

Authenticity by Typing for Security Protocols

Andrew D. Gordon
Microsoft Research

Alan Jeffrey
DePaul University

May 2001

Technical Report
MSR-TR-2001-49

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Publication History

A portion of this work will appear in the proceedings of the *14th IEEE Computer Security Foundations Workshop (CSFW 14)*, Cape Breton, June 11–13, 2001. The proceedings will be published by the IEEE Computer Society.

Affiliation

Alan Jeffrey is with DePaul University. The two authors completed part of this work at Microsoft Research in Cambridge and part at DePaul University in Chicago.

Abstract

We propose a new method to check authenticity properties of cryptographic protocols. First, code up the protocol in the spi-calculus of Abadi and Gordon. Second, specify authenticity properties by annotating the code with correspondence assertions in the style of Woo and Lam. Third, figure out types for the keys, nonces, and messages of the protocol. Fourth, check that the spi-calculus code is well-typed according to a novel type and effect system presented in this paper. Our main theorem guarantees that any well-typed protocol is robustly safe, that is, its correspondence assertions are true in the presence of any opponent expressible in spi. It is feasible to apply this method by hand to several well-known cryptographic protocols. It requires little human effort per protocol, puts no bound on the size of the opponent, and requires no state space enumeration. Moreover, the types for protocol data provide some intuitive explanation of how the protocol works. This paper describes our method and gives some simple examples. Our method has led us to the independent rediscovery of flaws in existing protocols and to the design of improved protocols.

Contents

| | | |
|----------|---|-----------|
| 1 | Verifying Correspondences by Typing Spi | 1 |
| 2 | Programming Protocols | 3 |
| 2.1 | Review of the Spi-Calculus | 4 |
| 2.2 | Programming an Example | 6 |
| 3 | Specifying Protocols | 7 |
| 3.1 | A Spi-Calculus with Correspondence Assertions | 7 |
| 3.2 | Specifying the Example | 9 |
| 3.3 | Fixing the Example | 10 |
| 4 | Typing Protocols | 11 |
| 4.1 | Types for Messages | 11 |
| 4.2 | Effects for Processes | 13 |
| 4.3 | Typing Rules | 18 |
| 4.4 | Typing the Example | 24 |
| 5 | Further Protocol Examples | 25 |
| 6 | Summary and Conclusion | 26 |
| A | Protocol Examples | 27 |
| A.1 | Abadi and Gordon's Variant of Wide Mouth Frog | 27 |
| A.2 | Woo and Lam's Authentication Protocol | 32 |
| A.3 | Otway and Rees's Key Exchange Protocol | 35 |
| A.4 | A Secure Message Stream | 36 |
| A.5 | Abbreviations Used in Examples | 38 |
| B | Formal Semantics of our Typed Spi-Calculus | 40 |
| B.1 | A Trace Semantics for our Spi-Calculus | 40 |
| B.2 | Correspondence Traces and Safe Processes | 44 |
| B.3 | Proof of Subject Reduction | 46 |
| B.4 | Proof of Safety | 55 |
| | References | 63 |

1 Verifying Correspondences by Typing Spi

We propose a new method for analysing authenticity properties of cryptographic protocols. Our proposal builds on and develops two existing ideas: Woo and Lam’s idea of correspondence assertions for specifying authentication properties of protocols [41], and Abadi’s idea of checking security properties of cryptographic protocols by type-checking [1].

Woo and Lam’s idea of correspondence assertions is very simple. Starting from some description of the sequence of messages exchanged by principals in a protocol, we annotate it with labelled events marking the progress of each principal through the protocol. Moreover, we divide these events into two kinds, begin-events and end-events. Event labels typically indicate the names of the principals involved and their roles in the protocol. For example, before running a protocol to authenticate its presence to another principal B , an initiator A asserts a begin-event labelled “initiator A authenticating itself to responder B ”. After satisfactory completion of the protocol, the principal B asserts an end-event with the same label. A protocol satisfies these assertions if in all protocol runs, and in the presence of a hostile opponent, every assertion of an end-event corresponds to a distinct, earlier assertion of a begin-event with the same label. The hostile opponent can capture, modify, and replay messages, but cannot forge assertions.

Woo and Lam’s paper [41] describes a formal semantics for correspondence assertions but suggests no verification techniques. Marrero, Clarke, and Jha [30] base a model-checker for security protocols on correspondence assertions. This paper formalises correspondence assertions as new commands in the spi-calculus [3], a concurrent programming language equipped with abstract forms of cryptographic primitives. We expect it would not be difficult to adapt the techniques of this paper to other concurrent languages.

There is a variety of different formulations of authenticity properties of protocols, and even a little controversy [6, 16, 27, 13]. Still, we adopt correspondence assertions because they are simple, precise, and flexible. They are simple annotations of a protocol expressed as a program. They have a precise semantics. They are flexible in the sense that by annotating a protocol in different ways we can express different authenticity intentions and guarantees. Correspondence assertions allow us to express what Lowe [27] calls injective agreement between protocol runs. In a formal comparison of authenticity properties, Focardi, Gorrieri, and Martinelli [14] formulate a property that systematically generalises the equational properties proved in the original work on spi [3], and show that this generalisation is strictly weaker than agreement. Therefore, there is some evidence that the authentication properties proved in this paper are at least as strong as in the original work.

By verifying suitable correspondence assertions, our method can rule out problems such as vulnerability to replay attacks or confusions of identity. Still, like most other formal methods for analysing authenticity protocols, our method deliberately abstracts from the details of the underlying encryption algorithms, and therefore cannot detect protocol weaknesses deriving from inadequacies in these algorithms.

Abadi’s idea of type-checking secrecy properties of cryptographic protocols in the spi-calculus is part of a surge of interest in types for security. Other work includes

type systems for checking untrusted mobile code [26, 32, 19], for checking access control [25, 37], and, most recently, other type systems for cryptographic primitives [35, 2]. Abadi’s original system establishes secrecy properties, and features some unusual constructs that allow any opponent to be type-checked. This paper develops some of the constructs in Abadi’s system, and proposes a new type and effect system [15, 29] for the spi-calculus. For a well-typed program containing correspondence assertions, a type safety theorem guarantees the program satisfies the assertions.

Our new method is the following. First, code up the protocol in the spi-calculus. Second, specify authenticity properties expected of the protocol by annotating the code with correspondence assertions. Third, figure out types for the keys, nonces, and messages of the protocol. Fourth, check that the spi-calculus code is well-typed. The type safety theorem guarantees the soundness of the authenticity properties specified in the second step. The theorem asserts these properties hold in the presence of an opponent represented by an arbitrary spi process. Therefore, a limitation of the theorem is that it does not rule out attacks that cannot be expressed in the spi-calculus. On the other hand, it does not limit the size of the attacker in any way. We have applied this method to several protocols by hand, and have re-discovered some known flaws.

Our method is one of only a few formal analyses that require little human effort per protocol, while putting no bound on the size of the protocol or opponent. Other examples include Song’s mechanisation [38] of strand spaces [39], Heather and Schneider’s algorithm [24, 22] for computing Schneider’s rank functions [36], and Cohen’s resolution-based theorem prover TAPS [10]. Non-examples include most approaches based on model-checking [28], which are automatic but require bounds on the size of the opponent or the protocols, and most approaches based on theorem-proving [8, 34], which impose no bound on opponent or protocol size, but require lengthy and expert human intervention.

Our method is also one of only a few where analysing a protocol involves no exploration or enumeration of the possible states or messages of the protocol, and so is decidable even for protocols with no bound on the size of the principals. The only other such methods we know of are those based on proof-checking belief logics [9, 17]. Like constructing a proof in a belief logic, the work of devising types for a protocol in our system amounts to writing down a formal argument explaining the protocol. Failing to find a proof or a typing can suggest possible attacks on the protocol. Unlike most belief logics, our method has a precise computational basis.

In this paper, we only consider type checking, not type synthesis. Type checking (where the computer checks user-defined typings) is easily seen to be decidable, and provides a straightforward top-down algorithm for protocol verification. Type synthesis (where the computer derives the typings itself) would be harder.

In summary, our new method enjoys a rare and attractive combination of strengths:

- It needs little human effort per protocol.
- It puts no bound on the size of the principals.
- It needs no state space enumeration per protocol.
- It has a precise computational foundation.
- It is decidable.

On the other hand, the type system on which our method is based has limitations. Like all type systems, it is incomplete in the sense that perfectly well-behaved code can fail to type-check. For example, we have found that certain uses of nonces cannot be type-checked. Our system is also limited to symmetric-key cryptography. We leave the study of types for other cryptographic primitives as future work.

The new technical contribution of this paper is a type and effect system for proving correspondence assertions that supports the cryptographic primitives of the spi-calculus. A series of examples supports its usefulness. In earlier work [18], we proposed a type system for proving correspondence assertions about non-cryptographic communication protocols in the π -calculus. The system of the present paper copes with untrusted opponents, encryption primitives, and synchronisation via nonce handshakes, additional features essential for cryptographic protocols.

Contents of this Paper

Section 2 presents the spi-calculus, and illustrates programming of security protocols. Section 3 extends the spi-calculus with correspondence assertions, and shows how they can specify authenticity properties. Section 4 describes our type and effect system. Section 5 discusses further examples. Section 6 concludes.

2 Programming Protocols

This section reviews the syntax and informal semantics of the spi-calculus, and explains how to express a simple protocol example as a spi-calculus program.

Milner, Parrow, and Walker's π -calculus [31] is a parsimonious formalism for concurrency. It explains many different kinds of computation by reducing them to exchanges of names on communication channels. An important constituent of the calculus is a name generation operator for generating fresh names, which identify communication channels.

Abadi and Gordon's spi-calculus [3] is an extension of the π -calculus with abstract forms of encryption and decryption, akin to the idealised versions introduced by Dolev and Yao [12]. The atomic names of the spi-calculus represent the random numbers of cryptographic protocols, such as encryption keys and nonces, as well as channels. The name generation operator abstractly represents the fresh generation of unguessable random numbers such as keys and nonces. We can describe cryptographic protocols by programming them in the spi-calculus.

There are several existing spi-calculus techniques, such as notions of bisimulation, for reasoning about authenticity properties. The new contribution of this paper is a type system for reasoning about authenticity. Our preliminary experience is that establishing authenticity properties by typing is much less labour intensive than constructing bisimulations.

2.1 Review of the Spi-Calculus

There are in fact several versions of spi. The main difference between the spi-calculus presented in this section and the original version [3] is that each binding occurrence of a name is annotated with a type, T . (We postpone defining the set of types till Section 4.) Choosing these type annotations is part of our verification method; they are needed for type-checking processes, but do not affect the runtime behaviour of processes.

We assume an infinite set of atomic names or variables, ranged over by m, n, x, y , and z . For the sake of simplicity in presenting our type system, this version of the spi-calculus, unlike the original, does not distinguish names from variables. The set of messages, which includes the set of names, is given by the grammar in the following table.

Names and Messages:

| | |
|-------------------|-------------------------------------|
| m, n, x, y, z | name: variable, channel, nonce, key |
| $L, M, N ::=$ | message |
| x | name |
| (M, N) | pair |
| $()$ | empty tuple |
| $\text{inl } (M)$ | left injection |
| $\text{inr } (M)$ | right injection |
| $\{M\}_N$ | encryption |

- A message (M, N) is a pair, and $()$ is an empty tuple. With these primitives we can describe any finite record.
- Messages $\text{inl } (M)$ and $\text{inr } (M)$ are tagged unions, differentiated by the distinct tags inl and inr . With these primitives we can encode any finite tagged union.
- A message $\{M\}_N$ is the ciphertext obtained by encrypting the plaintext M with the symmetric key N .

We regard messages as abstract representations of the bit strings manipulated by cryptographic protocols. We assume there is enough redundancy in the format that we can tell apart the different kinds of messages.

Free names of a message $fn(M)$:

| |
|---|
| $fn(x) \triangleq \{x\}$ |
| $fn(() \triangleq \emptyset$ |
| $fn(M, N) \triangleq fn(M) \cup fn(N)$ |
| $fn(\text{inl } (M)) \triangleq fn(M)$ |
| $fn(\text{inr } (M)) \triangleq fn(M)$ |
| $fn(\{M\}_N) \triangleq fn(M) \cup fn(N)$ |

We write $M\{x \leftarrow N\}$ for the outcome of a capture-avoiding substitution of the message N for each free occurrence of the name x in the message M .

The set of processes is defined by the grammar:

Processes:

| $O, P, Q, R ::=$ | process |
|--|-----------------|
| out $M N$ | output |
| inp $M (x:T); P$ | input |
| split M is $(x:T, y:U); P$ | pair splitting |
| case M is inl $(x:T) P$ is inr $(y:U) Q$ | union case |
| decrypt M is $\{x:T\}_N; P$ | decryption |
| check M is $N; P$ | name-check |
| new $(x:T); P$ | name generation |
| $P \mid Q$ | composition |
| repeat P | replication |
| stop | inactivity |

These processes are:

- Processes out $M N$ and inp $M (x:T); P$ are output and input, respectively, along an asynchronous, unordered channel M . If an output out $x N$ runs in parallel with an input inp $x (y); P$, the two can interact to leave the residual process $P\{y \leftarrow M\}$.
- A process split M is $(x:T, y:U); P$ splits the pair M into its two components. If M is (N, L) , the process behaves as $P\{x \leftarrow N\}\{y \leftarrow L\}$. Otherwise, it deadlocks, that is, does nothing.
- A process case M is inl $(x:T) P$ is inr $(y:U) Q$ checks the tagged union M . If M is inl (L) , the process behaves as $P\{x \leftarrow L\}$. If M is inr (N) it behaves as $Q\{y \leftarrow N\}$. Otherwise, it deadlocks.
- A process decrypt M is $\{x:T\}_N; P$ decrypts M using key N . If M is $\{L\}_N$, the process behaves as $P\{x \leftarrow L\}$. Otherwise, it deadlocks. We assume there is enough redundancy in the representation of ciphertexts to detect decryption failures.
- A process check M is $N; P$ checks the messages M and N are the same name before executing P . If the equality test fails, the process deadlocks.
- A process new $(x:T); P$ generates a new name x , whose scope is P , and then runs P .
- A process $P \mid Q$ runs processes P and Q in parallel.
- A process repeat P replicates P arbitrarily often. So repeat P behaves like $P \mid$ repeat P .
- The process stop is deadlocked.

Each binding occurrence of a name bears a type annotation. These types play a role in type-checking but have no role at runtime; they do not affect the operational behaviour of processes. In examples, for the sake of brevity, we sometimes omit type annotations.

Free names of a process $fn(P)$:

$$fn(\text{out } M N) \triangleq fn(M) \cup fn(N)$$

$$fn(\text{inp } M (x:T); P) \triangleq fn(M) \cup fn(T) \cup (fn(P) - \{x\})$$

$$\begin{aligned}
fn(\text{split } M \text{ is } (x:T, y:U); P) &\triangleq fn(M) \cup fn(T) \cup (fn(U) - \{x\}) \cup (fn(P) - \{x, y\}) \\
fn(\text{case } M \text{ is inl } (x:T) \ P \text{ is inr } (y:U) \ Q) &\triangleq fn(M) \cup fn(T) \cup (fn(P) - \{x\}) \cup \\
&\quad fn(U) \cup (fn(Q) - \{y\}) \\
fn(\text{decrypt } M \text{ is } \{x:T\}_N; P) &\triangleq fn(M) \cup fn(T) \cup fn(N) \cup (fn(P) - \{x\}) \\
fn(\text{check } M \text{ is } N; P) &\triangleq fn(M) \cup fn(N) \cup fn(P) \\
fn(\text{new } (x:T); P) &\triangleq fn(T) \cup (fn(P) - \{x\}) \\
fn(P \mid Q) &\triangleq fn(P) \cup fn(Q) \\
fn(\text{repeat } P) &\triangleq fn(P) \\
fn(\text{stop}) &\triangleq \emptyset
\end{aligned}$$

We write $P\{x \leftarrow N\}$ for the outcome of a capture-avoiding substitution of the message N for each free occurrence of the name x in the process P . We identify processes up to the consistent renaming of bound names, for example when $y \notin fn(P)$, we equate $\text{new } (x:T); P$ with $\text{new } (y:T); (P\{x \leftarrow y\})$. We will often elide stop from the end of processes, and we will write $\text{out } x \ M; P$ as shorthand for $\text{out } x \ M \mid P$.

2.2 Programming an Example

This section shows how to program a simple cryptographic protocol in spi. The protocol is intended to allow a fixed principal A to send a series of messages to another fixed principal B via a public channel, assuming they both share a secret key K . The idea is simply that A encrypts each message. Of course, for many purposes this protocol is actually far too simple: it is vulnerable to an attacker intercepting and replaying a message, so that B may accept the message twice though A sent it just once. In the next section, we introduce correspondence assertions to specify that B should accept a message M no more times than A sent M , and we discuss a standard guard against replay attacks, based on nonces.

In a common notation, we can summarise this flawed protocol as follows:

$$\text{Message 1 } A \rightarrow B : \{M\}_K$$

Although standard, this notation leaves implicit details of both protocol behaviour and security goals. One of the original purposes of the spi-calculus was to make protocol behaviour explicit in an executable format. We can program the protocol in spi as follows.

First, we describe the behaviour of the sender and receiver.

$$\begin{array}{ll}
\text{FlawedSender}(net, key) \triangleq & \text{FlawedReceiver}(net, key) \triangleq \\
\text{repeat} & \text{repeat} \\
\text{new } (msg); & \text{inp } net \ (c\text{text}); \\
\text{out } net \ \{msg\}_{key} & \text{decrypt } c\text{text} \text{ is } \{msg\}_{key}
\end{array}$$

These are:

- The process $FlawedSender(net, key)$ is the sender A , parameterized on net (the name of the public channel) and key (the shared secret key). It repeatedly generates a fresh name msg , and then sends the ciphertext $\{msg\}_{key}$ on the public net channel.

(In passing from the informal notation to the spi-calculus, we have determined that the plaintexts of the sent messages are freshly generated, rather than say being predetermined. It is easy to adapt this process to take a list of predetermined plaintexts as parameter.)

- The process $FlawedReceiver(net, key)$ is the receiver B , parameterized on net and key . It repeatedly receives a message on the public net channel, binds it to variable $ctxt$, and attempts to decrypt it with key key .

We specify the behaviour of the whole system running in the protocol by generating a fresh name key —the shared secret key—and then by placing the sender and receiver in parallel.

$$FlawedSystem(net, done) \triangleq \\ \text{new } (key); \\ (FlawedSender(net, key) \mid FlawedReceiver(net, key))$$

Most protocols analysed with the spi-calculus have been programmed in this style.

3 Specifying Protocols

Woo and Lam [41] introduce correspondence assertions, a method for specifying protocol authenticity properties, such as properties that are violated by replay or man-in-the-middle attacks. The method depends on principals asserting labelled begin- and end-events during the course of a protocol. The idea is that each end-event should correspond to a distinct, preceding begin-event with the same label. Otherwise there is an error in the protocol.

To formalize these ideas, in Section 3.1, we enrich our spi-calculus with assertions of begin- and end-events. Then, in Section 3.2, we illustrate how to specify an authenticity property of our example protocol, and show in fact that the protocol is flawed. In Section 3.3 we fix the flaw by adding a standard nonce handshake.

3.1 A Spi-Calculus with Correspondence Assertions

First, we introduce the following notation for events, using messages as labels.

Events:

| | |
|-----------|---------------------------------------|
| begin L | begin-event labelled with message L |
| end L | end-event labelled with message L |

Second, we add processes to assert begin- and end-events.

Processes:

| | |
|------------------|-------------------|
| $O, P, Q, R ::=$ | process |
| ... | as in Section 2.1 |
| begin $L; P$ | begin-assertion |
| end $L; P$ | end-assertion |

Assertions are autonomous in that they act independently without any synchronisation with other processes.

- The begin-assertion $\text{begin } L; P$ autonomously asserts a begin L event, and then behaves as P .
- The end-assertion $\text{end } L; P$ autonomously asserts an end L event, and then behaves as P .

Free names of a process $fn(P)$:

| |
|--|
| $fn(\text{begin } M; P) \triangleq fn(M) \cup fn(P)$ |
| $fn(\text{end } M) \triangleq fn(M)$ |

Given this informal semantics, we give an informal definition of process safety. (We formalize these definitions in Appendix B via a trace semantics for the spi-calculus.)

