

Validating a Web Service Security Abstraction by Typing

Andrew D. Gordon¹ and Riccardo Pucella²

¹Microsoft Research

²Cornell University

Abstract. An XML web service is, to a first approximation, an RPC service in which requests and responses are encoded in XML as SOAP envelopes, and transported over HTTP. We consider the problem of authenticating requests and responses at the SOAP-level, rather than relying on transport-level security. We propose a security abstraction, inspired by earlier work on secure RPC, in which the methods exported by a web service are annotated with one of three security levels: none, authenticated, or both authenticated and encrypted. We model our abstraction as an object calculus with primitives for defining and calling web services. We describe the semantics of our object calculus by translating to a lower-level language with primitives for message passing and cryptography. To validate our semantics, we embed correspondence assertions that specify the correct authentication of requests and responses. By appeal to the type theory for cryptographic protocols of Gordon and Jeffrey's Cryptyc, we verify the correspondence assertions simply by typing. Finally, we describe an implementation of our semantics via custom SOAP headers.

Keywords: Web services, remote procedure call, authentication, type systems

1. Introduction

It is common to provide application-level developers with security abstractions that hide detailed implementations at lower levels of a protocol stack. For example, the identity of the sender of a message may be exposed directly at the application-level, but computed via a hidden, lower level cryptographic protocol. The purpose of this paper is to explore how to build formal models of such security abstractions, and how to validate their correct implementation in terms of cryptographic primitives. Our setting is an experimental implementation of SOAP security headers for XML web services.

Correspondence and offprint requests to: Andrew D. Gordon, Microsoft Research, Roger Needham Building, 7 J J Thomson Ave, Cambridge CB3 0FB, UK. E-mail: adg@microsoft.com

1.1. Motivation: Web Services and SOAP

A crisp definition, due to the builders of the TerraService.NET service, is that “a web service is a web site intended for use by computer programs instead of human beings” [8]. Each request to or response from a web service is encoded in XML as a SOAP *envelope* [12]. An envelope consists of a *header*, containing perhaps routing or security information, and a *body*, containing the actual data of the request or response. A promising application for web services is to support direct retrieval of XML documents from remote databases, without resorting to unreliable “screen scraping” of data from HTML pages. For example, Google already offers programmatic access to its database via a web service [20]. Another major application is to support systems interoperability within an enterprise’s intranet.

The interface exported by a web service can be captured as an XML-encoded service description, in WSDL format [14], that describes the methods—and the types of their arguments and results—that make up the service. Tools exist for application-level developers to generate a WSDL description from the code of a service, and then to generate proxy code for convenient client access to the web service. Like tools for previous RPC mechanisms, these tools abstract from the details of the underlying messaging infrastructure. They allow us to regard calling a web service, for many if not all purposes, as if it were invoking a method on a local object. Our goal is to augment this abstraction with security guarantees.

There are many signs of fervour over web services: there is widespread tool support from both open source and commercial software suppliers, and frequent news of progress of web service standards at bodies such as OASIS and the W3C. Many previous systems support RPC, but one can argue that what’s new about web services is their combination of vendor-neutral interoperability, internet-scale, and toolsets for “mere mortals” [8]. Still, there are some reasons for caution. The XML format was not originally designed for messaging; it allows for interoperability but is inefficient compared to binary encodings. Moreover, it would be useful to use web services for inter-organisational communication, for example, for e-commerce, but SOAP itself does not define any security mechanisms.

In fact, there is already wide support for security at the transport-level, that is, for building secure web services using HTTPS and SSL. Still, SSL encrypts all traffic between the client and the web server, so that it is opaque to intermediaries. Hence, messages cannot be monitored by firewalls and cannot be forwarded by intermediate untrusted SOAP-level routers. There are proposals to avoid some of these difficulties by placing security at the SOAP-level, that is, by partially encrypting SOAP bodies and by including authenticators, such as signatures, in SOAP headers. In particular, the WS-Security [6] specification describes an XML syntax for including such information in SOAP envelopes.

Hence, the immediate practical goal of this work is to build and evaluate an exploratory system for SOAP-level security.

1.2. Background: Correspondences and Spi

Cryptographic protocols, for example, protocols for authenticating SOAP messages, are hard to get right. Even if we assume perfect cryptography, exposure to various replay and impersonation attacks may arise because of flaws in message formats. A common and prudent procedure is to invite expert analysis of any protocol, rather than relying on security through obscurity. Moreover, it is a useful discipline to specify and verify protocol goals using formal notations. Here, we specify authenticity goals of our protocol using Woo and Lam’s correspondence assertions [37], and verify them, assuming perfect cryptography in the sense of Dolev and Yao [17], using type theories developed as part of the Cryptyc project [22, 23, 21].

Woo and Lam’s correspondence assertions [37] are a simple and precise method for specifying authenticity properties. The idea is to specify labelled events that mark progress through the protocol. There are two kinds: begin-events and end-events. The assertion is that every end-event should correspond to a distinct, preceding begin-event with the same label. For example, Alice performs a begin-event with label “Alice sending Bob message M ” at the start of a session when she intends to send M to Bob. Upon receiving M and once convinced that it actually comes from Alice, Bob performs an end-event with the same label. If the correspondence assertion can be falsified, Bob can be manipulated into thinking a message comes from Alice when in fact it has been altered, or came from someone else, or is a replay. On the other hand, if the correspondence assertion holds, such attacks are ruled out.

