **let** $Kab : key = kgen(rKab)$ **in**
new $r : seed$; $\overline{c_3}\langle xA, \text{encrypt}(\text{concat}(c0, xB, Kab), KAs, r) \rangle$*

Note the bounded replication.

Analyzing the WMF protocol in CryptoVerif (2)

- We assume:
 - Encryption is IND-CPA and INT-CTXT.
 - The function concat returns bitstrings of constant length.
 - x, y, z can be computed from $\text{concat}(x, y, z)$ in PTIME.
 - All payloads have the same length.

Comparison of assumptions

- We do not assume that keys come only from the key-generation algorithm.
- We do not have any parsing assumption.
- We do not assume the absence of encryption cycles.
 - The success of the game-transformation sequence implies a key hierarchy.



Analyzing the WMF protocol in CryptoVerif (3)



- We have to tweak the code:
We manually make a case distinction between honest and dishonest interlocutors of A .
- Then CryptoVerif proves automatically the desired computational indistinguishability.

Lessons

- Soundness theorems often require more hypotheses.
- But, when the hypotheses are met, symbolic models and proofs suffice, and they are generally easier to obtain.




More work remains

- *Adapt symbolic tools to the hypotheses of computational-soundness theorems:*
allow length-revealing encryption,
prove the absence of cycles, ...
- *Extend computational-soundness theorems:*
allow private channels, nested replications, ...
- *Develop computationally sound provers:*
more automation and/or more user guidance,
more primitives and game transformations, ...

Outlook

- There has been a lot of progress!
- Several approaches now work reasonably well.
- They still require more research and a fair amount of engineering.
- Which is most effective and fruitful remains open to debate. For example:
 - how much trouble is it to get computational guarantees?
 - when is it worthwhile?



Reading

- My “Security Protocols: Principles and Calculi (Tutorial Notes)”, and its references
<http://users.soe.ucsc.edu/~abadi/Papers/fosad-protocols.pdf>
- “Modular Verification of Security Protocol Code by Typing”, by Bhargavan, Fournet, and Gordon (as sample of more recent work)

<http://research.microsoft.com/en-us/um/people/adg/publications/modular-verification-popl10.pdf>



G 163270