Safety:

| |
|---|
| A process P is <i>safe</i> if and only if for every run of the process and for every L , there is a distinct begin L event for every end L event. |
|---|

For example:

- Process $\text{begin } L; \text{end } L$ is safe.
- Process $\text{begin } L; \text{end } L; \text{end } L$ is unsafe because of the unmatched end L .
- Process $\text{begin } L; \text{begin } L; \text{end } L$ is safe; the unmatched begin L does not affect safety.
- Process $\text{begin } L; \text{begin } L; \text{end } L; \text{end } L$ is safe; here there are two correspondences, both named L .
- Process $\text{begin } L; \text{end } L; \text{begin } L'; \text{end } L'$ is safe.
- Process $\text{begin } L; \text{end } L'; \text{begin } L'; \text{end } L$ is unsafe.

Safety does not require begin- and end-assertions to be properly bracketed:

- Process $\text{begin } L; \text{begin } L'; \text{end } L'; \text{end } L$ is safe.
- Process $\text{begin } L; \text{begin } L'; \text{end } L; \text{end } L'$ is safe.

Finally, consider the parallel process $\text{begin } L \mid \text{end } L$. This process either asserts a $\text{begin } L$ event followed by an $\text{end } L$ event, or it asserts an $\text{end } L$ event followed by a $\text{begin } L$ event. Because of the latter run, the process is unsafe.

We are mainly concerned not just with safety, but with safety in the presence of an arbitrary hostile opponent, which we call robust safety. (This use of “robust” to describe a property invariant under composition with an arbitrary environment follows Grumberg and Long [20]). In the untyped spi-calculus [3], the opponent is modelled by an arbitrary process. In our typed spi-calculus, we do not consider completely arbitrary attacker processes, but restrict ourselves to *opponent* processes that satisfy two mild conditions:

- Opponents cannot assert events: otherwise, no process would be robustly safe, because of the opponent $\text{end } x$.
- Opponents are not required to be well-typed: we model this using a type Un for untyped, untrusted data. This is discussed further in Section 4

Opponents and Robust Safety:

| |
|---|
| <p>A process P is <i>assertion-free</i> if and only if it contains no begin- or end-assertions.</p> <p>A process P is <i>untyped</i> if and only if the only type occurring in P is Un.</p> <p>An <i>opponent</i> O is an assertion-free untyped process.</p> <p>A process P is <i>robustly safe</i> if and only if $P \mid O$ is safe for every opponent O.</p> |
|---|

3.2 Specifying the Example

Recall the protocol example of Section 2.2. Two fixed principals A and B share a key K with which A sends a sequence of messages to B . We introduce begin - and end -events labelled M for each message M . The sender asserts a begin -event labelled M before sending M , and the receiver asserts an end -event labelled M after successfully receiving a message M .

We express this idea informally as follows:

| | | |
|-----------|---------------------|-----------|
| Event 1 | A begins | M |
| Message 1 | $A \rightarrow B$: | $\{M\}_K$ |
| Event 2 | B ends | M |

We express the idea formally by inserting assertion processes into the spi-calculus de-

scriptions of the sender and receiver. We update our definitions as follows.

$$\begin{aligned}
 \text{CheckedSender}(net, key) &\triangleq & \text{CheckedReceiver}(net, key) &\triangleq \\
 \text{repeat} & & \text{repeat} & \\
 \text{new } (msg); & & \text{inp } net \ (ctext); & \\
 \text{begin } msg; & & \text{decrypt } ctext \text{ is } \{msg\}_{key}; & \\
 \text{out } net \ \{msg\}_{key} & & \text{end } msg & \\
 \\
 \text{CheckedSystem}(net) &\triangleq & & \\
 \text{new } (key); & & & \\
 (\text{CheckedSender}(net, key) \mid \text{CheckedReceiver}(net, key)) & & &
 \end{aligned}$$

Next, we precisely state the authenticity property we desire (but that is actually violated by the protocol).

Authenticity: The process $\text{CheckedSystem}(net)$ is robustly safe. (Breaks.)

If the protocol is safe, each $\text{end } msg$ has a distinct corresponding $\text{begin } msg$, and therefore B accepts each message no more times than A sent it. Moreover, if the protocol is robustly safe, no attacker can violate this property.

It is easy to prove that this protocol is safe, since the protocol itself never duplicates messages. Still, the protocol is not robustly safe since a suitable attacker can violate this safety property.

$$\begin{aligned}
 \text{Attacker}(net) &\triangleq \\
 \text{inp } net \ (ctext); \text{out } net \ (ctext); \text{out } net \ (ctext)
 \end{aligned}$$

Here is an unsafe run of the process $\text{CheckedSystem}(net) \mid \text{Attacker}(net)$. The sender $\text{CheckedSender}(net, key)$ generates a name msg , performs a single $\text{begin } msg$; event, and sends the ciphertext $\{msg\}_{key}$ on net . The attacker $\text{Attacker}(net)$ receives this message, and then sends two copies of on net . The receiver then receives one of these copies, successfully decrypts it, and asserts an $\text{end } msg$ event. So far so good. But now another iteration of the body of $\text{CheckedReceiver}(net, key)$ receives the second copy, successfully decrypts it, and asserts another $\text{end } msg$ event. Because of the second end-event is unmatched, the run breaks the authenticity property displayed above.

3.3 Fixing the Example

A standard countermeasure against replay attacks is to include a *nonce*, a randomly generated bit-string, in each ciphertext to ensure its uniqueness. The following variant of our protocol is now initiated by the receiver, who sends a new nonce N to the sender, to guard against replays of the encrypted form of the message M .

| | | |
|-----------|---------------------|--------------|
| Event 1 | A begins | M |
| Message 1 | $B \rightarrow A$: | N |
| Message 2 | $A \rightarrow B$: | $\{M, N\}_K$ |
| Event 2 | B ends | M |

In the spi-calculus, nonces are represented by names, and creation of fresh nonces by name generation. We program the revised protocol as follows:

$$\begin{array}{l}
 \text{FixedSender}(net, key) \triangleq \\
 \text{repeat} \\
 \text{inp } net \text{ (nonce);} \\
 \text{new (msg);} \\
 \text{begin msg;} \\
 \text{out } net \{msg, nonce\}_{key} \\
 \\
 \text{FixedReceiver}(net, key) \triangleq \\
 \text{repeat} \\
 \text{new (nonce);} \\
 \text{out } net \text{ nonce;} \\
 \text{inp } net \text{ (ciphertext);} \\
 \text{decrypt ciphertext} \\
 \text{is } \{msg, nonce'\}_{key}; \\
 \text{check nonce is nonce'}; \\
 \text{end msg}
 \end{array}$$

The process $\text{check nonce is nonce}'$; P checks that $nonce$ and $nonce'$ are the same name before executing P . For the sake of simplicity, in this example and others in the paper we omit error recovery code: upon receiving a ciphertext containing an unexpected nonce, an instance of the receiver just terminates. The whole system and its authenticity property are now:

$$\begin{array}{l}
 \text{FixedSystem}(net) \triangleq \\
 \text{new (key);} \\
 (\text{FixedSender}(net, key) \mid \text{FixedReceiver}(net, key))
 \end{array}$$

Authenticity: The process $\text{FixedSystem}(net)$ is robustly safe.

Given our modifications, this property is true. A direct proof is possible, but tricky, since we must quantify over all possible attackers. The original paper on the spi-calculus includes a verification via equational reasoning of a protocol similar to that embodied in $\text{FixedSystem}(net)$. The point of our type system, presented next, is to provide an efficient way of proving this specification, and others like it.

4 Typing Protocols

This section describes the heart of our method for analysing authenticity properties of protocols: a dependent type and effect system for statically verifying correspondence assertions by type-checking.

Section 4.1 and Section 4.2 explain informally how to type messages and how to ascribe effects to processes, respectively. We present the type and effect system formally in Section 4.3. Finally, in Section 4.4 we explain how to type the assertions in the example of the previous section.

4.1 Types for Messages

There is an objection in principle to a security analysis based on type-checking processes: it may be reasonable to assume that honest principals conform to typing rules, but it is imprudent to assume the same of the opponent. As previously discussed, our

general model of the opponent is any untyped, assertion-free process. The objection to a typed analysis is that we may miss attacks by ruling out processes that happen not to conform to our typing rules. On the internet, famously, nobody knows you're a dog. Likewise, nobody knows your code failed the type-checker.

To answer this objection, Abadi [1] introduces an *untrusted type* (which we call Un) for public messages, those exposed to the opponent. Every message and every opponent is typable if all their free variables are assigned the Un type. The type represents the unconstrained messages that an arbitrary process manipulates. Since any opponent can be typed in this trivial way we have not limited the power of opponents.

To illustrate this, here are some informal typing rules for messages and processes (for brevity, we elide some technical requirements on free names). Messages of the Un type may be output, input, paired, split apart, encrypted, and decrypted, with no constraints.

- If $M : \text{Un}$ and $N : \text{Un}$ then $\text{out } M N$ is well-typed.
- If $M : \text{Un}$ and P is well-typed then $\text{inp } M (x:\text{Un}); P$ is well-typed.
- If $M : \text{Un}$ and $N : \text{Un}$ then $(M, N) : \text{Un}$.
- If $M : \text{Un}$ and P is well-typed then $\text{split } M \text{ is } (x:\text{Un}, y:\text{Un}); P$ is well-typed.
- If $M : \text{Un}$ and $N : \text{Un}$ then $\{M\}_N : \text{Un}$.
- If $M : \text{Un}$ and $N : \text{Un}$ and P is well-typed then $\text{decrypt } M \text{ is } \{x:\text{Un}\}_N; P$ is well-typed.

When modelling protocols, we assume that all the names and messages exposed to the opponent—representing public data and channels—are of this type. Names and messages not publicly disclosed may be assigned other types, known as *trusted types*.

Messages of the trusted type $\text{Key}(T)$ are symmetric keys for encrypting messages of type T . When encrypting with a $\text{Key}(T)$, the plaintext must have type T , and the resulting ciphertext is given untrusted type. Using the rules above for Un , we can send and receive ciphertexts on untrusted channels. When decrypting with a $\text{Key}(T)$, if we succeed we know the plaintext must have been encrypted with the same key, and therefore our typing rules assign it type T .

- If $M : T$ and $N : \text{Key}(T)$ then $\{M\}_N : \text{Un}$.
- If $M : \text{Un}$ and $N : \text{Key}(T)$ and P is well-typed then $\text{decrypt } M \text{ is } \{x:T\}_N; P$ is well-typed.

The remaining trusted types are more standard. Messages of type $\text{Ch}(T)$ are channels communicating data of type T . Messages of type $(x:T, U)$ are dependent pairs where the first element has type T and the second element has type U . The variable x is bound, and has scope U . (The need for such dependent types arises later, when we introduce a type for nonces.) The only message of the empty tuple type $()$ is the empty tuple $()$. Messages of type $T + U$ are tagged unions. A union of type $T + U$ is either of the form $\text{inl } (M)$ where M has type T , or of the form $\text{inr } (N)$ where N has type U . As a technical convenience, to simplify some abbreviations introduced in Appendix A.5, we introduce the empty type, 0 . There are no messages of this type. Other base types

such as int or boolean could easily be added to this language: we expect they would produce no technical difficulties.

Types:

| $T, U ::=$ | type |
|--------------|---------------------|
| Un | untrusted type |
| Key(T) | shared-key type |
| Ch(T) | channel type |
| () | empty tuple type |
| ($x:T, U$) | dependent pair type |
| $T + U$ | variant type |
| 0 | empty type |

For example:

- Key(Un): key for encrypting untrusted data
- Ch(Un): channel for communicating untrusted data
- Key(Key(Un) + Ch(Un)): key for encrypting either a key for encrypting untrusted data or a channel for communicating untrusted data

4.2 Effects for Processes

Our effect system tracks the unmatched end-assertions of a process. In its most basic form, our main judgment

$$P : [\text{end } L_1, \dots, \text{end } L_n]$$

means that the effect $[\text{end } L_1, \dots, \text{end } L_n]$, is an upper bound on the multiset (or unordered list) of end-events that P may assert without asserting a matching begin-event. Hence, if $P : \square$ then every end-event in P has a matching begin-event, that is, P is safe.

Let e stand for an *atomic effect*. One kind of atomic effect is $\text{end } L$. The second kind is $\text{check } N$; we explain later its use to track nonce name-checking. Let es stand for an *effect*, that is, a multiset $[e_1, \dots, e_n]$ of atomic effects. We write $es + es'$ for the multiset union of the two multisets es and es' , that is, their concatenation. We write $es - es'$ for the multiset subtraction of es' from es , that is, the outcome of deleting an occurrence of each atomic effect in es' from es . If an atomic effect does not occur in an effect, then deleting the atomic effect leaves the effect unchanged.

Tracking Correspondences in Sequential Code

Given this notation, the typing rules for $\text{begin } L; P$ and $\text{end } L; P$ are essentially:

- If $P : es$ then $\text{begin } L; P : (es - [\text{end } L])$.
- If $P : es$ then $\text{end } L; P : (es + [\text{end } L])$.

These rules are enough to check correspondences in sequential code, for example:

- $\text{end } L : [\text{end } L]$

- $\text{begin } L; \text{end } L : []$
- $\text{end } L; \text{end } L : [\text{end } L, \text{end } L]$
- $\text{begin } L; \text{end } L; \text{end } L : [\text{end } L]$
- $\text{begin } L; \text{begin } L; \text{end } L; \text{end } L : []$

Transferring Effects between Parallel Processes

Our rules for assigning effects to communications and compositions are similar to those in previous work on effect systems for the π -calculus [11, 18].

- If $M : \text{Ch}(T)$ and $N : T$ then $\text{out } M N : []$.
- If $M : \text{Ch}(T)$ and $P : es$ then $\text{inp } M (x:T); P : es$.
- If $P : es_P$ and $Q : es_Q$ then $P \mid Q : (es_P + es_Q)$.

When computing the effect of the composition $P \mid Q$ of two processes, we simply compute the multiset union of the effects of the processes. This rule in itself does not allow a begin-assertion in P , say, to account for an end-assertion in Q . For example, the parallel composition $\text{begin } L \mid \text{end } L$ has effect $[\text{end } L]$, while in contrast the sequential composition $\text{begin } L; \text{end } L$ has effect $[\text{end } L]$. In the parallel case, we cannot assume that the begin-event precedes the end-event so we must conservatively assign the effect $[\text{end } L]$. In the sequential case, the syntax guarantees that the begin-event precedes the end-event so we can assign the effect $[\text{end } L]$. Somehow we need to be able to show that temporal precedences are established between parallel processes. Recall our *FixedSystem* example: we need to show that a distinct *begin msg* precedes each *end msg*, even though these assertions are running in parallel.

Typing Nonce Handshakes

A nonce handshake guarantees temporal precedence between events in parallel processes. In this paper, we consider a particular idiom for nonce handshakes, referred to by Guttman and Thayer as *incoming tests* [21]. Other idioms are possible, for example Guttman and Thayer's *outgoing tests*, but we leave these for future work. Incoming tests break down into several steps.

- (1) The receiver creates a fresh nonce and publishes it.
- (2) The sender embeds the nonce in a ciphertext.
- (3) The receiver looks for the nonce in a received ciphertext. Finding the nonce encrypted under a shared private key proves the sender recently generated the ciphertext. If this is the first and only time the nonce is found, there is a one-to-one correspondence between finding the nonce and the creation of the ciphertext by the sender.
- (4) To avoid vulnerability to replay of messages containing the nonce, the receiver subsequently discards the nonce and no longer looks for it.

We type-check these four steps as follows.

- (1) The receiver creates the nonce N in the untrusted type Un . This allows the nonce to be sent on an untrusted channel, and reflects that it can be received and copied by the opponent as well as the sender.
- (2) The sender embeds the nonce in a ciphertext as a message of a new trusted type $\text{Nonce } es$, where es is an effect. The sender casts the nonce $N : \text{Un}$ to this trusted type using the new process cast N is $(x:\text{Nonce } es);P$. At runtime, this process simply binds the message N to the variable x of type $\text{Nonce } es$, and then runs P . The sender uses the variable x to embed the nonce in the ciphertext.
- (3) After decrypting a ciphertext containing a nonce $N' : \text{Nonce } es$, the receiver uses a name-check $\text{check } N$ is $N';Q$ to check for the nonce $N : \text{Un}$ which it made public earlier. Only a cast can populate the type $\text{Nonce } es$. So the presence of the message $N' : \text{Nonce } es$ proves there was a preceding execution of a cast process. Our type system ensures that at most one name-check process checks for the presence of each nonce $N : \text{Un}$. Therefore, if the check succeeds, we are guaranteed a one-to-one correspondence between the check and the preceding process that cast N into the type $\text{Nonce } es$. Note that the safety of this step relies on global agreement between the trusted participants as to the types of each of the messages.
- (4) To guarantee that each nonce N is the subject of no more than one name-check, we introduce a new atomic effect, written $\text{check } N$. In general, our main judgment takes the form,

$$P : [\text{end } L_1, \dots, \text{end } L_m, \text{check } N_1, \dots, \text{check } N_n]$$

and means the multiset $[\text{end } L_1, \dots, \text{end } L_m]$ is an upper bound on the end-events P asserts without previously asserting a corresponding begin-event, and that the multiset $[\text{check } N_1, \dots, \text{check } N_n]$ is an upper bound on the multiset of free nonces name-checked by P . We include $\text{check } N$ in the effect of a name-check $\text{check } N$ is $N';Q$ on a nonce N . When checking name generation $\text{new } (N:\text{Un});P$, we check that $\text{check } N$ occurs at most once in the effect of P . This guarantees that each free name is the subject of no more than one name-check.

In summary, our type and effect system provides a solution to the problem of guaranteeing temporal precedences between parallel processes: for every successful execution of a process $\text{check } N$ is $N';Q$, where $N' : \text{Nonce } es$, there is a distinct preceding execution of a process cast N is $(x:\text{Nonce } es);P$, even if the name-check and the cast are in parallel processes.

The following rules for computing the effect of casts and name-checks exploit this temporal precedence. They allow us to guarantee by typing that those end-events following the name-check and listed in the effect es of the type $\text{Nonce } es$ are matched by distinct begin-events that precede the cast. This effect is transferred from the name-check to the cast; the effect es is added to the effect of a cast, and is subtracted from the effect of a name-check.

- If $N : \text{Un}$ and $P : es_P$ then cast N is $(x:\text{Nonce } es);P : (es_P + es)$.

- If $N : \text{Un}$ and $N' : \text{Nonce } es$ and $Q : es_Q$
then check N is N' ; $Q : ((es_Q - es) + [\text{check } N])$.
- If $P : es_P$ then new (N) ; $P : (es_P - [\text{check } N])$.

To illustrate these rules, we compute the effect of a nonce handshake that guarantees the safety of a correspondence between a begin-event labelled m in one process and an end-event with the same label in another. We consider fixed, global names m , n , and c . We assume $m:T$ for some type T . We assume $n:\text{Un}$ is the name of a nonce that somehow is already shared between the two processes. We assume $c:\text{Ch}((\text{Nonce } [\text{end } m]))$ is the name of a trusted channel shared by the processes. (To focus on casting and checking nonces, we communicate the nonce n over the trusted channel c ; in realistic examples, nonces are sent encrypted on untrusted channels.)

The first process

$$P = \begin{array}{l} \text{begin } m; \\ \text{cast } n \text{ is } (n':\text{Nonce } [\text{end } m]); \\ \text{out } c \ n' \end{array}$$

begins the correspondence, casts n into the type $\text{Nonce } [\text{end } m]$, and then sends it on c .

We have $\text{cast } n \text{ is } (n':\text{Nonce } [\text{end } m]); \text{out } c \ n' : [\text{end } m]$ and therefore $P : []$. The second process

$$Q = \begin{array}{l} \text{inp } c \ (x:\text{Nonce } [\text{end } m]); \\ \text{check } n \text{ is } x; \\ \text{end } m \end{array}$$

receives a name x off the channel c , checks that n equals x , and if so ends the correspondence.

We have $\text{end } m : [\text{end } m]$, and $\text{check } n \text{ is } x; \text{end } m : ([\text{end } m] - [\text{end } m]) + [\text{check } n]$, and therefore $Q : [\text{check } n]$. Now, by the rules for name generation and composition, we get that $R = \text{new } (n:\text{Un}); (P \mid Q) : []$. So R is safe.