There are several techniques for formally specifying and verifying correspondence assertions. Here, we model SOAP messaging within a process calculus, and model correspondence assertions by begin- and end-

statements within the calculus. We use a form of the spi-calculus [22], equipped with a type and effect system able to prove by typechecking that correspondence assertions hold in spite of an arbitrary attacker. Spi [5] is a small concurrent language with primitives for message passing and cryptography, derived from the π -calculus [32].

1.3. Contributions of this Paper

Our approach is as follows:

- Section 2 describes our high-level abstraction for secure messaging.
- Section 3 models the abstraction as an object calculus with primitives for creating and calling web services.
- Section 4 defines the semantics of our abstraction by translating to the spi-calculus. Correspondence assertions specify the authenticity guarantees offered to caller and callee, and are verified by typechecking.
- Section 5 describes a SOAP-based implementation using Visual Studio .NET.
- Section 6 shows how we can accommodate public-key infrastructures to implement the abstraction of Section 2.

Our main innovation is the idea of formalizing the authentication guarantees offered by a security abstraction by embedding correspondence assertions in its semantics. On the other hand, our high-level abstraction is fairly standard, and is directly inspired by work on secure network objects [35]. Although the rather detailed description of our model and its semantics may seem complex, the actual cryptographic protocol is actually quite simple. Still, we believe our framework and its implementation are a solid foundation for developing more sophisticated protocols and their abstractions.

Many formal details, as well as the proofs of our formal results, have been relegated to the appendices. Specifically, Appendix A gives sample messages exchanged during web service method calls using our abstractions, Appendix B gives a formal description of our object calculus, Appendix C gives a formal definition of the spi-calculus used in the paper, Appendix D gives the proofs of our formal results, and Appendix E describes an extension of our object calculus to capture a form of first-class web services.

A part of this article, in preliminary form, appears as a conference paper [24].

2. A Security Abstraction

We introduce a security abstraction for web services, where the methods exported by a web service are annotated by one of three security levels:

None	unauthenticated call
Auth	authenticated call
AuthEnc	authenticated and encrypted call

A *call* from a client to a web service is made up of two messages, the *request* from the client to the web service, and the *response* from the web service to the client. The inspiration for the security levels, and the guarantees they provide, comes from SRC Secure Network Objects [35]. An authenticated web method call provides a guarantee of *integrity* (that the request that the service receives is exactly the one sent by the client and that the response that the client receives is exactly the one sent by the service as a response to this request) and *at-most-once semantics* (that the service receives the request most once, and that the client receives the response at most once). An authenticated and encrypted web method call provides all the guarantees of an authenticated call, along with a guarantee of *secrecy* (that an eavesdropper does not obtain any part of the method name, the arguments, or the results of the call).

We use the language C[#] to present our security abstraction. (There is nothing specific to C[#] in our approach, although the implementation we describe in this section and in Section 5 takes advantage of some features of the language.) In C[#], where users can specify *attributes* on various entities, our security annotations take the form of an attribute on web methods, that is, the methods exported by a web service. The attribute is written `[SecurityLevel(level)]`, where *level* is one of `None`, `Auth`, or `AuthEnc`. For example, consider a simple interface to a banking service, where `[WebMethod]` is an attribute used to indicate a method exported by a web service:

```

class BankingServiceClass {

    string callerid;

    [WebMethod] [SecurityLevel(Auth)]
    public int Balance (int account);

    [WebMethod] [SecurityLevel(AuthEnc)]
    public string Statement (int account);

    [WebMethod] [SecurityLevel(Auth)]
    public void Transfer (int source,
                          int dest,
                          int amount);
}

```

The annotations get implemented by code to perform the authentication and encryption, at the level of SOAP envelopes, transparently from the user. The annotations on the web service side will generate a method on the web service that can be used to establish a security context. This method will never be invoked by the user, but automatically by the code implementing the annotations. For the purpose of this paper, we assume a simple setting for authentication and secrecy, namely that the principals involved possess shared keys. Specifically, we assume a distinct key K_{pq} shared between every pair of principals p and q . We use the key K_{pq} when p acts as the client and q as the web service. (Notice that K_{pq} is different from K_{qp} .) It is straightforward to extend our approach to different settings such as public-key infrastructures or certificate-based authentication mechanisms (see Section 6).

An authenticated call by p to a web method ℓ on a web service w owned by q with arguments u_1, \dots, u_n producing a result r uses the following protocol:

$$\begin{aligned}
 & p \rightarrow q : \text{request nonce} \\
 & q \rightarrow p : n_q \\
 & p \rightarrow q : p, \text{req}(w, \ell(u_1, \dots, u_n), s, n_q), n_p, \text{Hash}(\text{req}(w, \ell(u_1, \dots, u_n), s, n_q), K_{pq}) \\
 & q \rightarrow p : q, \text{res}(w, \ell(r), s, n_p), \text{Hash}(\text{res}(w, \ell(r), s, n_p), K_{pq})
 \end{aligned}$$

Here, Hash is a cryptographic hash function (a one-way message digest function such as MD5). We tag the request and the response messages to be able to differentiate them. We also tag the response with the name of the method that was originally called. We include a unique *session tag* s in both the request and response message to allow the caller p to match the response with the actual call that was performed.