On the other hand, consider the process $R' = \text{new } (n:\text{Un}); (P' \mid Q \mid Q)$ where we have duplicated Q and where the process

$$P' = \begin{array}{l} \text{begin } m; \\ \text{cast } n \text{ is } (n':\text{Nonce } [\text{end } m]); \\ (\text{out } c \ n' \mid \text{out } c \ n') \end{array}$$

is a variation of P' that duplicates the nonce. Now, R' is unsafe, because the two copies of Q can each receive one of the duplicate nonces sent by P' . Therefore both can assert an end-event, but only one is accounted for by the begin-assertion by P' . The process R' does not type-check, because it name-checks the nonce n more than once. We can derive $P' : []$, but the whole process R' fails the rule for name generation, because process $P' \mid Q \mid Q$ has effect $[\text{check } n, \text{check } n]$ so the condition $n \notin \text{fn}([\text{check } n, \text{check } n] - [\text{check } n])$ is false.

Effects and Atomic Effects

Given these motivations for and examples of assigning effects to processes, here is the grammar of effects and atomic effects.

Effects:

| | |
|---------------------|-------------------------------------|
| $e, f ::=$ | atomic effect |
| end L | end-event labelled with message L |
| check N | name-check for a nonce N |
| $es, fs ::=$ | effect |
| $[e_1, \dots, e_n]$ | multiset of atomic effects |

Free names, $fn(es)$, of an effect es :

| |
|--|
| $fn(\text{end } L) \triangleq fn(L)$ |
| $fn(\text{check } N) \triangleq fn(N)$ |
| $fn([e_1, \dots, e_n]) \triangleq fn(e_1) \cup \dots \cup fn(e_n)$ |

We write $es\{x \leftarrow M\}$ for the outcome of a capture-avoiding substitution of the message M for each free occurrence of the name x in the effect es .

Additional Types and Processes

We end this section by completing the grammars of types and processes with the new type and new processes we need for typing nonce handshakes. We add a type for nonces, and we give rules defining the set $fn(T)$ of any type T .

Types:

| | |
|------------|-------------------|
| $T, U ::=$ | type |
| ... | as in Section 4.1 |
| Nonce es | nonce type |

Free names, $fn(T)$, of a type T :

| |
|--|
| $fn(\text{Ch}(T)) \triangleq fn(T)$ |
| $fn(x:T, U) \triangleq fn(T) \cup (fn(U) - \{x\})$ |
| $fn() \triangleq \emptyset$ |
| $fn(T + U) \triangleq fn(T) \cup fn(U)$ |
| $fn(0) \triangleq \emptyset$ |
| $fn(\text{Un}) \triangleq \emptyset$ |
| $fn(\text{Key}(T)) \triangleq fn(T)$ |
| $fn(\text{Nonce } es) \triangleq fn(es)$ |

We write $T\{x \leftarrow M\}$ for the outcome of a capture-avoiding substitution of the message M for each free occurrence of the name x in the type T .

As we explained, we add a process to cast untrusted data into nonce type. Moreover, we add a new process for pattern matching pairs.

Processes:

| | |
|----------------------------|----------------------------|
| $O, P, Q, R ::=$ | process |
| ... | as in Sections 2.1 and 3.1 |
| cast M is $(x:T); P$ | cast to nonce type |
| match M is $(N, y:U); P$ | pair pattern matching |

In a process cast M is $(x:T); P$, the name x is bound; its scope is the process P . In a process match M is $(N, y:U); P$, the name y is bound; its scope of the process P .

- The process cast M is $(x:T); P$ casts the message M to the type T , by binding the variable x to M , and then running P . (This process can only be typed by our type system if T is of the form Nonce *es*.)
- The process match M is $(N, y); P$ is similar to split M is $(x, y); P$ except that it checks that the first component of M is equal to N before extracting the second component (which is bound to y in P). If the equality test fails, then the process deadlocks.

Pair pattern matching is a generalization of π -calculus name equality testing, since $[M = N]P$ can be written match $(M, ())$ is $(N, y); P$.

Free names of a process $fn(P)$:

| | |
|--|--|
| $fn(\text{cast } M \text{ is } (x:T); P) \triangleq$ | $fn(M) \cup fn(T) \cup (fn(P) - \{x\})$ |
| $fn(\text{match } M \text{ is } (N, y:U); P) \triangleq$ | $fn(M) \cup fn(N) \cup fn(U) \cup (fn(P) - \{y\})$ |

Pair pattern matching is used in the protocol examples in Appendix A.

4.3 Typing Rules

In this section, we formally define the judgments of our type and effect system.

These judgments all depend on an *environment*, E , that defines the types of all variables in scope. An environment takes the form $x_1:T_1, \dots, x_n:T_n$ and defines the type T_i for each variable x_i . The *domain*, $dom(E)$, of an environment E is the set of variables whose types it defines.

Environments:

| | |
|---|--------------------------|
| $D, E ::=$ | environment |
| \emptyset | empty |
| $E, x:T$ | entry |
| $dom(x_1:T_1, \dots, x_n:T_n) \triangleq$ | domain of an environment |
| $\{x_1, \dots, x_n\}$ | |

The following are the five judgments of our type and effect system. They are inductively defined by rules presented in the following tables.

Judgments $E \vdash j$:

| | |
|---------------------|-----------------------------------|
| $E \vdash \diamond$ | good environment |
| $E \vdash es$ | good effect es |
| $E \vdash T$ | good type T |
| $E \vdash M : T$ | good message M of type T |
| $E \vdash P : es$ | good process P with effect es |

Rules for Environments:

| | |
|--------------------------------------|--|
| (Env \emptyset) | (Env x) (where $x \notin \text{dom}(E)$) |
| $\frac{}{\emptyset \vdash \diamond}$ | $\frac{E \vdash T}{E, x:T \vdash \diamond}$ |

These standard rules define an environment $x_1:T_1, \dots, x_n:T_n$ to be well-formed just if each of the names x_1, \dots, x_n are distinct, and each of the types T_i is well-formed.

Rules for Effects:

| | | |
|--|--|--|
| (Effect \emptyset) | (Effect End) | (Effect Check) |
| $\frac{E \vdash \diamond}{E \vdash \emptyset}$ | $\frac{E \vdash es \quad E \vdash L : T}{E \vdash es + [\text{end } L]}$ | $\frac{E \vdash es \quad E \vdash N : \text{Un}}{E \vdash es + [\text{check } N]}$ |

These rules define an effect $[e_1, \dots, e_n]$ to be well-formed just if for each atomic effect $e_i = \text{end } L$, message L has type T for some type T , and for each atomic effect $e_i = \text{check } N$, message N has type Un .

Rules for Types:

| | | | |
|--|--|---|---|
| (Type Un) | (Type Chan) | (Type Pair) | (Type Unit) |
| $\frac{E \vdash \diamond}{E \vdash \text{Un}}$ | $\frac{E \vdash T}{E \vdash \text{Ch}(T)}$ | $\frac{E, x:T \vdash U}{E \vdash (x:T, U)}$ | $\frac{E \vdash \diamond}{E \vdash ()}$ |
| (Type Variant) | (Type Empty) | (Type Key) | (Type Nonce) |
| $\frac{E \vdash T \quad E \vdash U}{E \vdash T + U}$ | $\frac{E \vdash \diamond}{E \vdash 0}$ | $\frac{E \vdash T}{E \vdash \text{Key}(T)}$ | $\frac{E \vdash es}{E \vdash \text{Nonce } es}$ |

According to these rules a type is well-formed just if every effect occurring in the type is itself well-formed.

Next, we present the rules for deriving the judgment $E \vdash M : T$ that assigns a type T to a message M . We split the rules into three tables: first, the rule for variables; second, rules for manipulating data of trusted type; and third, rules for assigning the untrusted type to arbitrary messages.

Rule for Variables:

$$\frac{\text{(Msg } x)}{E', x:T, E'' \vdash \diamond} \\ \hline E', x:T, E'' \vdash x : T$$

Rules for Messages of Trusted Type:

$$\frac{\text{(Msg Pair)} \quad E \vdash M : T \quad E \vdash N : U \{x \leftarrow M\}}{E \vdash (M, N) : (x:T, U)} \quad \frac{\text{(Msg Unit)} \quad E \vdash \diamond}{E \vdash () : ()}$$

$$\frac{\text{(Msg Inl)} \quad E \vdash M : T \quad E \vdash U}{E \vdash \text{inl}(M) : T + U} \quad \frac{\text{(Msg Inr)} \quad E \vdash T \quad E \vdash N : U}{E \vdash \text{inr}(N) : T + U}$$

$$\frac{\text{(Msg Encrypt)} \quad E \vdash M : T \quad E \vdash N : \text{Key}(T)}{E \vdash \{M\}_N : \text{Un}}$$

Rules for Messages of Untrusted Type:

$$\frac{\text{(Msg Pair Un)} \quad E \vdash M : \text{Un} \quad E \vdash N : \text{Un}}{E \vdash (M, N) : \text{Un}} \quad \frac{\text{(Msg Unit Un)} \quad E \vdash \diamond}{E \vdash () : \text{Un}}$$

$$\frac{\text{(Msg Inl Un)} \quad E \vdash M : \text{Un}}{E \vdash \text{inl}(M) : \text{Un}} \quad \frac{\text{(Msg Inr Un)} \quad E \vdash N : \text{Un}}{E \vdash \text{inr}(N) : \text{Un}}$$

$$\frac{\text{(Msg Encrypt Un)} \quad E \vdash M : \text{Un} \quad E \vdash N : \text{Un}}{E \vdash \{M\}_N : \text{Un}}$$

Recall from Section 4.1 the principle that any message can be assigned the untrusted type Un , provided its free variables are also untrusted. Using just the rules in the first and third tables of message typing rules, we can prove:

Lemma 1 *If $\text{fn}(M) \subseteq \{x_1, \dots, x_n\}$ then $x_1:\text{Un}, \dots, x_n:\text{Un} \vdash M : \text{Un}$.*

Proof By structural induction on the message M . □

A message may be assigned both a trusted and an untrusted type. For example:

- $x:\text{Un}, y:\text{Un} \vdash (x, y):(z:\text{Un}, \text{Un})$ by (Msg Pair)
- $x:\text{Un}, y:\text{Un} \vdash (x, y):\text{Un}$ by (Msg Pair Un)

Finally, we present the rules for assigning effects to processes. To state the rule for name-generation we introduce the notion of a *generative type*. A type is generative if it is untrusted or if it is a key or channel type. A process $\text{new } (x:T);P$ is only well-typed if T is generative. This rule prevents the fresh generation of names of, for example, the *Nonce* es type; it is crucial to our system that the only way of populating this type is via a cast process.

Generative Types:

A type is *generative* if and only if
it takes the form $\text{Ch}(T)$, Un , or $\text{Key}(T)$.

Basic Rules for Processes:

| | |
|---|---|
| <p>(Proc Begin)</p> $\frac{E \vdash L : T \quad E \vdash P : es}{E \vdash \text{begin } L; P : es - [\text{end } L]}$ | <p>(Proc End)</p> $\frac{E \vdash L : T \quad E \vdash P : es}{E \vdash \text{end } L; P : es + [\text{end } L]}$ |
| <p>(Proc Par)</p> $\frac{E \vdash P : es \quad E \vdash Q : fs}{E \vdash P \mid Q : es + fs}$ | <p>(Proc Repeat)</p> $\frac{E \vdash P : []}{E \vdash \text{repeat } P : []}$ |
| <p>(Proc Stop)</p> $\frac{E \vdash \diamond}{E \vdash \text{stop} : []}$ | <p>(Proc Res) (where $x \notin \text{fn}(es - [\text{check } x])$)</p> $\frac{E, x:T \vdash P : es \quad T \text{ is generative}}{E \vdash \text{new } (x:T); P : es - [\text{check } x]}$ |
| <p>(Proc Subsum)</p> $\frac{E \vdash P : es \quad E \vdash es'}{E \vdash P : es + es'}$ | |

We discussed informal versions of the rules (Proc Begin), (Proc End), (Proc Par), and (Proc Res) previously. The rule (Proc Repeat) requires the effect of the replicated process P to be empty. If P had a non-empty effect, then somehow we might assign an infinite effect to repeat P but this would not be useful. Assigning an effect to a whole process is useful because if the effect is empty then the process is safe. Any process enclosing repeat P can only match a finite number of atomic effects arising from repeat P , and so must have a non-empty effect. So typing repeat P is only useful if P has an empty effect. The rule (Proc Stop) says the inactive process has empty effect. The effect of a process is an upper bound on the behaviour of a process; the rule (Proc Subsum) allows us to weaken this upper bound by enlarging the effect.

The rule (Proc Case), in the following table, uses an operator \vee defined as follows. Let the multiset ordering $es \leq es'$ mean there is an effect es'' such that $es + es'' = es'$. Then we write $es \vee es'$ for the least effect es'' in this ordering such that both $es \leq es''$ and $es' \leq es''$. Note that $(es \vee es') = ((es - es') + es')$.

Rules for Processes Manipulating Trusted Types:

(Proc Output)

$$\frac{E \vdash x : \text{Ch}(T) \quad E \vdash M : T}{E \vdash \text{out } x \ M : []}$$

(Proc Input) (where $y \notin \text{fn}(es)$)

$$\frac{E \vdash x : \text{Ch}(T) \quad E, y : T \vdash P : es}{E \vdash \text{inp } x (y : T); P : es}$$

(Proc Split) (where $x \notin \text{fn}(es)$ and $y \notin \text{fn}(es)$)

$$\frac{E \vdash M : (x : T, U) \quad E, x : T, y : U \vdash P : es}{E \vdash \text{split } M \text{ is } (x : T, y : U); P : es}$$

(Proc Match) (where $y \notin \text{fn}(es)$)

$$\frac{E \vdash M : (x : T, U) \quad E \vdash N : T \quad E, y : U \{x \leftarrow N\} \vdash P : es}{E \vdash \text{match } M \text{ is } (N, y : U \{x \leftarrow N\}); P : es}$$

(Proc Case) (where $x \notin \text{fn}(es)$ and $y \notin \text{fn}(fs)$)

$$\frac{E \vdash M : T + U \quad E, x : T \vdash P : es \quad E, y : U \vdash Q : fs}{E \vdash \text{case } M \text{ is inl } (x : T) \ P \text{ is inr } (y : U) \ Q : es \vee fs}$$

(Proc Decrypt) (where $x \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E \vdash y : \text{Key}(T) \quad E, x : T \vdash P : es}{E \vdash \text{decrypt } M \text{ is } \{x : T\}_y; P : es}$$

(Proc Cast) (where $x \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E, x : \text{Nonce } fs \vdash P : es}{E \vdash \text{cast } M \text{ is } (x : \text{Nonce } fs); P : es + fs}$$

(Proc Check)

$$\frac{E \vdash M : \text{Un} \quad E \vdash N : \text{Nonce } fs \quad E \vdash P : es}{E \vdash \text{check } M \text{ is } N; P : (es - fs) + [\text{check } M]}$$

We discussed informal versions of the rules (Proc Input), (Proc Output), (Proc Cast), and (Proc Check) previously. Rule (Proc Split) is a standard rule to allow a pair $M : (x : T, U)$ to be split into two components named $x : T$ and $y : U$, where x may occur free in the type U . The conditions $x \notin \text{fn}(es)$ and $y \notin \text{fn}(es)$ prevent the bound variables x and y from appearing out of scope in the effect es . In the rule (Proc Match), the message $N : T$ is meant to match the first component of the pair $M : (x : T, U)$, and the variable $y : U$ gets bound to the second component. Again, the condition $y \notin \text{fn}(es)$ prevents y from appearing out of scope in es . The rule (Proc Case) is a standard rule for checking inspections of tagged unions. In the rule (Proc Decrypt), the ciphertext M is of untrusted type, Un , the key y is of type $\text{Key}(T)$, and the plaintext, bound to x , has type T . The condition $x \notin \text{fn}(es)$ prevents x from appearing out of scope in the effect es .

Rules for Processes Manipulating Untrusted Types:

(Proc Output Un)

$$\frac{E \vdash M : \text{Un} \quad E \vdash N : \text{Un}}{E \vdash \text{out } M N : []}$$

(Proc Input Un) (where $y \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E, y:\text{Un} \vdash P : es}{E \vdash \text{inp } M (y:\text{Un}); P : es}$$

(Proc Split Un) (where $x \notin \text{fn}(es)$ and $y \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E, x:\text{Un}, y:\text{Un} \vdash P : es}{E \vdash \text{split } M \text{ is } (x:\text{Un}, y:\text{Un}); P : es}$$

(Proc Match Un) (where $y \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E \vdash N : \text{Un} \quad E, y:\text{Un} \vdash P : es}{E \vdash \text{match } M \text{ is } (N, y:\text{Un}); P : es}$$

(Proc Case Un) (where $x \notin \text{fn}(es)$ and $y \notin \text{fn}(fs)$)

$$\frac{E \vdash M : \text{Un} \quad E, x:\text{Un} \vdash P : es \quad E, y:\text{Un} \vdash Q : fs}{E \vdash \text{case } M \text{ is inl } (x:\text{Un}) P \text{ is inr } (y:\text{Un}) Q : es \vee fs}$$

(Proc Decrypt Un) (where $x \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E \vdash N : \text{Un} \quad E, x:\text{Un} \vdash P : es}{E \vdash \text{decrypt } M \text{ is } \{x:\text{Un}\}_N; P : es}$$

(Proc Cast Un) (where $x \notin \text{fn}(es)$)

$$\frac{E \vdash M : \text{Un} \quad E, x:\text{Un} \vdash P : es}{E \vdash \text{cast } M \text{ is } (x:\text{Un}); P : es}$$

(Proc Check Un)

$$\frac{E \vdash M : \text{Un} \quad E \vdash N : \text{Un} \quad E \vdash P : es}{E \vdash \text{check } M \text{ is } N; P : es}$$

These rules are similar to those in the previous table in how they compute effects of processes, but differ in that all messages are of untrusted type. These rules are needed to type-check opponents.

Our rules for processes conform to the principle, stated in Section 4.1, that any opponent can be typed if all its free variables are assigned the type Un.

Lemma 2 (Opponent Typability) *If O is an opponent, that is, an untyped, assertion-free process, and $\text{fn}(O) \subseteq \{x_1, \dots, x_n\}$ then $x_1:\text{Un}, \dots, x_n:\text{Un} \vdash O : []$.*

Proof By structural induction on O , with appeal to Lemma 1. \square

The following theorem, proved in Appendix B, says a process is safe if it can be assigned the empty effect.

Theorem 1 (Safety) *If $E \vdash P : []$ then P is safe.*

Combined, Lemma 2 (Opponent Typability) and Theorem 1 (Safety) establish our main result, that our type and effect system guarantees robust safety.

Theorem 2 (Robust Safety) *If $x_1:\text{Un}, \dots, x_n:\text{Un} \vdash P : []$ then P is robustly safe.*

Proof For any untyped, assertion-free O , find x_{n+1}, \dots, x_{n+m} such that $\text{fn}(O) \subseteq \{x_1, \dots, x_{n+m}\}$. By Lemma 2 (Opponent Typability), we have $x_1 : \text{Un}, \dots, x_{n+m} : \text{Un} \vdash O : []$. By a standard weakening lemma, proved in the full version, $x_1:\text{Un}, \dots, x_n:\text{Un} \vdash P : []$ implies $x_1:\text{Un}, \dots, x_{n+m}:\text{Un} \vdash P : []$. So by rule (Proc Par) we have $x_1 : \text{Un}, \dots, x_{n+m} : \text{Un} \vdash P \mid O : []$, and so by Theorem 1 (Safety), $P \mid O$ is safe. Thus, P is robustly safe. \square

4.4 Typing the Example

Our example *FixedSystem*(*net*) from Section 3.3 uses a nonce handshake over the public channel *net* to transfer messages from the sender to the receiver. Here we show how to prove the example's correspondence assertions by choosing suitable types and adding a cast process.

The sender receives a nonce *nonce* off the public channel *net*, performs a begin-event to indicate it is sending a message *msg*, embeds the nonce and the message in a ciphertext encrypted with the shared key *key*, and returns the ciphertext to the receiver on *net*. Any public channel should be accessible to the opponent, so we assign *net* the untrusted type Un , and since *nonce* is sent on these channels, they too must have the untrusted type. We fix some arbitrary type *Msg* and assume each *msg* is of this type. To type-check the correspondence between begin- and end-assertions made by the sender and receiver, respectively, we add a cast process to the sender to cast the nonce into the type $\text{Nonce} [\text{end } msg]$. Therefore, the shared key has type $\text{Key}(msg:\text{Msg}, nonce:\text{Nonce} [\text{end } msg])$; the first component of the ciphertext is the actual message, and the second component is a nonce proving it is safe to assert an end *msg* event.