An authenticated and encrypted call by p to a web method ℓ on a web service w owned by q with arguments u_1, \dots, u_n producing a result r uses a similar protocol, with the difference that the third and fourth messages are encrypted using the shared key instead of signed:

$$\begin{aligned}
 & p \rightarrow q : \text{request nonce} \\
 & q \rightarrow p : n_q \\
 & p \rightarrow q : p, \{\text{req}(w, \ell(u_1, \dots, u_n), s, n_q)\}_{K_{pq}}, n_p \\
 & q \rightarrow p : q, \{\text{res}(w, \ell(r), s, n_p)\}_{K_{pq}}
 \end{aligned}$$

To convince ourselves that the above protocols do enforce the guarantees prescribed by the security abstraction, we typically argue as follows. Let's consider the authenticated and encrypted case, the authenticated case being similar. When the web service w run by principal q receives a request $w, \ell(u_1, \dots, u_n), s, n_q$ encrypted with K_{pq} (q uses the identity p in the request to determine which key to use), it knows that only p could have created the message, assuming that the shared key K_{pq} is kept secret by both p and q . This enforces the integrity of the request. Since the message also contains the nonce n_q that the web service can check has never appeared in a previous message, it knows that the message is not a replayed message, hence enforcing at-most-once semantics. Finally, the secrecy of the shared key K_{pq} implies the secrecy of the request. A similar argument shows that the protocol satisfies integrity, at-most-once-semantics, and secrecy for the response.

What do we have at this point? We have an informal description of a security abstraction, we have an implementation of the abstraction in terms of protocols, and an informal argument that the guarantees

prescribed by the abstraction are enforced by the implementation. How do we make our security abstraction precise, and how do we ensure that the protocols do indeed enforce the required guarantees? In the next section, we give a formal model to make the abstraction precise. Then, we formalize the implementation by showing how to translate the abstractions into a lower level calculus that uses the above protocols. We use types to show that guarantees are formally met by the implementation, via correspondence assertions.

3. A Formal Model

We model the application-level view of authenticated messaging as an object calculus. Object calculi [1, 25, 29] are object-oriented languages in miniature, small enough to make formal proofs feasible, yet large enough to study specific features. As in FJ [29], objects are typed, class-based, immutable, and deterministic. As in some of Abadi and Cardelli’s object calculi [1], we omit subtyping and inheritance for the sake of simplicity. In spite of this simplicity, our calculus is Turing complete. We can define classes to implement arithmetic, lists, collections, and so on.

To model web services, we assume there are finite sets $Prin$ and $WebService$ of principal identifiers and web service identifiers, respectively. We think of each $w \in WebService$ as a URL referring to the service; moreover, $class(w)$ is the name of the class that implements the service, and $owner(w) \in Prin$ is the principal running the service.

To illustrate this model, we express the banking service interface introduced in the last section in our calculus. Suppose there are two principals $Alice, Bob \in Prin$, and a web service $w = http://bob.com/BankingService$, where we have $owner(w) = Bob$ and $class(w) = BankingServiceClass$. Suppose we wish to implement the *Balance* method so that given an account number, it checks that it has been called by the owner of the account, and if so returns the balance. If *Alice*’s account number is 12345, we might achieve this as follows:

```
class BankingServiceClass
  Id CallerId
  Num Balance(Num account)
  if account = 12345 then
    if this.CallerId = Alice then 100 else null
  else ...
```

There are a few points to note about this code. First, as in BIL [25], method bodies conform to a single applicative syntax, rather than there being separate grammars for statements and expressions. Second, while the C# code relies on attributes to specify exported methods and security levels, there are not attributes in our calculus. For simplicity, we assume that all the methods of a class implementing a web service are exported as web methods. Furthermore, we assume that all these exported methods are authenticated and encrypted, as if they had been annotated **AuthEnc**. (It is straightforward to extend our calculus to allow per-method annotations but it complicates the presentation of the translation in the next section.)

Every class implementing a web service has exactly one field, named *CallerId*, which exposes the identity of the caller, and allows application-level authorisation checks.

We write $w.Balance(12345)$ for a client-side call to method *Balance* of the service w . The semantics of such a web service call by *Alice* to a service owned by *Bob* is that *Bob* evaluates the local call $new\ BankingServiceClass(Alice).Balance(12345)$ as *Bob*. In other words, *Bob* creates a new object of the form $new\ BankingServiceClass(Alice)$ (that is, an instance of the class *BankingServiceClass* with *CallerId* set to *Alice*) and then calls the *Balance* method. This would terminate with 100, since the value of $this.CallerId$ is *Alice*. (For simplicity, we assume every class in the object calculus has a single constructor whose arguments are the initial values of the object’s fields.) This semantics guarantees to the server *Bob* that the field *CallerId* contains the identity of his caller, and guarantees to the client *Alice* that only the correct owner of the service receives the request and returns the result.

In a typical environment for web services, a client will not invoke web services directly. Rather, a client creates a proxy object corresponding to the web service, which encapsulates the remote invocations. Those proxy objects are generally created automatically by the programming environment. Proxy objects are easily expressible in our calculus, by associating with every web service w a proxy class $proxy(w)$. The class $proxy(w)$ has a method for every method of the web service class, the implementation for which simply calls the corresponding web service method. The proxy class also has a field *Id* holding the identity of the owner

of the web service. Here is the client-side proxy class for our example service:

```
class BankingServiceProxy
  Id Id()
  Bob
  Num Balance(Num account)
  w:Balance(account)
```

The remainder of this section details the syntax and informal semantics of our object calculus.

3.1. Syntax

In addition to *Prin* and *WebService*, we assume finite sets *Class*, *Field*, *Meth* of class, field, and method names, respectively.

Classes, Fields, Methods, Principals, Web Services:

$c \in \text{Class}$	class name
$f \in \text{Field}$	field name
$\ell \in \text{Meth}$	method name
$p \in \text{Prin}$	principal name
$w \in \text{WebService}$	web service name

There are two kinds of data type: *Id* is the type of principal identifiers, and $c \in \text{Class}$ is the type of instances of class *c*. A method signature specifies the types of its arguments and result.

Types and Method Signatures:

$A, B \in \text{Type} ::=$	type
Id	principal identifier
c	object
$sig \in \text{Sig} ::= B(A_1 x_1, \dots, A_n x_n)$	method signature (x_i distinct)

An execution environment defines the services and code available in the distributed system. In addition to *owner* and *class*, described above, the maps *fields* and *methods* specify the types of each field and the signature and body of each method, respectively. We write $X \rightarrow Y$ and $X \xrightarrow{\text{fin}} Y$ for the sets of total functions and finite maps, respectively, from *X* to *Y*.

Execution Environment: (*fields*, *methods*, *owner*, *class*)

$fields \in \text{Class} \rightarrow (\text{Field} \xrightarrow{\text{fin}} \text{Type})$	fields of a class
$methods \in \text{Class} \rightarrow (\text{Meth} \xrightarrow{\text{fin}} \text{Sig} \times \text{Body})$	methods of a class
$owner \in \text{WebService} \rightarrow \text{Prin}$	service owner
$class \in \text{WebService} \rightarrow \text{Class}$	service implementation

We complete the syntax by giving the grammars for *method bodies* and for *values*.

Values and Method Bodies:

x, y, z	name: variable, argument
$u, v \in \text{Value} ::=$	value
x	variable
$null$	null
$new\ c(v_1, \dots, v_n)$	object
p	principal identifier
$a, b \in \text{Body} ::=$	method body
v	value
$let\ x=a\ in\ b$	let-expression
$if\ u = v\ then\ a\ else\ b$	conditional

$v.f$	field lookup
$v.\ell(u_1, \dots, u_n)$	method call
$w:\ell(u_1, \dots, u_n)$	service call

The free variables $fv(a)$ of a method body are defined in the usual way, where the only binder is x being bound in b in the expression $let\ x=a\ in\ b$. We write $a\{x\leftarrow b\}$ for the outcome of a capture-avoiding substitution of b for each free occurrence of the variable x in method body a . We view method bodies as being equal up to renaming of bound variables. Specifically, we take $let\ x=a\ in\ b$ to be equal to $let\ x'=a\ in\ b\{x\leftarrow x'\}$, if $x' \notin fv(b)$.

Our syntax for bodies is in a reduced form that simplifies its semantics; in examples, it is convenient to allow a more liberal syntax. For instance, let the term $if\ a_1 = a_2\ then\ b_1\ else\ b_2$ be short for $let\ x_1=a_1\ in\ let\ x_2=a_2\ in\ if\ x_1 = x_2\ then\ b_1\ else\ b_2$. We already used this when writing $if\ this.CallerId = Alice\ then\ 100\ else\ null$ in our example. Similarly, we assume a class *Num* for numbers, and write integer literals such as 100 as shorthand for objects of that class.

Although objects are values, in this calculus, web services are not. This reflects the fact that current WSDL does not allow for web services to be passed as requests or results. We explore an extension of our model to account for web services as “first-class values” in Appendix E.

We assume all method bodies in our execution environment are well-typed. If $methods(c)(\ell) = (sig, b)$ and the signature $sig = B(A_1\ x_1, \dots, A_n\ x_n)$ we assume that the body b has type B given a typing environment $this:c, x_1:A_1, \dots, x_n:A_n$. The variable *this* refers to the object on which the ℓ method was invoked. The type system is given by a typing judgment $E \vdash a : A$, saying that a has type A in an environment E of the form $x_1:A_1, \dots, x_n:A_n$ that gives a type to the free variables in a . The domain $dom(E)$ of E is the set of variables $\{x_1, \dots, x_n\}$ given a type in E . The typing rules, which are standard, are given in Appendix B. We also assume the class $class(w)$ corresponding to each web service w has a single field *callerid*.

3.2. Informal Semantics of our Model

We explain informally the outcome of evaluating a method body b as principal p , that is, on a client or server machine controlled by p . (Only the semantics of web service calls depend on p .) A formal account of this semantics, as well as the typing rules of the calculus, can be found in Appendix B.

To evaluate a value v as p , we terminate at once with v itself.

To evaluate a let-expression $let\ x=a\ in\ b$ as p , we first evaluate a as p . If a terminates with a value v , we proceed to evaluate $b\{x\leftarrow v\}$, that is, b with each occurrence of the variable x replaced with v . The outcome of evaluating $b\{x\leftarrow v\}$ as p is the outcome of evaluating the whole expression.

To evaluate a conditional $if\ u = v\ then\ a\ else\ b$ as p , we evaluate a as p if u and v are the same; else we evaluate b as p .

To evaluate a field lookup $v.f$ as p , when v is an object value $new\ c(v_1, \dots, v_n)$, we check f is the j th field of class c for some $j \in 1..n$ (that is, that $fields(c) = f_i \mapsto A_i\ i \in 1..n$ and that $f = f_j$), and then return v_j . If v is null or if the check fails, evaluation has gone wrong.

To evaluate a method call $v.\ell(u_1, \dots, u_n)$ as p , when v is an object $new\ c(v_1, \dots, v_n)$, we check ℓ is a method of class c (that is, that $methods(c) = \ell_i \mapsto (sig_i, b_i)\ i \in 1..m$ and that $\ell = \ell_j$ for some $j \in 1..m$) and we check the arity of its signature is n (that is, that $sig_j = B(A_1\ x_1, \dots, A_n\ x_n)$) and then we evaluate the method body as p , but with the object v itself in place of the variable *this*, and actual parameters u_1, \dots, u_n in place of the formal parameters x_1, \dots, x_n (that is, we evaluate the expression $b_j\{this\leftarrow v, x_1\leftarrow u_1, \dots, x_n\leftarrow u_n\}$). If v is null or if either check fails, evaluation has gone wrong.