Therefore, we introduce the types

$$\begin{aligned} &Msg \text{ some arbitrary type} \\ &Network \triangleq \text{Un} \\ &MyNonce (msg) \triangleq \text{Nonce} [\text{end } msg] \\ &MyKey \triangleq \text{Key}(msg:\text{Msg}, nonce:MyNonce (msg)) \end{aligned}$$

and we type the sender as follows, where we display the effects of bracketed subprocesses to the right.

$$\begin{aligned} &TypedSender(net:Network, key:MyKey) : [] \triangleq \\ &\text{repeat} \\ &\quad \text{inp } net (nonce:\text{Un}); \\ &\quad \text{new } (msg:\text{Msg}); \\ &\quad \text{begin } msg; \\ &\quad \text{cast } nonce \\ &\quad \quad \left. \begin{array}{l} \text{is } (nonce':MyNonce (msg)); \\ \text{out } net \{msg, nonce'\}_{key} \end{array} \right\} [\text{end } msg] \quad \left. \right\} [] \end{aligned}$$

Next, we type the receiver. Like the sender, it is effect-free, that is, it can be assigned the empty effect.

$$\begin{array}{l} \text{TypedReceiver}(net:Network, key:MyKey) : [] \triangleq \\ \text{repeat} \\ \quad \text{new } (nonce:Un); \\ \quad \text{out } net \text{ } nonce; \\ \quad \text{inp } net \text{ } (c\text{text}:Un); \\ \quad \text{decrypt } c\text{text} \\ \quad \text{is } \{msg:Msg, nonce':MyNonce (msg)\}_{key}; \\ \quad \text{check } nonce \text{ is } nonce'; \\ \quad \text{end } msg \} [\text{end } msg] \} [\text{check } nonce] \end{array} \quad \left. \vphantom{\begin{array}{l} \text{TypedReceiver}(net:Network, key:MyKey) : [] \triangleq \\ \text{repeat} \\ \quad \text{new } (nonce:Un); \\ \quad \text{out } net \text{ } nonce; \\ \quad \text{inp } net \text{ } (c\text{text}:Un); \\ \quad \text{decrypt } c\text{text} \\ \quad \text{is } \{msg:Msg, nonce':MyNonce (msg)\}_{key}; \\ \quad \text{check } nonce \text{ is } nonce'; \\ \quad \text{end } msg \} [\text{end } msg] \} [\text{check } nonce] \end{array}} \right\} []$$

Since the sender and receiver are both effect-free, the whole system is also effect-free:

$$\begin{array}{l} \text{TypedSystem}(net:Network) : [] \triangleq \\ \text{new } (key:MyKey); \\ (\text{TypedSender}(net, key) \mid \text{TypedReceiver}(net, key)) \end{array}$$

By Theorem 2 (Robust Safety), it follows that $\text{TypedSystem}(net:Network)$ is robustly safe. This proves the following authenticity property by typing.

Authenticity: The process $\text{TypedSystem}(net)$ is robustly safe.

5 Further Protocol Examples

We have applied our method to several cryptographic protocols from the literature. We verified some protocols, found flaws in others, but also found at least one incompleteness in our method. Details are in an appendix, but we can summarise our experience as follows.

- Abadi and Gordon [3] propose a nonce-based variation of the Wide Mouth Frog key-exchange protocol [9]. We can verify authenticity properties of Abadi and Gordon's protocol by typing. Abadi and Gordon prove an equationally-specified authenticity property by constructing a bisimulation relation based on an elaborate invariant; our proof of correspondence assertions by typing took considerably less time.
- Woo and Lam [40] propose a nonce-based authentication protocol. Trying to type-check the protocol exposes known flaws in the protocol and suggests a known simplification [4, 5].
- Otway and Rees [33] propose another nonce-based key exchange protocol. The nonces used by the protocol to prove freshness are kept secret; hence the protocol does not fit the idiom that can be checked by our type system. Still, we can type-check a more efficient version of the protocol suggested by Abadi and Needham [4]. The typing suggests a further simplification.

In each case, there is a spi-calculus representation of the protocol in which there are arbitrarily many participant principals and arbitrarily many sessions.

6 Summary and Conclusion

To summarise, we reviewed the spi-calculus, a formalism for precisely describing the behaviour of security protocols based on cryptography. We embedded Woo and Lam's correspondence assertions in spi as a way of specifying authenticity properties. We devised a new type and effect system that proves authenticity properties, simply by type-checking.

To conclude, the examples in this paper, together with others we have investigated, suggest that this is a promising technique for checking protocols, since it requires little human effort to type a protocol, and the types of protocol data document how the protocol works.

Acknowledgements

Thanks to Martín Abadi, Gavin Lowe, Dusko Pavlovic, Simon Peyton Jones, Benjamin Pierce, Corin Pitcher, James Riely, and Andre Scedrov for discussions about this work. The anonymous referees for the *IEEE Computer Security Foundations Workshop* provided invaluable feedback. C.A.R. Hoare suggested several improvements to a draft. Alan Jeffrey was supported in part by Microsoft Research during some of the time we worked on this paper.

A Protocol Examples

In this appendix we describe details of the examples mentioned in Section 5. Section A.1 describes Abadi and Gordon’s version of Wide Mouth Frog. Section A.2 discusses Woo and Lam’s authentication protocol. Section A.3 discusses Otway and Rees’s key-exchange protocol. Finally, we present a new typed protocol for secure message streams in Section A.4.

Abbreviations Used in Examples

In these examples, we shall make use of the following syntax sugar:

- Dependent record types $(x_1:T_1, \dots, x_n:T_n)$, rather than just pairs. These come with a constructor (M_1, \dots, M_n) and a destructor match M is $(x_1:T_1, \dots, x_n:T_n);P$.
- Tagged union types $(\ell_1(T_1) \mid \dots \mid \ell_n(T_n))$ rather than just binary choice $T + U$. These come with a constructor $\ell_i(M)$ and a destructor match M is $\ell_i(x:T);P$.
- Dependent function types $(x:T) \rightarrow U$. These come with an appropriate function declaration and application syntax.

We show in Section A.5 that these constructs can be derived from our base language.

A.1 Abadi and Gordon’s Variant of Wide Mouth Frog

The original paper on the spi-calculus [3] includes a lengthy proof of authenticity and secrecy properties for a variation of the Wide Mouth Frog key distribution protocol [9] based on nonce handshakes instead of timestamps. In this section, we show how to type-check this protocol.

To begin with we look at an unsafe version of the protocol, to illustrate how attempting to type-check a protocol may expose flaws. This broken protocol consists of a sender (Alice), a receiver (Bob) and a server (Sam). Alice wishes to contact Bob, and asks Sam to establish her credentials:

| | | |
|-----------|-------------------|----------------------------------|
| Event 1 | A begins | “ A sending B key K_{AB} ” |
| Message 1 | $A \rightarrow S$ | A |
| Message 2 | $S \rightarrow A$ | N_S |
| Message 3 | $A \rightarrow S$ | $A, \{B, K_{AB}, N_S\}_{K_{AS}}$ |
| Message 4 | $S \rightarrow B$ | $()$ |
| Message 5 | $B \rightarrow S$ | N_B |
| Message 6 | $S \rightarrow B$ | $\{A, K_{AB}, N_B\}_{K_{BS}}$ |
| Event 2 | B ends | “ A sending B key K_{AB} ” |

(For the sake of readability, we use “ A sending B key K_{AB} ” as a shorthand for the message (A, B, K_{AB}) .)

This protocol can be compromised by an intruder I impersonating Sam, if Alice

acts both as a sender and a receiver:

| | | |
|--------------------|-------------------|----------------------------------|
| Event $\alpha.1$ | A begins | “ A sending B key K_{AB} ” |
| Message $\alpha.1$ | $A \rightarrow I$ | A |
| Message $\beta.4$ | $I \rightarrow A$ | $()$ |
| Message $\beta.5$ | $A \rightarrow I$ | N_A |
| Message $\alpha.2$ | $I \rightarrow A$ | N_A |
| Message $\alpha.3$ | $A \rightarrow I$ | $A, \{B, K_{AB}, N_A\}_{K_{AS}}$ |
| Message $\beta.6$ | $I \rightarrow A$ | $\{B, K_{AB}, N_A\}_{K_{AS}}$ |
| Event $\beta.2$ | A ends | “ B sending A key K_{AB} ” |

At this point, Alice believes that she has been contacted by Bob, when in fact she has been contacted by the intruder.

We can easily express this protocol in the spi-calculus, and use *begin* M and *end* M statements to specify the desired correspondence property.

We define $FlawedSender(net, alice, key)$ to be the sender, using net as the insecure communications medium, acting on behalf of $alice$ using secret key key (in order to bootstrap the system, we have the sender receive bob 's name from the network, so the attacker can create as many concurrent sessions as they like):

```

FlawedSender(
  net:Network, alice:Princ, key:WMFKey(alice)
)  $\triangleq$ 
  repeat
    in net (bob:Princ);
    new (sKey:SKey);
    begin “alice sending bob key sKey”;
    out net (alice);
    in net (nonce:Un);
    cast nonce is (nonce':WMFNonce(alice, bob, sKey));
    out net (alice, {bob, sKey, nonce'}key);

```

We define $FlawedReceiver(net, bob, key)$ to be the receiver, using net as the insecure

communications medium, acting on behalf of *bob*, using secret key *key*:

```

FlawedReceiver(
  net:Network,bob:Princ,key:WMFKey(bob)
)  $\triangleq$ 
repeat
  inp net ();
  new (nonce:Un);
  out net (nonce);
  inp net (ctxt:Un);
  decrypt ctxt
  is {alice:Princ,
      sKey : SKey,
      nonce' : WMFNonce(alice,bob,sKey) }key;
  check nonce is nonce';
  end “alice sending bob key sKey”

```

We define *FlawedServer*(*net*,*lookup*) to be the server, using *net* as the insecure communications medium, making use of a trusted database lookup function *lookup* to access the secret keys:

```

FlawedServer(net:Network,lookup:WMFLookup)  $\triangleq$ 
repeat
  inp net (alice:Princ);
  new (nonceA:Un);
  out net (nonceA);
  inp net (alice,ctxt:Un);
  let keyA : WMFKey(alice) = lookup (alice);
  decrypt ctxt
  is {bob:Princ,
      sKey : SKey,
      nonceA' : WMFNonce(alice,bob,sKey) }keyA;
  check nonceA is nonceA';
  out net ();
  inp net (nonceB:Un);
  cast nonceB
  is (nonceB':WMFNonce(alice,bob,sKey));
  let keyB : WMFKey(bob) = lookup (bob);
  out net {alice,sKey,nonceB'}keyB

```

Then we can try to define the types appropriately. For most of the types, it is fairly routine (for the *WMFLookup* type, we need to use an appropriate function type, and for the *SKey* type, we need an appropriate *Msg* type for the payload, but these do not play

an important role in the typing) :

$$\begin{aligned}
Network &\triangleq \text{Un} \\
Princ &\triangleq \text{Un} \\
WMF\text{Lookup} &\triangleq (princ:Princ) \rightarrow WMF\text{Key}(princ) \\
SKey &\triangleq \text{Key}(Msg) \\
WMF\text{Nonce}(alice, bob, sKey) &\triangleq \\
&\text{Nonce [end "alice sending bob key sKey"]} \\
WMF\text{Key}(princ) &\triangleq \text{Key}(WMF\text{Msg}(princ))
\end{aligned}$$

The problem comes when we try to give a definition for $WMF\text{Msg}$, which is the type of the plaintext of messages used in the WMF protocol. In order to type-check Message 3, we require:

$$\begin{aligned}
WMF\text{Msg}(alice) &= \\
&(bob:Princ, sKey:SKey, nonce:WMF\text{Nonce}(alice, bob, sKey))
\end{aligned}$$

and in order to type-check Message 6, we require:

$$\begin{aligned}
WMF\text{Msg}(bob) &= \\
&(alice:Princ, sKey:SKey, nonce:WMF\text{Nonce}(alice, bob, sKey))
\end{aligned}$$

Unfortunately, these requirements are inconsistent, since the roles of *alice* and *bob* have been swapped. This is the root of the attack on this broken WMF, which relies on the fact that the key for *alice* is being used in two incompatible ways, depending on whether *alice* is acting as the sender or the receiver.

This is an example of a type-flaw attack [23] and may be solved by the standard solution of adding tag information to messages. This is akin to the use of tagged union types in type-safe languages like ML or Haskell. In this case, we have the type for Message 3 of the protocol:

$$\begin{aligned}
WMF\text{Msg}_3(alice) &\triangleq \\
&(bob:Princ, sKey:SKey, nonce:WMF\text{Nonce}(alice, bob, sKey))
\end{aligned}$$

and the type for Message 6:

$$\begin{aligned}
WMF\text{Msg}_6(bob) &\triangleq \\
&(alice:Princ, sKey:SKey, nonce:WMF\text{Nonce}(alice, bob, sKey))
\end{aligned}$$

and we can define $WMF\text{Msg}(princ)$ as the tagged union of these two types:

$$\begin{aligned}
WMF\text{Msg}(princ) &\triangleq \\
&(msg_3(WMF\text{Msg}_3(princ)) \mid msg_6(WMF\text{Msg}_6(princ)))
\end{aligned}$$

We can then check that the safe versions of the principals are effect-free. The sender, receiver, and server are given in Figure 1.

The key database has to implement the *lookup* function, and be effect-free. In practice, an implementation would require access to a secure database, but in this example,

we can just hard-wire in the principal names and keys, and use pattern-matching to define the database:

$$\begin{aligned} & \text{KeyDB}(\text{lookup}:\text{WMFLookup}, \text{princ}_1:\text{Princ}, \text{key}_1:\text{WMFKey}(\text{princ}_1), \dots, \\ & \quad \text{princ}_n:\text{Princ}, \text{key}_n:\text{WMFKey}(\text{princ}_n)) \triangleq \\ & \text{function} \\ & \quad \text{lookup}(\text{princ}_1) : \text{WMFKey}(\text{princ}_1) \text{ is return } \text{key}_1 \\ & \quad \vdots \\ & \quad \text{lookup}(\text{princ}_n) : \text{WMFKey}(\text{princ}_n) \text{ is return } \text{key}_n \end{aligned}$$

We define a *Wide Mouth Frog configuration* to be a process of the form:

$$\begin{aligned} & \text{new } (\text{lookup}:\text{WMFLookup}); \\ & \text{new } (\text{princ}_1:\text{Princ}); \dots \\ & \text{new } (\text{princ}_n:\text{Princ}); \\ & \text{new } (\text{key}_1:\text{WMFKey}(\text{princ}_1)); \dots \\ & \text{new } (\text{key}_n:\text{WMFKey}(\text{princ}_n)); \\ & \text{FixedSender}(\text{net}, \text{princ}_1, \text{key}_1) \mid \dots \mid \\ & \text{FixedSender}(\text{net}, \text{princ}_n, \text{key}_n) \mid \\ & \text{FixedReceiver}(\text{net}, \text{princ}_1, \text{key}_1) \mid \dots \mid \\ & \text{FixedReceiver}(\text{net}, \text{princ}_n, \text{key}_n) \mid \\ & \text{FixedServer}(\text{net}, \text{lookup}) \mid \\ & \text{KeyDB}(\text{lookup}, \text{princ}_1, \text{key}_1, \dots, \text{princ}_n, \text{key}_n) \end{aligned}$$

We can then apply the results of this paper to get:

- Any Wide Mouth Frog configuration is effect-free, and hence robustly safe.

Thus, we have shown the Wide Mouth Frog protocol to satisfy this particular safety property for an arbitrary number of principals, sessions, and in the presence of an arbitrary attacker and well-typed database implementation.

The use of tagged unions to represent the different message types which are sent in a protocol is a common technique, and corresponds to the final phrase of Principle 10 of Abadi and Needham [4]:

If an encoding is used to present the meaning of a message, then it should be possible to tell which encoding is being used. In the common case where the encoding is protocol dependent, it should be possible to deduce that the message belongs to this protocol, and in fact to a particular run of the protocol, and to know its number in the protocol.

Many protocols use ad hoc techniques such as incrementing timestamps, or juggling the order of participant names to encode message numbers implicitly. Our type system makes these ad hoc solutions formal, as an instance of the standard technique of using tagged union types.

A.2 Woo and Lam’s Authentication Protocol

Woo and Lam [40] propose a server-based symmetric-key authentication protocol. Alice wishes to authenticate herself to Bob, and does so by responding to a nonce challenge with a message which Bob can ask the trusted server to decrypt:

| | | |
|-----------|---------------------|--|
| Event 1 | A begins | “A authenticates to B” |
| Message 1 | $A \rightarrow B :$ | A |
| Message 2 | $B \rightarrow A :$ | N_B |
| Message 3 | $A \rightarrow B :$ | $\{msg_3(N_B)\}_{K_{AS}}$ |
| Message 4 | $B \rightarrow S :$ | $\{msg_4(A, \{msg_3(N_B)\}_{K_{AS}})\}_{K_{BS}}$ |
| Message 5 | $S \rightarrow B :$ | $\{msg_5(N_B)\}_{K_{BS}}$ |
| Event 2 | B ends | “A authenticates to B” |

(In the original protocol, the messages were untagged, but we have provided tags for the reasons discussed in the previous section.) Abadi and Needham [4] demonstrate that this protocol is not robustly safe, because message 5 does not mention A.