To evaluate a service call $w:\ell(u_1, \dots, u_n)$ as p , we evaluate the local method call $new\ c(p).\ell(u_1, \dots, u_n)$ as q , where $c = class(w)$ is the class implementing the service, and $q = owner(w)$ is the principal owning the service. (By assumption, c 's only field is *CallerId* of type *Id*.) This corresponds directly to creating a new object on q 's web server to process the incoming request.

4. A Spi-Calculus Semantics

We confer a formal semantics on our calculus by translation to the spi-calculus [5, 22], a lower-level language with primitives for message-passing (to model SOAP requests and responses) and cryptography (to model encryption and decryption of SOAP headers and bodies).

4.1. A Typed Spi-Calculus (Informal Review)

To introduce the spi-calculus, we formalize the situation where Alice sends a message to Bob using a shared key, together with a correspondence assertion concerning authenticity of the message, as outlined in Section 1. A *name* is an identifier that is atomic as far as our analysis is concerned. In this example, the names *Alice* and *Bob* identify the two principals, the name *K* represents a symmetric key known only to *Alice* and *Bob*, and the name *n* represents a public communication channel. A *message*, *M* or *N*, is a data structure such as a name, a tuple (M_1, \dots, M_n) , a tagged message $t(M)$, or a ciphertext $\{M\}_N$ (that is, a message *M* encrypted with a key *N*, which is typically a name). A *process*, *P* or *Q*, is a program that may perform local computations such as encryptions and decryptions, and may communicate with other processes by message-passing on named channels. For example, the process $P_{Alice} = \text{begin } \text{sending}(Alice, Bob, M); \text{out } n \{M\}_K$ defines Alice's behaviour. First, she performs a begin-event labelled by the tagged tuple $\text{sending}(Alice, Bob, M)$, and then she sends the ciphertext $\{M\}_K$ on the channel *n*. The process $P_{Bob} = \text{in } n(x); \text{decrypt } x \text{ is } \{y\}_K; \text{end } \text{sending}(Alice, Bob, y)$ defines Bob's behaviour. He blocks till a message *x* arrives on the channel *n*. Then he attempts to decrypt the message with the key *K*. We assume there is sufficient redundancy, such as a checksum, in the ciphertext that we can tell whether it was encrypted with *K*. If so, the plaintext message is bound to *y*, and he performs an end-event labelled $\text{sending}(Alice, Bob, y)$. The process $\text{new}(K); (P_{Alice} \mid P_{Bob})$ defines the complete system. The composition $P_{Alice} \mid P_{Bob}$ represents Alice and Bob running in parallel, and able to communicate on shared channels such as *n*. The binder $\text{new}(K)$ restricts the scope of the key *K* to the process $P_{Alice} \mid P_{Bob}$ so that no external process may use it. Appendix C contains the grammar of spi messages and processes. The grammar includes the type annotations that are required to appear in spi terms.

We include begin- and end-events in processes simply to specify correspondence assertions. We say a process is *safe* to mean that in every run, and for every *L*, there is a distinct, preceding **begin** *L* event for every **end** *L* event. Our example is safe, because Bob's end-event can only happen after Alice's begin-event.

For correspondence assertions to be interesting, we need to model the possibility of malicious attacks. Let an *opponent* be a spi-calculus process *O*, arbitrary except that *O* itself cannot perform begin- or end-events. We say a process *P* is *robustly safe* if and only if $P \mid O$ is safe for every opponent *O*. Our example system $\text{new}(K); (P_{Alice} \mid P_{Bob})$ is not robustly safe. The opponent cannot acquire the key *K* since its scope is restricted, but it can intercept messages on the public channel *n* and mount a replay attack. The opponent $\text{in } n(x); \text{out } n x; \text{out } n x$ duplicates the encrypted message so that Bob may mistakenly accept *M* and perform the end-event $\text{sending}(Alice, Bob, M)$ twice. To protect against replays, and to achieve robust safety, we can add a nonce handshake to the protocol.

In summary, spi lets us precisely represent the behaviour of protocol participants, and specify authenticity guarantees by process annotations. Robust safety is the property that no opponent at the level of the spi-calculus may violate these guarantees. We omit the details here, but a particular type and effect system verifies robust safety: if a process can be assigned the empty effect, then it is robustly safe. The example above is simple, but the general method works for a wide range of protocol examples [22, 21].

For the sake of clarity, we defer some of the technical details to the appendices. Specifically, Appendix C contains more details on the spi-calculus and the type and effect system, as well as a formal definition of robust safety; Appendix D gives a proof of our technical results.

4.2. A Semantics for Local Computation

We translate the types, values, and method bodies of our object calculus to types, messages, and processes, respectively, of the spi calculus. To begin with, we omit web services. Many computational models can be studied by translation to process calculi; our translation of local computation follows a fairly standard pattern.

We use the notation $\llbracket \cdot \rrbracket$ to represent the translation of the types and terms of our object calculus to appropriate types, messages, and processes in the spi calculus. In many places, we also define abbreviations in the spi calculus (for instance, we define $\text{let } x = \text{call}_w(p, \text{args}); P$ as shorthand for a more complex spi calculus process); these do not use the $\llbracket \cdot \rrbracket$ notation.

We assume that Prin are spi-calculus names, and that $\text{Field} \cup \text{Meth} \cup \text{Class} \cup \{\text{null}\}$ are message tags. The translations for types is straightforward. Since principal identifiers are presumably known to the opponent, the type of identifiers corresponds to the spi type Un . A value of class c is either the value null , or a tagged tuple $\text{new } c(v_1, \dots, v_n)$. As we shall see below, we translate null to a tagged empty tuple $\text{null}()$, and an object to a tagged tuple $c(v_1, \dots, v_n)$. Thus, a class c translates to a tagged union type with components $\text{null}(\text{Un})$ and $c(\text{Un})$. (The types Un indicate that the content of the tuples are presumably known to the opponent.)

Type Translation:

$$\begin{aligned} \text{Prin} &\triangleq \text{Un} \\ \llbracket \text{Id} \rrbracket &\triangleq \text{Prin} \\ \llbracket c \rrbracket &\triangleq \text{Union}(\text{null}(\text{Un}), c(\text{Un})) \end{aligned}$$

Environment Translation:

$$\llbracket x_1:A_1, \dots, x_n:A_n \rrbracket \triangleq x_1:\llbracket A_1 \rrbracket, \dots, x_n:\llbracket A_n \rrbracket$$

If $\text{As} = A_1, \dots, A_n$ and $\text{xs} = x_1, \dots, x_n$ we sometimes write $B(\text{As } \text{xs})$ as shorthand for the signature $B(A_1 x_1, \dots, A_n x_n)$. We define two shorthands for types corresponding to web method calls. The type $\text{Req}(w)$ represents the type of possible calls to web methods provided by the service w ; the type of a call is simply the translated type of the arguments of the web method, tagged with the name of the method. Similarly, the type $\text{Res}(w)$ represents the type of the results of web methods provided by the service w ; the type of a result of a call is simply the translated type of the result of the web method, tagged once again with the name of the method.

Request and Response Types:

$$\begin{aligned} \llbracket A_1, \dots, A_m \rrbracket &\triangleq \llbracket A_1 \rrbracket, \dots, \llbracket A_m \rrbracket \\ \text{Req}(w) &\triangleq \text{Union}(\ell_i(\llbracket \text{As}_i \rrbracket)_{i \in 1..n}) \\ &\quad \text{where } \text{class}(w) = c \text{ and } \text{methods}(c) = \ell_i \mapsto (B_i(\text{As}_i \text{xs}_i), b_i)_{i \in 1..n} \\ \text{Res}(w) &\triangleq \text{Union}(\ell_i(\llbracket B_i \rrbracket)_{i \in 1..n}) \\ &\quad \text{where } \text{class}(w) = c \text{ and } \text{methods}(c) = \ell_i \mapsto (B_i(\text{As}_i \text{xs}_i), b_i)_{i \in 1..n} \end{aligned}$$

The translation of expressions really acts on the type derivation of an expression, not just the expression itself. This means that during the translation of an expression, we have access to the types of the subexpressions appearing in the expression. To reduce clutter, we write the translation as though it is acting on the expression itself, except that when we need access to the type of a subexpression, we annotate the appropriate subexpression with its type. For example, the translation of $\text{let } x = a \text{ in } b$ depends on the type of a , which is available through the type derivation of $E \vdash \text{let } x = a \text{ in } b : B$. We write $\text{let } x = a_A \text{ in } b$ to indicate that the type of a is A , according to the type derivation. Values translate easily; in particular, an object translates to a tagged tuple containing the values of its fields.

Translation of a Value v to a Message $\llbracket v \rrbracket$:

$$\begin{aligned} \llbracket x \rrbracket &\triangleq x \\ \llbracket \text{null} \rrbracket &\triangleq \text{null}() \\ \llbracket \text{new } c(v_1, \dots, v_n) \rrbracket &\triangleq c(\llbracket v_1 \rrbracket, \dots, \llbracket v_n \rrbracket) \\ \llbracket p \rrbracket &\triangleq p \end{aligned}$$

We translate a body b to a process $\llbracket b \rrbracket_k^p$ that represents the evaluation of b as principal p . The name k is a continuation, a communications channel on which we send $\llbracket v \rrbracket$ to represent termination with value v . Since our focus is representing safety rather than liveness properties, we represent an evaluation that goes wrong

simply by the inactive process **stop**; it would be easy—but a complication—to add an exception mechanism. We use standard **split** and **case** statements to analyse tuples and tagged messages, respectively. To call a method ℓ of an object v of class c , with arguments u_1, \dots, u_n we send the tuple $(p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k)$ on the channel $c.\ell$. The name p is the caller, and channel k is the continuation for the call. We translate method ℓ of class c to a process that repeatedly awaits such messages, and triggers evaluations of its body. We defer the translation of web method calls until Section 4.3. Our translation depends in part on type information; we write v_c in the translation of field lookups and method calls to indicate that c is the type of v .

Translation of a Method Body b to a Process $\llbracket b \rrbracket_k^p$:

$$\begin{aligned} \llbracket v \rrbracket_k^p &\triangleq \text{out } k \llbracket v \rrbracket \\ \llbracket \text{let } x = a_A \text{ in } b \rrbracket_k^p &\triangleq \text{new } (k':\text{Un}); (\llbracket a \rrbracket_{k'}^p \mid \text{in } k' (x:\text{Un}); \llbracket b \rrbracket_k^p) \\ \llbracket \text{if } u = v \text{ then } a \text{ else } b \rrbracket_k^p &\triangleq \text{if } \llbracket u \rrbracket = \llbracket v \rrbracket \text{ then } \llbracket a \rrbracket_k^p \text{ else } \llbracket b \rrbracket_k^p \\ \llbracket v_c.f_j \rrbracket_k^p &\triangleq \text{case } \llbracket v \rrbracket \text{ is } \text{null}(y:\text{Un}); \text{stop} \\ &\quad \text{is } c(y:\text{Un}); \text{split } y \text{ is } (x_1:\llbracket A_1 \rrbracket, \dots, x_n:\llbracket A_n \rrbracket); \text{out } k \ x_j \\ &\quad \text{where } \text{fields}(c) = f_i \mapsto A_i \text{ } i \in 1..n, \text{ and } j \in 1..n \\ \llbracket v_c.\ell(u_1, \dots, u_n) \rrbracket_k^p &\triangleq \text{case } \llbracket v \rrbracket \text{ is } \text{null}(y:\text{Un}); \text{stop} \\ &\quad \text{is } c(y:\text{Un}); \text{out } c.\ell (p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k) \end{aligned}$$