The possibility of this attack is made clear when we try to type-check the protocol. We have types:

$$\begin{aligned}
 WLKey(princ) &\triangleq \text{Key}((WLMsg(princ))) \\
 WLMsg(princ) &\triangleq (msg_3(WLMsg_3(princ)) \mid \\
 &\quad msg_4(WLMsg_4(princ)) \mid \\
 &\quad msg_5(WLMsg_5(princ))) \\
 WLMsg_3(alice) &\triangleq (nonce:WLNonce(alice, bob)) \\
 WLMsg_4(bob) &\triangleq (alice:Princ, ctext:Un) \\
 WLMsg_5(bob) &\triangleq (nonce:WLNonce(alice, bob)) \\
 WLNonce(alice, bob) &\triangleq \text{Nonce} [\text{end } \text{“}alice \text{ authenticates to } bob\text{”}] \\
 WLLookup &\triangleq (princ:Princ) \rightarrow WLKey(princ)
 \end{aligned}$$

At this point it becomes clear that the protocol is not well-typed, since the types are not well-formed: $WLMsg_3(alice)$ contains an unbound occurrence of bob and $WLMsg_5(bob)$ contains an unbound occurrence of $alice$. Abadi and Needham observe that Message 5 should be changed to:

$$\text{Message 5'} \quad S \rightarrow B : \quad \{msg_5(A, N_B)\}_{K_{BS}}$$

but did not make any similar observation for Message 3. Their strengthened protocol allows Bob to know that Alice is talking to somebody, but does not allow Bob to know that Alice is talking to Bob. For example, one possible run, where Alice begins a

dialogue with Charlie, but is authenticated to Bob is:

| | | |
|--------------------|-------------------|--|
| Event $\alpha.1$ | A begins | “ A authenticates to C ” |
| Message $\alpha.1$ | $A \rightarrow I$ | A |
| Message $\beta.1$ | $I \rightarrow B$ | A |
| Message $\beta.2$ | $B \rightarrow I$ | N_B |
| Message $\alpha.2$ | $I \rightarrow A$ | N_B |
| Message $\alpha.3$ | $A \rightarrow I$ | $\{msg_3(N_B)\}_{K_{AS}}$ |
| Message $\beta.3$ | $I \rightarrow B$ | $\{msg_3(N_B)\}_{K_{AS}}$ |
| Message $\beta.4$ | $B \rightarrow S$ | $\{msg_4(A, \{msg_3(N_B)\}_{K_{AS}})\}_{K_{BS}}$ |
| Message $\beta.5$ | $S \rightarrow B$ | $\{msg_5(N_B)\}_{K_{BS}}$ |
| Event $\beta.2$ | B ends | “ A authenticates to B ” |

This attack is noted by Anderson and Needham [5], and is stopped by a similar change to the protocol:

Message 3' $A \rightarrow B$: $\{msg_3(B, N_B)\}_{K_{AS}}$

Finally, our type system makes clear that the encryption of message 4 is unnecessary, since all the data is of type Un , and so can safely be sent in plaintext, as suggested by Abadi and Needham [4]:

Message 4' $B \rightarrow S$: $A, B, \{msg_3(B, N_B)\}_{K_{AS}}$

The resulting protocol can be type-checked, using types:

$$\begin{aligned}
 WLMsg(princ) &\triangleq \\
 &(msg_3(WLMsg_3(princ)) \mid msg_5(WLMsg_5(princ))) \\
 WLMsg_3(alice) &\triangleq \\
 &(bob:Princ, nonce:WLNonce(alice, bob)) \\
 WLMsg_5(bob) &\triangleq \\
 &(alice:Princ, nonce:WLNonce(alice, bob))
 \end{aligned}$$

To see that the sender is effect-free, we calculate:

$$\begin{aligned}
 &FixedSender(net:Network, alice:Princ, key:WLKey(alice)) \triangleq \\
 &repeat \\
 &\quad inp \ net \ (bob:Princ); \\
 &\quad begin \ "alice \ authenticates \ to \ bob"; \\
 &\quad out \ net \ (alice) \\
 &\quad inp \ net \ (nonce:Un); \\
 &\quad cast \ nonce \ is \ (nonce':WLNonce(alice, bob)); \\
 &\quad out \ net \ \{msg_3(bob, nonce')\}_{key} \} [end \ \dots] \} \square
 \end{aligned}$$

To see that the receiver is effect-free, we calculate:

$$\text{FixedReceiver}(\text{net}:\text{Network}, \text{bob}:\text{Princ}, \text{key}:\text{WLKey}(\text{bob})) \triangleq$$

```

repeat
  inp net (alice:Princ);
  new (nonce:Un);
  out net (nonce)
  inp net (cctx:Un);
  out net (alice, bob, cctx)
  inp net ({msg5(alice, nonce':WLNonce(alice, bob))}_{key});
  check nonce is nonce';
  end "alice authenticates to bob" [end ...] } [check nonce]

```

To see that the server is effect-free, notice that the server makes no use of any process check N is N' ; P , cast N is (N') ; P or end M , and so is automatically effect-free:

$$\text{FixedServer}(\text{net}:\text{Network}, \text{lookup}:\text{WLLookup}) \triangleq$$

```

repeat
  inp net (alice:Princ, bob:Princ, cctx:Un);
  let keyA:WLKey(alice) = lookup(alice);
  let keyB:WLKey(bob) = lookup(bob);
  decrypt cctx is {msg3(bob, nonce':WLNonce(alice, bob))}_{keyA};
  out net {msg5(alice, nonce')}_{keyB}

```

We define a *Woo and Lam configuration* to be a process of the form:

```

new (lookup:WLLookup);
new (princ1:Princ); ... new (princn:Princ);
new (key1:WLKey(princ1)); ... new (keyn:WLKey(princn));
FixedSender(net, princ1, key1) | ... | FixedSender(net, princn, keyn)
  | FixedReceiver(net, princ1, key1) | ... | FixedReceiver(net, princn, keyn)
  | FixedServer(net, lookup) | KeyDB(lookup, princ1, key1, ..., princn, keyn)

```

for any effect-free *KeyDB*. We can then apply the results of this paper to get:

- Any Woo and Lam configuration is effect-free, and hence robustly safe.

This example has shown that in our type system, it is important that all messages contain the names of the principals involved. Our type system enforces Principle 3 of Abadi and Needham [4]:

If the identity of a principal is essential to the meaning of a message, it is prudent to mention the principal's name explicitly in the message.

This requirement is enforced through the usual requirement for variables in a program to be correctly scoped: violations of Principle 3 may be caught because a variable is used when it is not in scope.

A.3 Otway and Rees's Key Exchange Protocol

Otway and Rees [33] propose a server-based symmetric-key key exchange protocol. We cannot verify their protocol using the type system of this paper, even though (as far as we are aware) it is correct, since it relies on using nonces to stand for principal names, which are kept secret, as well as for freshness. Still, it may be possible to adapt our type system to deal with this use of nonces; we leave this for future work.

Abadi and Needham [4] propose a simplification of the protocol, which we verify here:

| | | |
|-----------|-------------------|---|
| Message 1 | $A \rightarrow B$ | A, B, N_A |
| Message 2 | $B \rightarrow S$ | A, B, N_A, N_B |
| Event 1 | S begins | “initiator A shares K_{AB} with B ” |
| Event 2 | S begins | “responder B shares K_{AB} with A ” |
| Message 3 | $S \rightarrow B$ | $\{msg_4(A, B, K_{AB}, N_A)\}_{K_{AS}},$ $\{msg_3(A, B, K_{AB}, N_B)\}_{K_{BS}}$ |
| Event 3 | B ends | “responder B shares K_{AB} with A ” |
| Message 4 | $B \rightarrow A$ | $\{msg_4(A, B, K_{AB}, N_A)\}_{K_{AS}}$ |
| Event 4 | A ends | “initiator A shares K_{AB} with B ” |

At the end of this dialogue, Alice and Bob both know that K_{AB} was generated by Sam for their private use. Alice does not know that Bob actually received K_{AB} , since this protocol does not ensure that Alice and Bob actually receive K_{AB} , just that nobody else does.

We can allocate types to this protocol:

$$\begin{aligned}
ORKey(princ) &\triangleq \\
&\text{Key}((msg_3(ORMsg_3(princ)) \mid msg_4(ORMsg_4(princ)))) \\
ORMsg_3(bob) &\triangleq \\
&(alice:Princ, bob':Princ, sKey:SKey, \\
&\text{nonce:ORNonce}_3(alice, bob, sKey)) \\
ORMsg_4(alice) &\triangleq \\
&(alice':Princ, bob:Princ, sKey:SKey, \\
&\text{nonce:ORNonce}_3(alice, bob, sKey)) \\
ORNonce_3(alice, bob, sKey) &\triangleq \\
&\text{Nonce}[\text{end “responder } bob \text{ shares } sKey \text{ with } alice”] \\
ORNonce_4(alice, bob, sKey) &\triangleq \\
&\text{Nonce}[\text{end “initiator } alice \text{ shares } sKey \text{ with } bob”] \\
ORLookup &\triangleq \\
&(princ:Princ) \rightarrow ORKey(princ)
\end{aligned}$$

and then type-check Alice:

$$\begin{array}{l}
 \text{FixedSender}(net:Network, alice:Princ, key:ORKey(alice)) \triangleq \\
 \text{repeat} \\
 \quad \text{inp } net \ (bob:Princ); \\
 \quad \text{new } (nonceA:Un); \\
 \quad \text{out } net \ (alice, bob, nonceA); \\
 \quad \text{inp } net \ (\{msg_4(alice, bob, sKey:SKey, \\
 \quad \quad \quad nonceA':ORNonce_4(alice, bob, sKey))\}_{keyA}); \\
 \quad \text{check } nonceA \text{ is } nonceA'; \\
 \quad \text{end "initiator } alice \text{ shares } sKey \text{ with } bob" \} [end \dots] \} [check \ nonceA] \} \square
 \end{array}$$

type-check Bob:

$$\begin{array}{l}
 \text{FixedReceiver}(net : Network, bob:Princ, key:ORKey(bob)) \triangleq \\
 \text{repeat} \\
 \quad \text{inp } net \ (alice:Princ, bob, nonceA:Un); \\
 \quad \text{new } (nonceB:Un); \\
 \quad \text{out } net \ (alice, bob, nonceA, nonceB); \\
 \quad \text{inp } net \ (ctxt:Un, \{msg_3(alice, bob, sKey:SKey, \\
 \quad \quad \quad nonceB:ORNonce_3(alice, bob, sKey))\}_{key}); \\
 \quad \text{check } nonceB \text{ is } nonceB'; \\
 \quad \text{end "responder } bob \text{ shares } sKey \text{ with } alice" \} [end \dots] \} [check \ nonceB] \} \square
 \end{array}$$

and type-check Sam:

$$\begin{array}{l}
 \text{FixedServer}(net:Network, lookup:ORLookup) \triangleq \\
 \text{repeat} \\
 \quad \text{inp } net \ (alice:Princ, bob:Princ, nonceA:Un, nonceB:Un); \\
 \quad \text{let } key_A:ORKey(alice) = lookup(alice); \\
 \quad \text{let } key_B:ORKey(bob) = lookup(bob); \\
 \quad \text{new } (sKey:SKey); \\
 \quad \text{begin "initiator } alice \text{ shares } sKey \text{ with } bob"; \\
 \quad \text{begin "responder } bob \text{ shares } sKey \text{ with } alice"; \\
 \quad \text{cast } nonceA \text{ is } (nonceA':ORNonce_4(alice, bob, sKey)); \\
 \quad \text{cast } nonceB \text{ is } (nonceB':ORNonce_3(alice, bob, sKey)); \\
 \quad \text{out } net \ (\{msg_4(alice, bob, sKey, nonceA')\}_{key_A}, \\
 \quad \quad \{msg_3(alice, bob, sKey, nonceB')\}_{key_B}) \} \square \} [end \dots] \} [end \dots] \} \square
 \end{array}$$

We can then apply the techniques of this paper to show that this modified protocol is robustly safe. This typing makes it clear that Bob's name is not required in Message 3 and Alice's name is not required in Message 4, and these names could be dropped without compromising the correspondence assertions.

A.4 A Secure Message Stream

In Section 4.4 we showed how we can verify a simple two-message protocol to ensure the authenticity of messages. The protocol relied on Alice to send Bob a nonce

challenge for every message Bob sends:

| | | |
|-----------|---------------------|--------------|
| Event 1 | A begins | A sent M |
| Message 1 | $B \rightarrow A :$ | N |
| Message 2 | $A \rightarrow B :$ | $\{M, N\}_K$ |
| Event 2 | B ends | A sent M |

This is rather inefficient, since it requires an acknowledgement message for every message. Instead, we could use *message identifiers* to ensure the freshness of messages without Alice having to send constant acknowledgements. Our language does not support message identifiers directly, but they can be coded in messages of nonces: each time Bob sends Alice a message, he sends *two* nonces: the nonce for the current message, and the nonce for the next message. This is enough for Alice to ensure freshness of messages:

| | | |
|-------------|---------------------|---------------------------|
| Message 0 | $B \rightarrow A :$ | N_1 |
| Event 1a | A begins | A sent M_1 |
| Message 1 | $A \rightarrow B :$ | $\{M_1, N_2, N_1\}_K$ |
| Event 1b | B ends | A sent M_1 |
| ... | | |
| Event na | A begins | A sent M_n |
| Message n | $A \rightarrow B :$ | $\{M_n, N_{n+1}, N_n\}_K$ |
| Event nb | B ends | A sent M_n |

In order to check this protocol, we need to make use of latent nonce effects, since nonce N_n is being used to ensure the freshness of nonce N_{n+1} . The types we use are:

$$\begin{aligned} \text{MidKey} &\triangleq \text{Key}((\text{msg}:\text{Msg}, \text{nonce}B:\text{Un}, \text{nonce}A:\text{MidNonce}(\text{msg}, \text{nonce}B))) \\ \text{MidNonce}(\text{msg}, \text{nonce}B) &\triangleq \text{Nonce} [\text{end} \text{“Sender sent } \text{msg}\text{”, check } \text{nonce}B] \end{aligned}$$

The receiver is type-checked:

$$\begin{aligned} \text{FixedReceiver}(\text{net}:\text{Network}, \text{key}:\text{MidKey}) &\triangleq \\ \left. \begin{array}{l} \text{new } (\text{nonce}A:\text{Un}); \\ \text{out } \text{net } \text{nonce}A \\ \text{FixedReceiver}(\text{net}, \text{key}, \text{nonce}A) \} [\text{check } \text{nonce}A] \end{array} \right\} [] \end{aligned}$$

where we use the recursive function:

$$\begin{aligned} \text{FixedReceiver}(\text{net}:\text{Network}, \text{key}:\text{MidKey}, \text{nonce}A:\text{Un}) &\triangleq \\ \text{inp } \text{net } (\{(\text{msg}:\text{Msg}, \text{nonce}B:\text{Un}, \text{nonce}A':\text{MidNonce}(\text{msg}, \text{nonce}B))\}_{\text{key}}); \\ \text{check } \text{nonce}A \text{ is } \text{nonce}A'; \\ \text{end} \text{“Sender sent } \text{msg}\text{”}; \\ \text{FixedReceiver}(\text{net}, \text{key}, \text{nonce}B) \end{aligned}$$

The sender is type-checked similarly. This example shows that it is useful for participants in a protocol to be able to pass nonces and nonce effects, as allowed by our effect system.

A.5 Abbreviations Used in Examples

We shall now show that the abbreviations we used in our examples can be defined in our type system. We made use of types for dependent records, tagged unions, and dependent function types:

Syntax sugar for use in types:

| | |
|---|----------------------------|
| $T, U ::=$ | type |
| ... | as in Sections 4.1 and 4.2 |
| $(x_1:T_1, x_2:T_2, \dots, x_n:T_n)$ | dependent record |
| $(\ell_1(T_1) \mid \dots \mid \ell_n(T_n))$ | tagged union |
| $(x:T) \rightarrow U$ | dependent function |

We allowed the construction of messages of record or tagged union type:

Syntax sugar for use in messages:

| | |
|---------------------|-------------------|
| $L, M, N ::=$ | message |
| ... | as in Section 2.1 |
| (M_1, \dots, M_n) | record |
| $\ell_i(M)$ | tagged union |

In processes, we can make use of function declaration, function call, function return, and pattern-matching:

Syntax sugar for use in processes:

| | |
|--|---------------------------------|
| $O, P, Q, R ::=$ | process |
| ... | as in Sections 2.1, 3.1 and 4.2 |
| function $f(X_1) : T_1$ is $P_1 \dots f(X_n) : T_n$ is P_n | function declaration |
| let $x:U = f(M); P$ | function call |
| return M | function return |
| match M is $X; P$ | pattern match |
| out $M P;$ | output with residual |
| inp $M (X); P$ | pattern matching input |
| decrypt M is $\{X\}_P;$ | pattern matching decrypt |

where X ranges over a grammar of patterns:

Patterns:

| | |
|---------------------|--------------|
| $X, Y, Z ::=$ | patterns |
| $x:T$ | variable |
| M | constant |
| (X_1, \dots, X_n) | tuple |
| $\ell_i(X)$ | tagged union |
| $\{X\}_M$ | cyphertext |

We will now give definitions for each of these extensions, beginning with types. Dependent records and tagged unions are routine, since we already have pairs and variants types. Dependent records use a variant of the translation of functions into the π -calculus [31]; this is explored in more detail in [18]).

Abbreviations for types:

$$\begin{aligned} (x_1:T_1, x_2:T_2, \dots, x_n:T_n) &\triangleq (x_1:T_1, (x_2:T_2, (\dots (x_n:T_n, ()) \dots))) \\ (\ell_1(T_1) \mid \dots \mid \ell_n(T_n)) &\triangleq (T_1 + (T_2 + (\dots (T_n + 0) \dots))) \\ (x:T) \rightarrow U &\triangleq \text{Ch}(x:T, \text{Ch}(U)) \end{aligned}$$

The translations of messages are similarly straightforward.

Abbreviations for messages:

$$\begin{aligned} (M_1, M_2, \dots, M_n) &\triangleq (M_1, (M_2, (\dots (M_n, ()) \dots))) \\ \ell_i(M) &\triangleq \text{in}_i(M) \\ \text{in}_1(M) &\triangleq \text{inl}(M) \\ \text{in}_{n+1}(M) &\triangleq \text{inr}(\text{in}_n(M)) \end{aligned}$$

We write out $x(M);P$ as a simple shorthand for out $x M \mid P$:

Abbreviations out $M N;P$:

$$\text{out } M N;P \triangleq (\text{out } M N) \mid P$$

We use a variant of Milner's translation of the λ -calculus into the π -calculus, extended to deal with pattern-matching.

Abbreviations for functions, where $f : (x:T) \rightarrow U$:

$$\begin{aligned} \text{function } f(X_1) : U_1 \text{ is } P_1 \cdots f(X_n) : U_n \text{ is } P_n &\triangleq \\ \text{repeat inp } f(\text{request}:(x:T, \text{Ch}(U))) & \\ (\text{match } \text{request} \text{ is } (X_1, \text{return}:\text{Ch}(U_1)); P_1 \mid \dots \mid) & \\ \text{match } \text{request} \text{ is } (X_n, \text{return}:\text{Ch}(U_n)); P_n & \\ \text{return } M &\triangleq \\ \text{out } \text{return } N & \\ \text{let } x:U = f(M);P &\triangleq \\ \text{new } (k:\text{Ch}(U)); \text{out } f(M, k); \text{inp } k(x:U);P & \end{aligned}$$

where we define pattern-matching as:

Abbreviations for pattern matching:

$$\begin{aligned} \text{inp } M(X);P &\triangleq \text{inp } M(x); \text{match } x \text{ is } X;P \\ \text{decrypt } M \text{ is } \{X\}_N;P &\triangleq \text{decrypt } M \text{ is } \{x\}_N; \text{match } x \text{ is } X;P \\ \text{match } M \text{ is } x:T;P &\triangleq P\{x \leftarrow M\} \end{aligned}$$

$\text{match } M \text{ is } (); P \triangleq P$
 $\text{match } M \text{ is } (N, X_1, \dots, X_n); P \triangleq \text{match } M \text{ is } (N, y); \text{match } y \text{ is } (X_1, \dots, X_n); P$
 $\text{match } M \text{ is } (X_0, X_1, \dots, X_n); P \triangleq \text{split } M \text{ is } (x, y);$
 $\text{match } x \text{ is } X_0; \text{match } y \text{ is } (X_1, \dots, X_n); P$
 $\text{match } M \text{ is } \text{in}_1(X); P \triangleq \text{case } M \text{ is } \text{inl}(x) \text{ match } x \text{ is } X; P \text{ is } \text{inr}(x) \text{ stop}$
 $\text{match } M \text{ is } \text{in}_{n+1}(X); P \triangleq \text{case } M \text{ is } \text{inl}(x) \text{ stop is } \text{inr}(x) \text{ match } x \text{ is } \text{in}_n(X); P$
 $\text{match } M \text{ is } \{X\}_N P; \triangleq \text{decrypt } M \text{ is } \{x\}_N; \text{match } x \text{ is } X; P$
 $\text{match } M \text{ is } N; P \triangleq \text{match } (M, ()) \text{ is } (N, x); P$

Thus we have demonstrated that our core language is powerful enough to describe the examples in this section.

B Formal Semantics of our Typed Spi-Calculus

This appendix develops a formal operational semantics for the spi-calculus. Hence, we make precise the informal definition of process safety stated in Section 3.1, and prove the type safety result, Theorem 1 (Safety), stated in Section 4.3.

We begin in Appendix B.1 by defining a trace semantics for the spi-calculus, and use it to define safety in Appendix B.2. In Appendix B.3, we state and prove a subject reduction property (that is, a type preservation property). Finally, in Appendix B.4 we exploit subject reduction to prove Theorem 1 (Safety).

B.1 A Trace Semantics for our Spi-Calculus

We use a trace semantics based on the Chemical Abstract Machine [7]. First, we define a structural equivalence $P \equiv Q$ on processes, and then we define the trace semantics in terms of structural equivalence. This is the same technique as Milner [31] uses in the presentation of the π -calculus, and Abadi and Gordon [3] use in the presentation of the spi-calculus.

Structural Equivalence: $P \equiv Q$

| | |
|--|--------------------|
| $P \equiv P$ | (Struct Refl) |
| $Q \equiv P \Rightarrow P \equiv Q$ | (Struct Symm) |
| $P \equiv Q, Q \equiv R \Rightarrow P \equiv R$ | (Struct Trans) |
| $P \equiv Q \Rightarrow \text{new } (x:T); P \equiv \text{new } (x:T); Q$ | (Struct Res) |
| $P \equiv Q \Rightarrow P \mid R \equiv Q \mid R$ | (Struct Par) |
| $P \mid \text{stop} \equiv P$ | (Struct Par Zero) |
| $P \mid Q \equiv Q \mid P$ | (Struct Par Comm) |
| $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$ | (Struct Par Assoc) |
| $\text{repeat } P \equiv P \mid \text{repeat } P$ | (Struct Repl Par) |
| $x \notin \text{fn}(P) \Rightarrow P \mid \text{new } (x:T); Q \equiv \text{new } (x:T); (P \mid Q)$ | (Struct Par Res) |

$$x \neq y, x \notin fn(U), y \notin fn(T) \Rightarrow \text{(Struct Res Res)}$$

$$\text{new } (x:T); \text{new } (y:U); P \equiv \text{new } (y:U); \text{new } (x:T); P$$

A trace of a process is a finite sequence of events. The set of possible events includes the begin- and end-events defined in Section 3.1, as well as other events representing various actions of processes.

Each process is given a trace semantics, where a trace is a sequence of events performed by the process. Events take the following forms.

Events:

| $\alpha, \beta ::=$ | events |
|---------------------|---------------------------------------|
| begin L | begin-event labelled with message L |
| end L | end-event labelled with message L |
| cast $x:T$ | cast-event of name x to type T |
| check x | check-event for nonce x |
| gen $x:T$ | fresh-event for name x |
| τ | internal-event |

Events may contain free names. For example, $fn(\text{end (Sender sent } msg)) = \{msg\}$.

Free names, $fn(\alpha)$, of an event α

$$fn(\tau) \triangleq \emptyset$$

$$fn(\text{cast } x:T) \triangleq \{x\} \cup fn(T)$$

$$fn(\text{check } x) \triangleq \{x\}$$

$$fn(\text{begin } M) \triangleq fn(M)$$

$$fn(\text{end } M) \triangleq fn(M)$$

$$fn(\text{gen } x:T) \triangleq \{x\} \cup fn(T)$$

Events may also contain generated names. For example, $gn(\text{gen } msg:Msg) = \{msg\}$.

Generated names, $gn(\alpha)$, of an event α

$$gn(\alpha) \triangleq \begin{cases} \{x\} & \text{if } \alpha = \text{gen } x:T \\ \emptyset & \text{otherwise} \end{cases}$$

We interpret events as follows:

- An event begin L arises from a process begin $L;P$, and represents the beginning of a correspondence.
- An event end L arises from a process end $L;P$, and represents the end of a correspondence.
- An event cast $N:T$ arises from a process cast N is $(x:T);P$, and represents the cast of an untrusted message into the type T , which the type system requires to be of the specific form Nonce es .

- An event check N arises from a process check N is $N;P$, and represents a successful check for the presence of a nonce.
- An event gen $x:T$ arises from a process new $(x:T);P$, and represents the generation of a fresh name x .
- An event τ arises from an internal computational step of a process.

For example, in the *FixedSystem(net)* example from Section 3.3, one possible sequence of events is:

- gen *nonce:Un*: the receiver generates a fresh untrusted name *nonce*.
- gen *msg:Msg*: the sender generates a new message *msg*.
- begin (Sender sent *msg*): the sender begins a correspondence.
- cast *nonce:MyNonce (msg)*: the sender casts the untrusted message *nonce* to the type *MyNonce*.
- check *nonce*: the receiver checks that the received nonce is *nonce*.
- end (Sender sent *msg*): the receiver ends a correspondence.

On the other hand, in the compromised system *FlawedSystem(net) | Attacker(net)* one possible sequence of events is:

- gen *msg:Msg*: the sender generates a new message *msg*.
- begin (Sender sent *msg*): the sender begins a correspondence.
- end (Sender sent *msg*): the receiver ends a correspondence.
- end (Sender sent *msg*): the receiver mistakenly ends the same correspondence twice.

Next, we give a formal definition of the events a process is capable of, using a *labelled transition system* semantics $P \xrightarrow{\alpha} P'$, meaning “ P can perform event α and become P' ”.

Labelled transitions: $P \xrightarrow{\alpha} P'$

| | |
|---|------------------|
| $\text{out } x M \mid \text{inp } x (y:T); P \xrightarrow{\tau} P\{y \leftarrow M\}$ | (Trans Comm) |
| $\text{split } (M, N) \text{ is } (x:T, y:U); P \xrightarrow{\tau} P\{x \leftarrow M\}\{y \leftarrow N\}$ | (Trans Split) |
| $\text{match } (M, N) \text{ is } (M, y:U); P \xrightarrow{\tau} P\{y \leftarrow N\}$ | (Trans Match) |
| $\text{case inl } (M) \text{ is inl } (x:T) P \text{ is inr } (y:U) Q \xrightarrow{\tau} P\{x \leftarrow M\}$ | (Trans Case Inl) |
| $\text{case inr } (M) \text{ is inl } (x:T) P \text{ is inr } (y:U) Q \xrightarrow{\tau} Q\{y \leftarrow M\}$ | (Trans Case Inr) |
| $\text{decrypt } \{M\}_N \text{ is } \{x:T\}_N; P \xrightarrow{\tau} P\{x \leftarrow M\}$ | (Trans Decrypt) |
| $\text{cast } x \text{ is } (y:T); P \xrightarrow{\text{cast } x:T} P\{y \leftarrow x\}$ | (Trans Cast) |
| $\text{check } x \text{ is } x; P \xrightarrow{\text{check } x} P$ | (Trans Check) |
| $\text{begin } M; P \xrightarrow{\text{begin } M} P$ | (Trans Begin) |
| $\text{end } M; P \xrightarrow{\text{end } M} P$ | (Trans End) |

$\text{new } (x:T); P \xrightarrow{\text{gen } x} P$ (Trans Gen)

$\text{gn}(\alpha) \cap \text{fn}(Q) = \emptyset \Rightarrow P \xrightarrow{\alpha} P' \Rightarrow P \mid Q \xrightarrow{\alpha} P' \mid Q$ (Trans Par)

$x \notin \text{fn}(\alpha) \Rightarrow P \xrightarrow{\alpha} P' \Rightarrow \text{new } (x:T); P \xrightarrow{\alpha} \text{new } (x:T); P'$ (Trans Res)

$P \equiv Q, Q \xrightarrow{\alpha} Q', Q' \equiv P' \Rightarrow P \xrightarrow{\alpha} P'$ (Trans \equiv)

A trace is a sequence of events which the process may perform.

Traces:

$s, t ::= \alpha_1, \dots, \alpha_n$ trace (written ε if $n = 0$)

We extend the definition of free and generated names to traces:

Free names, $\text{fn}(s)$, and generated names, $\text{gn}(s)$, of trace s

$\text{fn}(a_1, \dots, a_n) \triangleq \text{fn}(a_1) \cup \dots \cup \text{fn}(a_n)$

$\text{gn}(a_1, \dots, a_n) \triangleq \text{gn}(a_1) \cup \dots \cup \text{gn}(a_n)$

The traces of a process P are defined using a trace-labelled transition system $P \xrightarrow{s} P'$ meaning ‘ P performs trace s and becomes P' .’

Traced transitions: $P \xrightarrow{s} P'$

$P \equiv P' \Rightarrow P \xrightarrow{\varepsilon} P'$ (Trace \equiv)

$P \xrightarrow{\alpha} P'', P'' \xrightarrow{s} P' \Rightarrow P \xrightarrow{\alpha, s} P'$ (Trace Event) (where $\text{fn}(a) \cap \text{gn}(s) = \emptyset$)

For example one trace of $\text{FixedSystem}(\text{net})$ is (ignoring τ actions):

gen $\text{nonce}:\text{Un}$,
gen $\text{msg}:\text{Msg}$,
begin (Sender sent msg),
cast $\text{nonce}:\text{MyNonce}(\text{msg})$,
check nonce ,
end (Sender sent msg)

One trace of $\text{FlawedSystem}(\text{net}) \mid \text{Attacker}(\text{net})$ is:

gen $\text{msg}:\text{Msg}$,
begin (Sender sent msg),
end (Sender sent msg),
end (Sender sent msg)

B.2 Correspondence Traces and Safe Processes

We now define our notion of safety, through correspondence assertions. To do so, we need to recall some standard notation for unordered collections of data, or *multisets*. If x ranges over elements of some given set, we let xs range over multisets of elements of that set.

Multiset of elements

| | |
|---------------------|----------------------------------|
| $xs ::=$ | multiset |
| $[x_1, \dots, x_n]$ | unordered collection of elements |

We identify multisets up to permuting elements, so $[x, y] = [y, x]$ but not up to copying elements, so $[x] \neq [x, x]$. We define some standard operations on multisets.

Multiset algebra $xs + xs', xs \leq xs', xs - xs', x \in xs, xs \vee xs'$

| |
|---|
| $[x_1, \dots, x_m] + [y_1, \dots, y_n] \triangleq [x_1, \dots, x_m, y_1, \dots, y_n]$ |
| $xs \leq xs'$ if and only if $xs + xs'' = xs'$ for some xs'' |
| $xs - xs' \triangleq$ the smallest xs'' such that $xs \leq xs'' + xs'$ |
| $x \in xs$ if and only if $[x] \leq xs$ |
| $xs \vee xs' \triangleq$ the smallest xs'' such that $xs \leq xs''$ and $xs' \leq xs''$ |

For example:

- $[x, y] + [y, z] = [x, y, y, z]$.
- $[x, y] \leq [x, y, z]$ but $[x, y, y] \not\leq [x, y, z]$.
- $[x, y, y, z] - [w, x, y] = [y, z]$.
- $x \in [x, y]$ but $z \notin [x, y]$.
- $[x, y] \vee [y, z] = [x, y, z]$.

For example, we use M_s to range over multisets of messages.

Multisets of messages M_s :

| | |
|---------------------|----------------------------------|
| $M_s ::=$ | multiset of messages |
| $[M_1, \dots, M_n]$ | unordered collection of messages |

We define the *beginnings* and *endings* of a trace s as the multiset of event labels begun and ended, respectively, in s .

Beginnings, $begins(s)$, and endings, $ends(s)$, of a trace s

| |
|--|
| $begins(a_1, \dots, a_n) \triangleq begins(a_1) + \dots + begins(a_n)$ |
| where $begins(a) \triangleq \begin{cases} [M] & \text{if } a = \text{begin } M \\ [] & \text{otherwise} \end{cases}$ |

$$\begin{aligned}
\text{ends}(a_1, \dots, a_n) &\triangleq \text{ends}(a_1) + \dots + \text{ends}(a_n) \\
\text{where } \text{ends}(a) &\triangleq \begin{cases} [M] & \text{if } a = \text{end } M \\ [] & \text{otherwise} \end{cases}
\end{aligned}$$

Next, we say that a trace is a correspondence if its beginnings dominate its endings; that is, for each end-event labelled L , there is a corresponding begin-event labelled L .

Correspondence:

A trace s is a *correspondence* if and only if $\text{ends}(s) \leq \text{begins}(s)$.

For the example trace of $\text{FixedSystem}(\text{net})$ we have:

$$\begin{aligned}
&\text{begins}(\text{gen } \text{nonce:Un}, \text{gen } \text{msg:Msg}, \text{begin}(\text{Sender sent } \text{msg}), \\
&\quad \text{cast } \text{nonce:MyNonce}(\text{msg}), \text{check } \text{nonce}, \text{end}(\text{Sender sent } \text{msg})) \\
&\quad = [\text{Sender sent } \text{msg}] \\
&\text{ends}(\text{gen } \text{nonce:Un}, \text{gen } \text{msg:Msg}, \text{begin}(\text{Sender sent } \text{msg}), \\
&\quad \text{cast } \text{nonce:MyNonce}(\text{msg}), \text{check } \text{nonce}, \text{end}(\text{Sender sent } \text{msg})) \\
&\quad = [\text{Sender sent } \text{msg}]
\end{aligned}$$

Therefore, since this trace s satisfies $\text{ends}(s) \leq \text{begins}(s)$, it is a correspondence.

For the example trace of $\text{FlawedSystem}(\text{net}) \mid \text{Attacker}(\text{net})$ we have:

$$\begin{aligned}
&\text{begins}(\text{gen } \text{msg:Msg}, \text{begin}(\text{Sender sent } \text{msg}), \\
&\quad \text{end}(\text{Sender sent } \text{msg}), \text{end}(\text{Sender sent } \text{msg})) \\
&\quad = [\text{Sender sent } \text{msg}] \\
&\text{ends}(\text{gen } \text{msg:Msg}, \text{begin}(\text{Sender sent } \text{msg}), \\
&\quad \text{end}(\text{Sender sent } \text{msg}), \text{end}(\text{Sender sent } \text{msg})) \\
&\quad = [\text{Sender sent } \text{msg}, \text{Sender sent } \text{msg}]
\end{aligned}$$

Since this trace has $\text{ends}(s) \not\leq \text{begins}(s)$, it is not a correspondence.

We can now restate, precisely, the notions of safety and robust safety introduced informally in Section 3.1.

Safety and Robust Safety:

A process P is *safe* if and only if for all traces s and processes P' ,

if $P \xrightarrow{s} P'$ then s is a correspondence.

A process P is *robustly safe* if and only if for all opponent processes O , $P \mid O$ is safe

For example, since $\text{FlawedSystem}(\text{net}) \mid \text{Attacker}(\text{net})$ has a trace that is not a correspondence, it follows that $\text{FlawedSystem}(\text{net}) \mid \text{Attacker}(\text{net})$ is not safe. Since the process $\text{Attacker}(\text{net})$ is an opponent process, it follows that $\text{FlawedSystem}(\text{net})$ is not robustly safe.

B.3 Proof of Subject Reduction

In this section, we prove a subject reduction property for the labelled transition system, that transitions preserve typings. To do so, however, we need to extend the type system to accommodate the fact that cast-processes can change the type of a name after a transition.

We can illustrate some of the subtleties introduced by casting by considering three processes that are well-typed with respect to the typing environment defined by $E = x:\text{Un}, y:\text{Un}, z:\text{Ch}(\text{Nonce} [\text{end } y])$.

Firstly, the following example illustrates that a well-typed process can cast the name x , originally of type Un , into the distinct type $\text{Nonce} [\text{end } y]$.

$$\begin{aligned} P_1 &\triangleq \text{cast } x \text{ is } (x':\text{Nonce} [\text{end } y]); \text{out } z \ x' \\ E \vdash P_1 &: [\text{end } y] \\ P_1 &\xrightarrow{\text{cast } x:\text{Nonce} [\text{end } y]} \text{out } z \ x \end{aligned}$$

Secondly, the following example illustrates that the name y , originally of type Un , can be cast into the type $\text{Nonce} [\text{end } y]$, that depends on the name y itself.

$$\begin{aligned} P_2 &\triangleq \text{cast } y \text{ is } (y':\text{Nonce} [\text{end } y]); \text{out } z \ y' \\ E \vdash P_2 &: [\text{end } y] \\ P_2 &\xrightarrow{\text{cast } y:\text{Nonce} [\text{end } y]} \text{out } z \ y \end{aligned}$$

Thirdly, the following example illustrates that the name x can be cast to two distinct types, $\text{Nonce} [\text{end } x]$ and $\text{Nonce} [\text{end } y]$.

$$\begin{aligned} P_3 &\triangleq \text{cast } x \text{ is } (x':\text{Nonce} [\text{end } x]); \text{cast } x \text{ is } (x'':\text{Nonce} [\text{end } y]); \text{out } z \ x'' \\ E \vdash P_3 &: [\text{end } x, \text{end } y] \\ P_3 &\xrightarrow{\text{cast } x:\text{Nonce} [\text{end } y]} \xrightarrow{\text{cast } x:\text{Nonce} [\text{end } x]} \text{out } z \ x \end{aligned}$$

Moreover, the possibility that a name can come to inhabit multiple distinct types arises in the setting of our running example. Recall from Section 3.3 that we have:

$$net:\text{Un} \vdash \text{FixedSystem}(net) : []$$

Now, consider the attacker:

$$\text{Attacker}(net) \triangleq \text{inp } net \ (nonce:\text{Un}); \text{out } net \ (nonce); \text{out } net \ (nonce)$$

We can derive

$$net:\text{Un} \vdash \text{FixedSystem}(net) \mid \text{Attacker}(net) : []$$

but $\text{FixedSystem}(net) \mid \text{Attacker}(net)$ has the trace:

- $\text{gen } (nonce)$: receiver generates a nonce (initially $nonce:\text{Un}$).
- $\text{gen } (msg_1)$: sender generates a message $msg_1:\text{Msg}$.

- begin (Sender sent msg_1): sender begins correspondence 1.
- cast $nonce:Nonce$ [begin (Sender sent msg_1): sender casts $nonce$ (so now $nonce:Nonce$ [begin (Sender sent msg_1))].
- gen (msg_2): sender generates a message $msg_2:Msg$.
- begin (Sender sent msg_2): sender begins correspondence 2.
- cast $nonce:Nonce$ [begin (Sender sent msg_2): sender casts $nonce$ again (so now $nonce:Nonce$ [begin (Sender sent msg_2))].

At the end of this trace, $nonce$ has been given three incompatible types:

- Of an untrusted message, $nonce:Un$.
- Of a nonce for correspondence 1, $nonce:Nonce$ [begin (Sender sent msg_1)].
- Of a nonce for correspondence 2, $nonce:Nonce$ [begin (Sender sent msg_2)].

If we are going to allow names to have more than one type, we need to extend the definition of an environment to allow this.

To accommodate the possibility that a name of type Un can be cast to additional types of the form $Nonce\ es$, we allow additional entries of the form $+x:Nonce\ es$ to be added to environments.

Extended environments:

| | |
|-----------|--|
| $E ::=$ | environment |
| \dots | as before |
| $E, +x:T$ | extended entry (T always takes the form $Nonce\ es$) |

For example, we now allow the environment:

$$\begin{aligned}
 &nonce:Un, \\
 &msg_1:Msg, \\
 &+nonce:Nonce\ [begin\ (Sender\ sent\ msg_1)], \\
 &msg_2:Msg, \\
 &+nonce:Nonce\ [begin\ (Sender\ sent\ msg_2)]
 \end{aligned}$$

which records that $nonce$ originally had type Un , but has since been cast to two other nonce types.

We extend the definitions of $dom(E)$ and $fn(E)$:

Free names $fn(E)$ of an extended environment:

$$fn(E, +x:T) \triangleq fn(E) \cup \{x\} \cup fn(T)$$

Domain $dom(E)$ of an extended environment:

$$dom(E, +x:T) \triangleq dom(E) \cup \{x\}$$

We also extend the rules for typing with extended environments. The rule rule (Env $+x$) allows the formation of an extended environment $E, +x:T$ only when x originally had type Un , and now also has nonce type. This matches the type rule (Proc Cast). The rule (Msg $+x$) extracts type information from such extended environments.

Typing with extended environments:

$$\frac{\text{(Env } +x) \quad E \vdash x : Un \quad E \vdash es}{E, +x:\text{Nonce } es \vdash \diamond} \quad \frac{\text{(Msg } +x) \quad E, +x:T, E' \vdash \diamond}{E, +x:T, E' \vdash x : T}$$

We now show some standard properties about our extended type and effect system.

Lemma 3 (Environment) *If $E \vdash j$ then $E \vdash \diamond$.*

Proof Show by induction on the derivation of $E, E' \vdash j$ that if $E, E' \vdash j$ then $E \vdash \diamond$. \square

Lemma 4 (Weakening) *If $E, E'' \vdash j$ and $E, E', E'' \vdash \diamond$ then $E, E', E'' \vdash j$.*

Proof An induction on the derivation of $E, E'' \vdash j$. \square

Lemma 5 (Substitutivity) *If $E, x:T, E' \vdash j$ and $E \vdash M : T$ and $x \notin \text{dom}(E')$ then we have $E, (E'\{x \leftarrow M\}) \vdash j\{x \leftarrow M\}$.*

Proof First show by case analysis that if $E \vdash M : T$ and T is $\text{Ch}(U)$, $\text{Key}(U)$ or $\text{Nonce } es$ then M is a name. The result then follows by induction on the derivation of $E, x:T, E' \vdash j$. \square

Lemma 6 (Subsumption Elimination) *If $E \vdash P : es$ then $E \vdash P : fs$ can be derived without rule (Proc Subsum), where $fs \leq es$.*

Proof First show that if $E \vdash es$ and $E \vdash fs$ then $E \vdash es + fs$, $E \vdash es - fs$ and $E \vdash es \vee fs$. Then show by induction on derivation that if $E \vdash P : es$ then $E \vdash es$. The result then follows by induction on the derivation of $E \vdash P : es$. \square

Next, we show a standard property of our labelled transition system, that we can move every use of structural equivalence up to top level. To state this lemma, we use the shorthand $\text{new } (D); P$ where:

$$\text{new } (x_1:T_1, \dots, x_n:T_n); P \triangleq \text{new } (x_1:T_1); \dots \text{new } (x_n:T_n); P$$

This construct enjoys the derived type rule:

Derived type rule for $\text{new } (D); P$:

$$\frac{\text{(Proc Res } D) \text{ (where } \text{dom}(D) \cap \text{fn}(es - \text{checks}(D)) = \emptyset) \quad E, D \vdash P : es \quad D \text{ is generative}}{E \vdash \text{new } (D); P : es - \text{checks}(D)}$$

where $\text{checks}(x_1:T_1, \dots, x_n:T_n) = [\text{check } x_1, \dots, \text{check } x_n]$ and $x_1:T_1, \dots, x_n:T_n$ is *generative* if and only if T_1, \dots, T_n are all generative.