Translation of Method ℓ of Class c :

$$\begin{aligned} I_{\text{class}}(c, \ell) &\triangleq \text{repeat in } c.\ell (z:\text{Un}); \\ &\quad \text{split } z \text{ is } (p:\text{Prin}, \text{this}:\text{Un}, x_1:\llbracket A_1 \rrbracket, \dots, x_n:\llbracket A_n \rrbracket, k:\text{Un}); \llbracket b \rrbracket_k^p \\ &\quad \text{where } \text{methods}(c)(\ell) = (B(A_1 \ x_1, \dots, A_n \ x_n), b) \end{aligned}$$

4.3. A Semantics for Web Services

We complete the semantics for our object calculus by translating our cryptographic protocol for calling a web service to the spi-calculus. A new idea is that we embed begin- and end-events in the translation to represent the abstract authenticity guarantees offered by the object calculus.

We assume access to all web methods is at the highest security level **AuthEnc** from Section 2, providing both authentication and secrecy. Here is the protocol, for p making a web service call $w:\ell(u_1, \dots, u_n)$ to service w owned by q , including the names of continuation channels used at the spi level. Recall that the protocol assumes that the client has a way to query the web service for a nonce. Therefore, we assume that in addition to the methods of $\text{class}(w)$, each web service also supports a method getnonce , which we implement specially.

$$\begin{aligned} p &\rightarrow q \text{ on } w : \text{req}(\text{getnonce}()), k_1 \\ q &\rightarrow p \text{ on } k_1 : \text{res}(\text{getnonce}(n_q)) \\ p &\rightarrow q \text{ on } w : p, \{\text{req}(w, \ell(u_1, \dots, u_n), t, n_q)\}_{K_{pq}}, n_p, k_2 \\ q &\rightarrow p \text{ on } k_2 : q, \{\text{res}(w, \ell(r), t, n_p)\}_{K_{pq}} \end{aligned}$$

We are assuming there is a shared key K_{pq} for each pair of principals $p, q \in \text{Prin}$. For the sake of brevity, we omit the formal description of the type and effect system [21] we rely on, but see Appendix C for a detailed overview. Still, to give a flavour, we can define the type of a shared key K_{pq} as follows:

Type of Key Shared Between Client p and Server q :

$$\begin{aligned} \text{CSKey}(p, q) &\triangleq \\ &\text{SharedKey}(\text{Union}(\\ &\quad \text{req}(w:\text{Un}, a:\text{Un}, t:\text{Un}, \\ &\quad \quad n_q:\text{Public Response} [\text{end } \text{req}(p, q, w, a, t)]), \\ &\quad \text{res}(w:\text{Un}, r:\text{Un}, t:\text{Un}, \\ &\quad \quad n_p:\text{Public Response} [\text{end } \text{res}(p, q, w, r, t)])) \end{aligned}$$

The type says we can use the key in two modes. First, we may encrypt a plaintext tagged req containing

four components: a public name w of a service, an argument a suitable for the service, a session tag t , and a nonce n_q proving that a begin-event labelled $req(p, q, w, a, t)$ has occurred, and therefore that an end-event with that label would be safe. Second, we may encrypt a plaintext tagged res containing four components: a service w , a result r from that service, the session tag t , and a nonce n_p proving that a begin-event labelled $res(p, q, w, r, t)$ has occurred.

We translate a service call to the client-side of our cryptographic protocol as follows. We start by embedding a begin-event labelled $req(p, q, w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t)$ to record the details of client p 's call to server $q = owner(w)$. We request a nonce n_q , and use it to freshen the encrypted request, which we send with our own nonce n_p , which the server uses to freshen its response. If the response indeed contains our nonce, we embed an end-event to record successful authentication. For the sake of brevity, we rely on some standard shorthands for pattern-matching.

Translation of Web Method Call:

```

 $\llbracket w:\ell(u_1, \dots, u_n) \rrbracket_k^p \triangleq$ 
  new ( $k_1:\text{Un}, k_2:\text{Un}, t:\text{Un}, n_p:\text{Public Challenge} []$ );
  begin  $req(p, q, w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t)$ ;
  out  $w$  ( $req(getnonce()), k_1$ );
  inp  $k_1$  ( $res(getnonce(n_q:\text{Un}))$ );
  cast  $n_q$  is ( $n'_q:\text{Public Response} [end\ req(p, q, w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t)]$ );
  out  $w$  ( $p, \{req(w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t, n'_q)\}_{K_{pq}}, n_p, k_2$ );
  inp  $k_2$  ( $q':\text{Un}, bdy:\text{Un}$ ); decrypt  $bdy$  is  $\{res(plain)\}_{K_{pq}}$ ;
  match  $plain$  is ( $w, rest:(r:\text{Res}(w), t':\text{Un}, \text{Public Response} [end\ res(p, q, w, r, t')])$ );
  split  $rest$  is ( $r:\text{Res}(w), rest':(t':\text{Un}, n'_p:\text{Public Response} [end\ res(p, q, w, r, t')])$ );
  match  $rest'$  is ( $t, n'_p:\text{Public Response} [end\ res(p, q, w, r, t)]$ );
  check  $n_p$  is  $n'_p$ ; end  $res(p, q, w, r, t)$ ; case  $r$  is  $\ell(x)$ ; out  $k$   $x$ 
  where  $q = owner(w)$ 

```

Our server semantics relies on a shorthand notation defined below; $\text{let } x = \text{call}_w(p, \ell(u_1, \dots, u_n)); P$ runs the method ℓ of the class $class(w)$ implementing the service w , with arguments u_1, \dots, u_n , and with its *CallerId* field set to p , binds the result to x and runs P .