Lemma 7 (\equiv Elimination) *If $P \xrightarrow{\alpha} P'$ then:*

$$P \equiv \text{new } (D); (Q \mid R) \quad P' \equiv \text{new } (D); (Q' \mid R) \quad \text{fn}(\alpha) \cap \text{dom}(D) = \emptyset$$

and $Q \xrightarrow{\alpha} Q'$ can be derived without rules (Trans Par) (Trans Res) or (Trans \equiv).

Proof An induction on the derivation of $P \xrightarrow{\alpha} P'$. □

We now show that the effect judgment for processes is preserved by structural equivalence.

Proposition 8 (Subject Equivalence) *If $E \vdash P : es$ and $P \equiv Q$ then $E \vdash Q : es$.*

Proof Prove by induction on the derivation of \equiv that if $P \equiv Q$ or $Q \equiv P$ then if $E \vdash P : es$ then $E \vdash Q : es$. □

Next, we state the main result of this section, a subject reduction property for our labelled transition system.

Proposition 9 (Subject Reduction) *Suppose $E \vdash P : es$.*

- (1) *If $P \xrightarrow{\tau} P'$ then $E \vdash P' : es$.*
- (2) *If $P \xrightarrow{\text{cast } x:T} P'$ then either:*
 - (a) *$E \vdash P' : es$, or*
 - (b) *$E, +x:T \vdash P' : es - fs$ where $T = \text{Nonce } fs$ and $fs \leq es$.*
- (3) *If $P \xrightarrow{\text{check } x} P'$ then either:*
 - (a) *$E \vdash P' : es$, or*
 - (b) *$E \vdash x : \text{Nonce } fs$, $E \vdash P' : (es + fs) - [\text{check } x]$ and $\text{check } x \in es$.*
- (4) *If $P \xrightarrow{\text{begin } M} P'$ then $E \vdash P' : es + [\text{end } M]$.*
- (5) *If $P \xrightarrow{\text{end } M} P'$ then $E \vdash P' : es - [\text{end } M]$, and $\text{end } M \in es$.*
- (6) *If $P \xrightarrow{\text{gen } x:T} P'$ then (up to appropriate α -conversion of x) either:*
 - (a) *$E, x:T \vdash P' : es$ and T is generative, or*
 - (b) *$E, x:T \vdash P' : es + [\text{check } x]$ and T is Un.*

Proof

- (1) *If $P \xrightarrow{\tau} P'$ then such that $E \vdash P' : es$.*

A case analysis on the derivation of $P \xrightarrow{\tau} P'$.

Case (Trans Comm): By Lemma 6 (Subsumption Elimination) and Lemma 7 (\equiv Elimination); and Rules (Proc Par), (Proc Res D), (Proc Output), (Proc Input), (Proc Output Un) and (Proc Input Un), we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{out } x M \mid \text{inp } x (y:T); Q \mid R) \\
P' &\equiv \text{new } (D); (Q\{y \leftarrow M\} \mid R) \\
E, D &\vdash x : U \\
E, D &\vdash M : T \\
E, D, y:T &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq (es_Q + es_R) - \text{checks}(D)
\end{aligned}$$

where $T = \text{Un}$ and $U = \text{Un}$, or $U = \text{Ch}(T)$; D is generative; $\text{dom}(D) \cap \text{fn}((es_Q + es_R) - \text{checks}(D)) = \emptyset$; and $y \notin \text{fn}(es_Q)$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

Case (Trans Split): By Lemma 6 (Subsumption Elimination) and Lemma 7 (\equiv Elimination); and Rules (Proc Par), (Proc Res D), (Msg Pair), (Proc Split), (Msg Pair Un), (Proc Split Un) we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{split } (M, N) \text{ is } (x:T, y:U); Q \mid R) \\
P' &\equiv \text{new } (D); (Q\{x \leftarrow M\}\{y \leftarrow N\} \mid R) \\
E, D &\vdash M : T \\
E, D &\vdash N : U\{x \leftarrow M\} \\
E, D, x:T, y:U &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq (es_Q + es_R) - \text{checks}(D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}((es_Q + es_R) - \text{checks}(D)) = \emptyset$; $x \notin \text{fn}(es_Q)$; and $y \notin \text{fn}(es_Q)$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

Case (Trans Match): By Lemma 6 (Subsumption Elimination) and Lemma 7 (\equiv Elimination); and Rules (Proc Par), (Proc Res D), (Msg Pair), (Proc

Match), (Msg Pair Un), (Proc Match Un) we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{match } (M, N) \text{ is } (M, y:U\{x \leftarrow M\}); Q \mid R) \\
P' &\equiv \text{new } (D); (Q\{y \leftarrow N\} \mid R) \\
E, D &\vdash M : T \\
E, D &\vdash N : U\{x \leftarrow M\} \\
E, D, y:U\{x \leftarrow M\} &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq (es_Q + es_R) - \text{checks}(D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}((es_Q + es_R) - \text{checks}(D)) = \emptyset$; and $y \notin \text{fn}(es_Q)$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

Case (Trans Case Inl): By Lemma 6 (Subsumption Elimination) and Lemma 7 (\equiv Elimination); and Rules (Proc Par), (Proc Res D), (Msg Inl), (Proc Case), (Msg Inl Un), (Proc Case Un), we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{case inl } (M) \text{ is inl } (x:T) Q \text{ is inr } (y:U) R \mid S) \\
P' &\equiv \text{new } (D); (Q\{x \leftarrow M\} \mid S) \\
E, D &\vdash M : T \\
E, D, x:T &\vdash Q : es_Q \\
E, D, y:U &\vdash R : es_R \\
E, D &\vdash S : es_S \\
es &\geq (es_Q \vee es_R) + es_S
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}((es_Q \vee es_R) + es_S) - \text{checks}(D) = \emptyset$; $x \notin \text{fn}(es_Q)$; and $y \notin \text{fn}(es_R)$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

Case (Trans Case Inr): As for Case (Trans Case Inl).

Case (Trans Decrypt): By Lemma 6 (Subsumption Elimination) and Lemma 7 (\equiv Elimination); and Rules (Proc Par), (Proc Res D), (Msg Encrypt), (Proc

Decrypt), (Msg Encrypt Un), (Proc Decrypt Un), we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{decrypt } \{M\}_N \text{ is } \{x:T\}_N; Q \mid R) \\
P' &\equiv \text{new } (D); (Q\{x \leftarrow M\} \mid R) \\
E, D &\vdash M : T \\
E, D, x:T &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq (es_Q + es_R) - \text{checks}(D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}((es_Q + es_R) - \text{checks}(D)) = \emptyset$; and $x \notin \text{fn}(es_Q)$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

(2) If $P \xrightarrow{\text{cast } x:T} P'$ then either:

- (a) $E \vdash P' : es$, or
- (b) $E, +x:T \vdash P' : es - fs$ where $T = \text{Nonce } fs$ and $fs \leq es$.

By Lemmas 6 (Subsumption Elimination) and 7 (\equiv Elimination); and Rules (Proc Par), (Proc Cast) and (Proc Cast Un) we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{cast } x \text{ is } (y:T); Q \mid R) \\
P' &\equiv \text{new } (D); (Q\{y \leftarrow x\} \mid R) \\
E &\vdash x : \text{Un} \\
E, D, y:T &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq (es_Q + fs + es_R) - \text{checks}(D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}(\text{cast } x:T) = \emptyset$; $\text{dom}(D) \cap \text{fn}((es_Q + fs + es_R) - \text{checks}(D)) = \emptyset$; $y \notin \text{fn}(es_Q)$; and either $T = \text{Un}$ and $fs = []$ or $T = \text{Nonce } fs$.

Case ($T = \text{Un}$ and $fs = []$): By Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

Case ($T = \text{Nonce } fs$): Since $E, D, y:T \vdash Q : es_Q$, by Lemma 3 (Environment) and Rules (Env x) we have $E, D \vdash T$. Since $\text{fn}(T) \cap \text{dom}(D) = \emptyset$, and $x \notin \text{dom}(D)$, by repeated use of Lemma 5 (Substitutivity) we have $E \vdash T$ and

$E \vdash x : \text{Un}$. Then we use Rule (Env $+x$) to get $E, +x:T \vdash \diamond$, so we can apply Lemma 4 (Weakening) to get $E, +x:T, D, y:T \vdash Q : es_Q$. By Rule (Msg $+x$) we have $E, +x:T, D \vdash x : T$ and so we can apply Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) to get:

$$E, +x:T \vdash P' : es - fs$$

and $fs \leq es$ as required.

(3) If $P \xrightarrow{\text{check } x} P'$ then either:

(a) $E \vdash P' : es$, or

(b) $E \vdash x : \text{Nonce } fs$, $E \vdash P' : (es + fs) - [\text{check } x]$ and $\text{check } x \in es$.

By Lemmas 6 (Subsumption Elimination) and 7 (\equiv Elimination); and rules (Proc Par), (Proc Res D), (Proc Check), and (Proc Check Un) we have:

$$\begin{aligned} P &\equiv \text{new } (D); (\text{check } x \text{ is } x; Q \mid R) \\ P' &\equiv \text{new } (D); (Q \mid R) \\ E, D &\vdash x : T \\ E, D &\vdash Q : es_Q \\ E, D &\vdash R : es_R \\ es &\geq ((es_Q - fs) + fs' + es_R) - \text{checks}(D) \end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}(\text{check } x:T) = \emptyset$; $\text{dom}(D) \cap \text{fn}(((es_Q - fs) + fs' + es_R) - \text{checks}(D)) = \emptyset$; and either $T = \text{Un}$ and $fs = fs' = []$ or $T = \text{Nonce } fs$ and $fs' = [\text{check } x]$.

Case ($T = \text{Un}$ and $fs = fs' = []$) By Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es$$

as required.

Case ($T = \text{Nonce } fs$ and $fs' = [\text{check } x]$) By Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$\begin{aligned} E &\vdash x : \text{Nonce } fs \\ E &\vdash P' : (es + fs) - [\text{check } x] \end{aligned}$$

and $\text{check } x \in es$ as required.

(4) If $P \xrightarrow{\text{begin } M} P'$ then $E \vdash P' : es + [\text{end } M]$.

By Lemmas 6 (Subsumption Elimination) and 7 (\equiv Elimination); and Rules (Trans Begin), (Proc Par), (Proc Res D), and (Proc Begin) we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{begin } M; Q \mid R) \\
P' &\equiv \text{new } (D); (Q \mid R) \\
E, D &\vdash M : T \\
E, D &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq ((es_Q - [\text{end } M]) + es_R) - \text{checks } (D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}(\text{begin } M) = \emptyset$; and $\text{dom}(D) \cap \text{fn}((es_Q - [\text{end } M]) + es_R) - \text{checks } (D) = \emptyset$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es + [\text{end } M]$$

as required.

- (5) If $P \xrightarrow{\text{end } M} P'$ then $E \vdash P' : es - [\text{end } M]$, and $\text{end } M \in es$.

By Lemmas 6 (Subsumption Elimination) and 7 (\equiv Elimination); and Rules (Trans End), (Proc Par), (Proc Res D) and (Proc End) we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{end } M \mid Q) \\
P' &\equiv \text{new } (D); Q \\
E, D &\vdash M : T \\
E, D &\vdash Q : es_Q \\
es &\geq (es_Q + [\text{end } M]) - \text{checks } (D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}(\text{end } M) = \emptyset$; and $\text{dom}(D) \cap \text{fn}((es_Q + [\text{end } M]) - \text{checks } (D)) = \emptyset$. Then by Lemma 5 (Substitutivity), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E \vdash P' : es - [\text{end } M]$$

and $\text{end } M \in es$ as required.

- (6) If $P \xrightarrow{\text{gen } x:T} P'$ then (up to appropriate α -conversion of x) either:

- (a) $E, x:T \vdash P' : es$ and T is generative, or
- (b) $E, x:T \vdash P' : es + [\text{check } x]$ and T is Un.

By Lemmas 6 (Subsumption Elimination) and 7 (\equiv Elimination); and Rules (Trans Gen), (Proc Par), (Proc Res D) and (Proc Res) we have:

$$\begin{aligned}
P &\equiv \text{new } (D); (\text{new } (x:T); Q \mid R) \\
P' &\equiv \text{new } (D); (Q \mid R) \\
E, D, x:T &\vdash Q : es_Q \\
E, D &\vdash R : es_R \\
es &\geq ((es_Q - [\text{check } x]) + es_R) - \text{checks } (D)
\end{aligned}$$

where D is generative; $\text{dom}(D) \cap \text{fn}(\text{gen } x:T) = \emptyset$; $\text{dom}(D) \cap \text{fn}((es_Q + [\text{end } M]) - \text{checks}(D)) = \emptyset$; $x \notin \text{fn}(es_Q)$; and T is generative.

Case ($T = \text{Un}$) By Lemma 5 (Substitutivity), Lemma 4 (Weakening), and Rules (Proc Par), (Proc Res D) and (Proc Subsum) we have:

$$E, x:\text{Un} \vdash P' : es + [\text{check } x]$$

as required.

Case ($T \neq \text{Un}$) Since $E, D, x:T \vdash Q : es_Q$, we have by Lemma 3 (Environment) $E, D, x:T \vdash es_Q$ and so (since $E, D, x:T \not\vdash [\text{check } x]$) we have $\text{check } x \notin es_Q$. By Lemma 5 (Substitutivity), Lemma 4 (Weakening), and Rules (Proc Par) (Proc Res D) and (Proc Subsum) we have:

$$E, x:T \vdash P' : es$$

as required. □

B.4 Proof of Safety

The purpose of this appendix is to prove the type safety result, Theorem 1 (Safety). It asserts that a process assigned the empty effect is safe. This theorem is a key fact in the proof of the main result of the paper, Theorem 2 (Robust Safety), in Section 4.3.

To prove Theorem 1 (Safety), we actually prove a stronger invariant, Proposition 18 (Transition Safety), about processes with non-empty effects. To state the invariant we introduce a function $\text{ends}(E \vdash es)$ which computes the multiset of end-events represented by an effect es . With this notation, we can roughly state the invariant as follows:

- If $E \vdash P : es$ and $P \xrightarrow{s} P'$ then we can find E' and es' such that $E' \vdash P' : es'$ where $\text{ends}(E \vdash es) + \text{begins}(s) \geq \text{ends}(E' \vdash es') + \text{ends}(s)$.

From this, we deduce that every process $E \vdash P : []$ is safe.

However, we have some work ahead of us, in particular in defining the function $\text{ends}(es)$. A naïve definition would just be to count all of the end M effects in es , but this ignores the latent effect of nonces. Consider the following typing:

$$x:\text{Un}, +x:\text{Nonce} [\text{end } M] \vdash (\text{check } x \text{ is } x; \text{end } M) : [\text{check } x]$$

The process has trace $\text{check } x, \text{end } M$, which has an unbalanced end M , even though the effect of the process only contains a check x effect. So, in addition to counting the end-events, we need also to compute the end-events that may be unleashed by nonce effects.

Another problem is that we have to make sure that nonces are used *linearly*, that is, at most once. For example we need to ban processes such as:

$$x:\text{Un}, +x:\text{Nonce} [\text{end } M] \vdash (\text{check } x \text{ is } x; \text{check } x \text{ is } x; \text{end } M; \text{end } M) : [\text{check } x, \text{check } x]$$

which use a nonce more than once, or even worse:

$$x:\text{Un}, +x:\text{Nonce} [\text{end } M, \text{check } x] \vdash \\ (\text{check } x \text{ is } x; \text{end } M; \text{check } x \text{ is } x; \text{end } M; \dots) : [\text{check } x]$$

where we have a self-certifying nonce with the *cyclic* type $x : \text{Nonce} [\text{end } M, \text{check } x]$, which allows an unbounded number of unbalanced assertions.

We first define the *latent effects* of a well-typed message $E \vdash M : \text{Un}$. If M is anything other than a name, then the latent effects are empty. Otherwise if $M = x$, we find all the occurrences of $x:\text{Nonce}$ es in E , and sum them. For example:

$$\text{effects } (x:\text{Un}, +x:\text{Nonce} [\text{end } M], +x:\text{Nonce} [\text{check } N] \vdash x : \text{Un}) \\ = [\text{end } M, \text{check } N]$$

Effects $\text{effects } (E \vdash M : \text{Un})$ of a typed message $E \vdash M : \text{Un}$:

$$\begin{aligned} \text{effects } (E, x:T \vdash x : \text{Un}) &\triangleq [] \\ \text{effects } (E, +x:\text{Nonce } es \vdash x : \text{Un}) &\triangleq \text{effects } (E \vdash x : \text{Un}) + es \\ \text{effects } (E, x:T \vdash y : \text{Un}) &\triangleq \text{effects } (E \vdash y : \text{Un}) \quad (\text{when } x \neq y) \\ \text{effects } (E, +x:T \vdash y : \text{Un}) &\triangleq \text{effects } (E \vdash y : \text{Un}) \quad (\text{when } x \neq y) \\ \text{effects } (E \vdash M : \text{Un}) &\triangleq [] \quad (\text{when } M \text{ is not a name}) \end{aligned}$$

As discussed above, we maintain an invariant for well-typed systems, which is that they are *nonce linear*, so they only allow each nonce to be checked once. We define this in terms of a predicate $Ms \sqsubseteq \text{checks } (E \vdash es)$ which can be read as ‘ Ms is a lower bound on the nonce checks allowed by $E \vdash es$ ’. For example:

$$\begin{aligned} [x] &\sqsubseteq \text{checks } (x:\text{Un} \vdash [\text{check } x]) \\ [x, y] &\sqsubseteq \text{checks } (x:\text{Un}, y:\text{Un} \vdash [\text{check } x, \text{check } y]) \\ [x, y] &\sqsubseteq \text{checks } (x:\text{Un}, y:\text{Un}, +x:\text{Nonce} [\text{check } y] \vdash [\text{check } x]) \\ [x, x] &\sqsubseteq \text{checks } (x:\text{Un} \vdash [\text{check } x, \text{check } x]) \\ [x, x] &\sqsubseteq \text{checks } (x:\text{Un}, +x:\text{Nonce} [\text{check } x] \vdash [\text{check } x]) \end{aligned}$$

When we calculate the lower bound on the nonces allowed by $E \vdash es$, we include the latent effects of es . In particular, the last example shows that we have to be careful about cyclic uses of nonces.

Lower bound $Ms \sqsubseteq \text{checks } (E \vdash es)$ of the nonces of a typed effect $E \vdash es$:

$$\frac{(\text{Nonces } []) \quad (\text{Nonces check } M)}{Ms - [M] \sqsubseteq \text{checks } (E \vdash es + \text{effects } (E \vdash M : \text{Un}))} \\ \frac{[] \sqsubseteq \text{checks } (E \vdash es)}{Ms \sqsubseteq \text{checks } (E \vdash es + [\text{check } M])}$$

Having defined $Ms \sqsubseteq \text{checks } (E \vdash es)$, we can define the nonce linear and nonce acyclic effects:

Nonce linear typed effects $E \vdash es$

A typed effect $E \vdash es$ is *nonce linear* if and only if there is no M such that $[M, M] \sqsubseteq checks (E \vdash es)$.

Nonce acyclic typed effects $E \vdash es$

A typed effect $E \vdash es$ is *nonce acyclic* if and only if there is no $[M] \sqsubseteq checks (E \vdash es)$ such that $[M, M] \sqsubseteq checks (E \vdash [check M])$.

For example:

| | |
|--|-------------------------------|
| $x:Un \vdash [check x]$ | is linear and acyclic |
| $x:Un, y:Un \vdash [check x, check y]$ | is linear and acyclic |
| $x:Un, y:Un, +x:Nonce \checkmark y \vdash [check x]$ | is linear and acyclic |
| $x:Un \vdash [check x, check x]$ | is acyclic but not linear |
| $x:Un, +x:Nonce \checkmark x \vdash [check x]$ | is neither linear nor acyclic |

We can now show some properties about nonce linear effects, and nonce acyclic effects, in particular that every nonce linear effect is nonce acyclic.

Lemma 10 (Nonce monotonicity) *If $Ms \sqsubseteq checks (E \vdash es)$ and $E \vdash fs$ then $Ms \sqsubseteq checks (E \vdash es + fs)$.*

Proof An induction on the proof of $Ms \sqsubseteq checks (E \vdash es)$. □

Lemma 11 (Nonce transitivity) *If we have $[M] \sqsubseteq checks (E \vdash es)$ and also that $Ms \sqsubseteq checks (E \vdash [check M])$ then $Ms \sqsubseteq checks (E \vdash es)$.*

Proof An induction on the proof of $[M] \sqsubseteq checks (E \vdash es)$. To get $[M] \sqsubseteq checks (E \vdash es)$ we must have used Rule (Nonces check M) and so either:

- $check M \in es$, so by Lemma 10 (Nonce monotonicity) we have $Ms \sqsubseteq checks (E \vdash es)$, or
- $es = fs + [check N]$ and $[M] \sqsubseteq checks (E \vdash fs + effects (E \vdash N : Un))$, so by induction $Ms \sqsubseteq checks (E \vdash fs + effects (E \vdash N : Un))$, and so by Rule (Nonces check M) $Ms \sqsubseteq checks (E \vdash es)$.

The result follows. □

Lemma 12 (Linear implies acyclic) *If $E \vdash es$ is nonce linear then $E \vdash es$ is nonce acyclic.*

Proof Follows from Lemma 11 (Nonce transitivity). □

We can now define *ends* ($E \vdash es$) for a nonce acyclic effect $E \vdash es$. This is used to set up the invariant for our type safety result.

Endings $ends(E \vdash es)$ of a nonce acyclic effect $E \vdash es$:

$$\begin{aligned}
 ends(E \vdash []) &\triangleq [] \\
 ends(E \vdash es + [end M]) &\triangleq ends(E \vdash es) + [M] \\
 ends(E \vdash es + [check M]) &\triangleq ends(E \vdash es) + ends(E \vdash effects(E \vdash M : \text{Un}))
 \end{aligned}$$

Note that $ends(E \vdash es)$ is not well-defined for nonce cyclic effects, for example if:

$$E = x:\text{Un}, +x:\text{Nonce} [check\ x, end\ M]$$

then:

$$\begin{aligned}
 ends(E \vdash [check\ x]) & \\
 &= ends(E \vdash [check\ x, end\ M]) \\
 &= ends(E \vdash [check\ x, end\ M, end\ M]) \\
 &= \dots
 \end{aligned}$$

However, they are well-defined for nonce acyclic effects, which is enough for our purposes.

Lemma 13 (End Definedness) *If $E \vdash es$ is nonce acyclic then $ends(E \vdash es)$ is well-defined.*

Proof For any finite set of names X , let $ends(E \vdash es)_X$ be defined:

$$\begin{aligned}
 ends(E \vdash [])_X & \\
 &= [] \\
 ends(E \vdash es + [end M])_X & \\
 &= ends(E \vdash es)_X + [M] \\
 ends(E \vdash es + [check M])_X & \\
 &= \begin{cases} ends(E \vdash es)_X + ends(E \vdash effects(E \vdash M : \text{Un}))_{X - \{M\}} & \text{if } M \in X \\ ends(E \vdash es)_X & \text{otherwise} \end{cases}
 \end{aligned}$$

It is routine to see that $ends(E \vdash es)_X$ is well-defined, by induction first on X then on es . We then show by induction on the definition of $ends(E \vdash es)_X$ that:

$$\text{if } \forall [x] \trianglelefteq checks(E \vdash es) . x \in X \text{ then } ends(E \vdash es)_X = ends(E \vdash es)$$

In particular, we have that $ends(E \vdash es)_{dom(E)} = ends(E \vdash es)$, and so $ends(E \vdash es)$ is well-defined. \square

We can now prove some lemmas, leading up to the type safety results we need to show that effect-free processes are safe.

Lemma 14 (End Homomorphism) $ends(E \vdash es + fs) = ends(E \vdash es) + ends(E \vdash fs)$.

Proof An induction on es . □

Lemma 15 (End $+x$) *If $[x] \not\triangleleft$ checks $(E \vdash es)$ then $ends(E, +x:T \vdash es) = ends(E \vdash es)$.*

Proof An induction on es . □

Lemma 16 (End Nonce) *If $E \vdash x : \text{Nonce } es$ and $E \vdash x : \text{Un}$ then $ends(E \vdash [\text{check } x]) \geq ends(E \vdash es)$.*

Proof Show by induction on E that $es \leq effects(E \vdash x : \text{Un})$. The result then follows by Lemma 14 (End Homomorphism). □

Lemma 17 (End Add Nonce) *If $E \vdash es + fs$ is nonce linear then $ends(E \vdash es + fs) \geq ends(E, +x:\text{Nonce } fs \vdash es)$*

Proof An induction on es . The only interesting case is when:

$$es = es' + [\text{check } x]$$

Since $E \vdash es + fs$ is nonce linear, we have $[x] \not\triangleleft$ checks $(E \vdash es')$ and $[x] \triangleleft$ checks $(E \vdash effects(E, +x:\text{Nonce } fs \vdash x : \text{Un}))$ and so by Lemmas 14 (End Homomorphism) and 15 (End $+x$):

$$\begin{aligned} & ends(E \vdash es + fs) \\ &= ends(E \vdash es' + [\text{check } x] + fs) \\ &= ends(E \vdash es' + [\text{check } x]) + ends(E \vdash fs) \\ &= ends(E \vdash es') + ends(E \vdash effects(E \vdash x : \text{Un})) + ends(E \vdash fs) \\ &= ends(E \vdash es') + ends(E \vdash effects(E \vdash x : \text{Un}) + fs) \\ &= ends(E, +x:\text{Nonce } fs \vdash es') + \\ &\quad ends(E, +x:\text{Nonce } fs \vdash effects(E, +x:\text{Nonce } fs \vdash x : \text{Un})) \\ &= ends(E, +x:\text{Nonce } fs \vdash es' + [\text{check } x]) \\ &= ends(E, +x:\text{Nonce } fs \vdash es) \end{aligned}$$

as required.

Proposition 18 (Transition Safety) *If $E \vdash es$ is nonce linear, $E \vdash P : es$ and $P \xrightarrow{\alpha} P'$ then $E' \vdash P' : es'$ for some nonce linear $E' \vdash es'$ such that $ends(E \vdash es) + begins(\alpha) \geq ends(E' \vdash es') + ends(\alpha)$.*

Proof A case analysis on α :

Case ($\alpha = \text{cast } x:T$) By Proposition 9 (Subject Reduction), we have one of the following cases:

Subcase ($E \vdash P' : es$) Immediate.

Subcase $(E, +x:T \vdash P' : es - fs$ **where** $T = \text{Nonce } fs$ **and** $fs \leq es)$ Then, using Lemma 17 (End Add Nonce) we have:

$$\begin{aligned}
& ends(E \vdash es) + begins(\alpha) \\
&= ends(E \vdash es) \\
&\geq ends(E, +x:T \vdash es - fs) \\
&= ends(E, +x:T \vdash es - fs) + ends(\alpha)
\end{aligned}$$

and $E, +x:T \vdash es - fs$ is nonce linear as required.

Case $(\alpha = \text{check } x)$ By Proposition 9 (Subject Reduction), we have one of the following cases:

Subcase $(E \vdash P' : es)$ Immediate.

Subcase $(E \vdash x : \text{Nonce } fs, E \vdash P' : (es + fs) - [\text{check } x], \text{check } x \in es)$ Given Lemmas 16 (End Nonce) and 14 (End Homomorphism) we have:

$$\begin{aligned}
& ends(E \vdash es) + begins(\alpha) \\
&= ends(E \vdash es) \\
&= ends(E \vdash (es - [\text{check } x]) + [\text{check } x]) \\
&= ends(E \vdash (es - [\text{check } x])) + ends(E \vdash [\text{check } x]) \\
&\geq ends(E \vdash (es - [\text{check } x])) + ends(E \vdash fs) \\
&= ends(E \vdash (es + fs) - [\text{check } x]) \\
&= ends(E \vdash (es + fs) - [\text{check } x]) + ends(\alpha)
\end{aligned}$$

and $E \vdash (es + fs) - [\text{check } x]$ is nonce linear as required.

Case $(\alpha = \text{begin } M)$ Follows directly from Proposition 9 (Subject Reduction).

Case $(\alpha = \text{end } M)$ Follows directly from Proposition 9 (Subject Reduction).

Case $(\alpha = \text{gen } x:T)$ Follows directly from Proposition 9 (Subject Reduction).

Case $(\alpha = \tau)$ Follows directly from Proposition 9 (Subject Reduction). \square

Proposition 19 (\xrightarrow{s} **Safety**) *If $E \vdash es$ is nonce linear, $E \vdash P : es$ and $P \xrightarrow{s} P'$ then $E' \vdash P' : es'$ for some nonce linear $E' \vdash es'$ such that $ends(E \vdash es) + begins(s) \geq ends(E' \vdash es') + ends(s)$.*

Proof An induction on the derivation of $P \xrightarrow{s} P'$.

Case (Trace \equiv) We have:

$$\begin{aligned}
s &= \epsilon \\
P &\equiv P'
\end{aligned}$$

By Proposition 8 (Subject Equivalence):

$$E \vdash P' : es$$

and so (since $\text{begins}(\varepsilon) = \text{ends}(\varepsilon) = []$) we have:

$$\text{ends}(E \vdash es) + \text{begins}(s) \geq \text{ends}(E \vdash es) + \text{ends}(s)$$

as required.

Case (Trace Event) We have:

$$\begin{aligned} P &\xrightarrow{\alpha} P'' \\ P'' &\xrightarrow{t} P' \\ s &= \alpha, t \end{aligned}$$

By Proposition 18 (Transition Safety) we can find nonce linear $E'' \vdash es''$ such that:

$$\begin{aligned} E'' &\vdash P'' : es'' \\ \text{ends}(E \vdash es) + \text{begins}(\alpha) &\geq \text{ends}(E'' \vdash es'') + \text{ends}(\alpha) \end{aligned}$$

By induction, we can find nonce linear $E' \vdash es'$ such that:

$$\begin{aligned} E' &\vdash P' : es' \\ \text{ends}(E'' \vdash es'') + \text{begins}(t) &\geq \text{ends}(E' \vdash es') + \text{ends}(t) \end{aligned}$$

and so:

$$\begin{aligned} &\text{ends}(E \vdash es) + \text{begins}(s) \\ &= \text{ends}(E \vdash es) + \text{begins}(\alpha) + \text{begins}(t) \\ &\geq \text{ends}(E'' \vdash es'') + \text{ends}(\alpha) + \text{begins}(t) \\ &\geq \text{ends}(E' \vdash es') + \text{ends}(\alpha) + \text{ends}(t) \\ &= \text{ends}(E' \vdash es') + \text{ends}(s) \end{aligned}$$

as required. □

We have now done all the work required to show our main theorem: any effect-free process is safe.

Proof of Theorem 1 (Safety) *If $E \vdash P : []$ then P is safe.*

Proof If $P \xrightarrow{s} P'$ then we use Proposition 19 (\xrightarrow{s} Safety) to get:

$$\begin{aligned} &\text{begins}(s) \\ &= \text{ends}(E \vdash []) + \text{begins}(s) \\ &\geq \text{ends}(E' \vdash es') + \text{ends}(s) \\ &\geq \text{ends}(s) \end{aligned}$$

Thus, P is safe. □

$$\text{FixedSender}(\text{net}:\text{Network}, \text{alice}:\text{Princ}, \text{key}:\text{WMFKey}(\text{alice})) \triangleq$$

```

repeat
  inp net (bob:Princ);
  new (sKey:SKey);
  begin "alice sending bob key sKey";
  out net (alice);
  inp net (nonceA:Un);
  cast nonceA is (nonceA':WMFNonce(alice, bob, sKey));
  out net (alice, {msg3(bob, sKey, nonceA')key});
  [end ...]

```

$$\text{FixedReceiver}(\text{net}:\text{Network}, \text{bob}:\text{Princ}, \text{key}:\text{WMFKey}(\text{bob})) \triangleq$$

```

repeat
  inp net ();
  new (nonceB:Un);
  out net (nonceB);
  inp net (cText:Un);
  decrypt cText is {msg6(alice, sKey, nonceB')key};
  check nonceB is nonceB';
  end "alice sending bob key sKey" [end ...]

```

$$\text{FixedServer}(\text{net}:\text{Network}, \text{lookup}:\text{WMFLookup}) \triangleq$$

```

repeat
  inp net (alice:Princ);
  new (nonceA:Un);
  out net (nonceA);
  inp net (alice, cText:Un);
  let keyA : WMFKey(alice) = lookup(alice);
  decrypt cText is {msg3(bob, sKey, nonceA')keyA};
  check nonceA is nonceA';
  out net ();
  inp net (nonceB:Un);
  cast nonceB is (nonceB':WMFNonce(alice, bob, sKey));
  let keyB : WMFKey(bob) = lookup(bob);
  out net {msg6(alice, sKey, nonceB')keyB};
  [end ...]

```

Figure 1: Type checked participants in the Wide Mouth Frog protocol

References

- [1] M. Abadi. Secrecy by typing in security protocols. *Journal of the ACM*, 46(5):749–786, September 1999.
- [2] M. Abadi and B. Blanchet. Secrecy types for asymmetric communication. In *Foundations of Software Science and Computation Structures (FoSSaCS 2001)*, volume 2030 of *Lectures Notes in Computer Science*, pages 25–41. Springer, 2001.
- [3] M. Abadi and A.D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148:1–70, 1999.
- [4] M. Abadi and R. Needham. Prudent engineering practice for cryptographic protocols. *IEEE Transactions on Software Engineering*, 22(1):6–15, 1996.
- [5] R. Anderson and R. Needham. Programming Satan’s computer. In J. van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lectures Notes in Computer Science*, pages 426–440. Springer, 1995.
- [6] M. Bellare and P. Rogaway. Entity authentication and key distribution. In *Advances in Cryptology: CRYPTO’93*, volume 773 of *Lectures Notes in Computer Science*, pages 232–249. Springer, 1994.
- [7] G. Berry and G. Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96(1):217–248, April 1992.
- [8] D. Bolignano. An approach to the formal verification of cryptographic protocols. In *Third ACM Conference on Computer and Communications Security*, pages 106–118, 1996.
- [9] M. Burrows, M. Abadi, and R.M. Needham. A logic of authentication. *Proceedings of the Royal Society of London A*, 426:233–271, 1989.
- [10] E. Cohen. TAPS: A first-order verifier for cryptographic protocols. In *13th Computer Security Foundations Workshop*, pages 144–158. IEEE Computer Society Press, 2000.
- [11] S. Dal Zilio and A.D. Gordon. Region analysis and a π -calculus with groups. In *Mathematical Foundations of Computer Science 2000 (MFCS2000)*, volume 1893 of *Lectures Notes in Computer Science*, pages 1–21. Springer, 2000.
- [12] D. Dolev and A.C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, IT-29(2):198–208, 1983.
- [13] A. Durante, R. Focardi, and R. Gorrieri. A compiler for analysing cryptographic protocols. *ACM Transactions on Software Engineering and Methodology*, 2000. To appear.
- [14] R. Focardi, R. Gorrieri, and F. Martinelli. Message authentication through non-interference. In *International Conference on Algebraic Methodology And Software Technology (AMAST2000)*, volume 1816 of *Lectures Notes in Computer Science*, pages 258–272. Springer, 2000.

- [15] D.K. Gifford and J.M. Lucassen. Integrating functional and imperative programming. In *ACM Conference on Lisp and Functional Programming*, pages 28–38, 1986.
- [16] D. Gollmann. What do we mean by entity authentication? In *1995 IEEE Computer Society Symposium on Research in Security and Privacy*, pages 46–54, 1995.
- [17] L. Gong, R. Needham, and R. Yahalom. Reasoning about beliefs in cryptographic protocols. In *1990 IEEE Computer Society Symposium on Research in Security and Privacy*, 1990.
- [18] A.D. Gordon and A. Jeffrey. Typing correspondence assertions for communication protocols. In *Mathematical Foundations of Programming Semantics 17*, Electronic Notes in Theoretical Computer Science. Elsevier, 2001. To appear.
- [19] A.D. Gordon and D. Syme. Typing a multi-language intermediate code. In *28th ACM Symposium on Principles of Programming Languages (POPL'01)*, pages 248–260, 2001.
- [20] O. Grumberg and D.E. Long. Model checking and modular verification. *ACM Transactions on Programming Languages and Systems*, 16(3):843–871, 1994.
- [21] J.D. Guttman and F.J. Thayer Fábrega. Authentication tests. In *2000 IEEE Computer Society Symposium on Research in Security and Privacy*, 2000.
- [22] J. Heather. ‘Oh! ... Is it really you?’ *Using rank functions to verify authentication protocols*. PhD thesis, Royal Holloway, University of London, 2000.
- [23] J. Heather, G. Lowe, and S. Schneider. How to prevent type flaw attacks on security protocols. In *13th Computer Security Foundations Workshop*, pages 255–268. IEEE Computer Society Press, 2000.
- [24] J. Heather and S. Schneider. Towards automatic verification of authentication protocols on an unbounded network. In *13th Computer Security Foundations Workshop*, pages 132–143. IEEE Computer Society Press, 2000.
- [25] M. Hennessy and J. Riely. Resource access control in systems of mobile agents. In *3rd International Workshop on High-Level Concurrent Languages*, volume 16(3) of *Electronic Notes in Theoretical Computer Science*. Elsevier, 1998.
- [26] T. Lindholm and F. Yellin. *The Java™ Virtual Machine Specification*. Addison-Wesley, 1997.
- [27] G. Lowe. A hierarchy of authentication specifications. In *10th Computer Security Foundations Workshop*, pages 31–43. IEEE Computer Society Press, 1995.
- [28] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using CSP and FDR. In T. Margaria and B. Steffen, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS'96)*, volume 1055 of *Lectures Notes in Computer Science*, pages 147–166. Springer, 1996.
- [29] J.M. Lucassen. *Types and effects, towards the integration of functional and imperative programming*. PhD thesis, MIT, 1987. Available as Technical Report MIT/LCS/TR-408, MIT Laboratory for Computer Science.

- [30] W. Marrero, E.M. Clarke, and S. Jha. Model checking for security protocols. In *DIMACS Workshop on Design and Formal Verification of Security Protocols*, 1997. Preliminary version appears as Technical Report TR-CMU-CS-97-139, Carnegie Mellon University, May 1997.
- [31] R. Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999.
- [32] G. Necula. Proof-carrying code. In *24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 106–119. ACM Press, 1997.
- [33] D. Otway and O. Rees. Efficient and timely “mutual authentication”. *Operating Systems Review*, 21(1):8–10, 1987.
- [34] L.C. Paulson. The inductive approach to verifying cryptographic protocols. *Journal of Computer Security*, 6:85–128, 1998.
- [35] B. Pierce and E. Sumii. Relating cryptography and polymorphism. Available from the authors, 2000.
- [36] S.A. Schneider. Verifying authentication protocols in CSP. *IEEE Transactions on Software Engineering*, 24(9), September 1998.
- [37] C. Skalka and S. Smith. Static enforcement of security with types. In P. Wadler, editor, *2000 ACM International Conference on Functional Programming*, pages 34–45, 2000.
- [38] D.X. Song. Athena: a new efficient automatic checker for security protocol analysis. In *12th Computer Security Foundations Workshop*. IEEE Computer Society Press, 1999.
- [39] F.J. Thayer Fábrega, J.C. Herzog, and J.D. Guttman. Strand spaces: Why is a security protocol correct? In *1998 IEEE Computer Society Symposium on Research in Security and Privacy*, 1998.
- [40] T.Y.C. Woo and S.S. Lam. Authentication for distributed systems. *Computer*, 25(1):39–52, 1992.
- [41] T.Y.C. Woo and S.S. Lam. A semantic model for authentication protocols. In *IEEE Symposium on Security and Privacy*, pages 178–194, 1993.