Server-Side Invocation of a Web Method:

```

let  $x = \text{call}_w(p, args); P \triangleq$ 
  new ( $k$ );
  (case  $args$  (is  $\ell_i(xs_i)$ ; new ( $k'$ ); (out  $c\ell_i(q, c(p), xs_i, k')$  | inp  $k'$  ( $r$ ); out  $k$   $\ell_i(r)$ ))  $i \in 1..n$ 
  | inp  $k$  ( $x$ );  $P$ )
  where  $c = class(w)$ ,  $q = owner(w)$ ,
  and  $methods(c) = \ell_i \mapsto (B_i(As_i, xs_i), b_i) \ i \in 1..n$ 

```

Finally, we implement each service w by a process $I_{ws}(w)$. We repeatedly listen for nonce requests, reply with one, and then await a web service call freshened by the nonce. If we find the nonce, it is safe to perform an end-event labelled $req(p, q, w, a, t)$, where p is the caller, $q = owner(w)$ is the service owner, a is the received method request, and t is the session tag. We use the shorthand above to invoke a . If r is the result, we perform a begin-event labelled $res(p, q, w, r, t)$ to record we are returning a result, and then send a response, freshened with the nonce we received from the client. In general, the notation $\prod_{i \in 1..n} P_i$ means $P_1 \mid \dots \mid P_n$.

Web Service Implementation:

$$\begin{aligned}
I_{ws}(w) \triangleq & \text{repeat in } w \text{ (} bdy:\text{Un}, k_1:\text{Un}\text{);} \\
& \text{case } bdy \text{ is } req(\text{getnonce}()); \\
& \text{new } (n_q:\text{Public Challenge } []); \\
& \text{out } k_1 \text{ (} res(\text{getnonce}(n_q))\text{);} \\
& \text{in } w \text{ (} p':\text{Un}, cipher:\text{Un}, n_p:\text{Un}, k_2:\text{Un}\text{);} \\
& \prod_{p \in Prin} \text{if } p = p' \text{ then} \\
& \text{decrypt } cipher \text{ is } \{req(plain)\}_{K_{pq}}; \\
& \text{match } plain \text{ is } (w, rest: \\
& \quad (a:Req(w), t:\text{Un}, \text{Public Response } [\text{end } req(p, q, w, a, t)])); \\
& \text{split } rest \text{ is } (a:Req(w), t:\text{Un}, n'_q: \\
& \quad \text{Public Response } [\text{end } req(p, q, w, a, t)]); \\
& \text{check } n_q \text{ is } n'_q; \text{end } req(p, q, w, a, t); \\
& \text{let } r:Res(w) = \text{call}_w(p, a); \\
& \text{begin } res(p, q, w, r, t); \\
& \text{cast } n_p \text{ is } (n'_p:\text{Public Response } [\text{end } res(p, q, w, r, t)]); \text{out } k_2 \text{ (} q, \{res(w, r, t, n'_p)\}_{K_{pq}}\text{)} \\
& \text{where } q = \text{owner}(w)
\end{aligned}$$

This semantics is subject to more deadlocks than a realistic implementation, since we do not have a single database of outstanding nonces. Still, since we are concerned only with safety properties, not liveness, it is not a problem that our semantics is rather more nondeterministic than an actual implementation.

4.4. Security Properties of a Complete System

We define the process $Sys(b, p, k)$ to model a piece of code b being run by principal p (with continuation k) in the context of implementations of all the classes and web services in $Class$ and $WebService$. The implementation of the classes and web services are given as follows.

Implementation of Classes and Web Services:

$$\begin{aligned}
ClMeth & \triangleq \{(c, \ell) : c \in Class, \ell \in dom(methods(c))\} \\
I_{class} & \triangleq \prod_{(c, \ell) \in ClMeth} I_{class}(c, \ell) \\
I_{ws} & \triangleq \prod_{w \in WebService} I_{ws}(w)
\end{aligned}$$

The process $Sys(b, p, k)$ is defined with respect to an environment that specifies the type of its free variables, such as the names of the web services, principals, classes and methods, and keys.

Top-Level Environments:

$$\begin{aligned}
E_{class} & \triangleq (c.\ell:\text{Un})_{(c, \ell) \in ClMeth} \\
E_{keys} & \triangleq (K_{pq}:\text{CSKey}(p, q))_{p, q \in Prin} \\
E_{ws} & \triangleq (w:\text{Un})_{w \in WebService} \\
E_{prin} & \triangleq p_1:\text{Prin}, \dots, p_n:\text{Prin} \quad \text{where } Prin = \{p_1, \dots, p_n\} \\
E_0 & \triangleq E_{ws}, E_{prin}, E_{class}, E_{keys}
\end{aligned}$$

The process $Sys(b, p, k)$ is defined as follows:

$$Sys(b, p, k) \triangleq \text{new } (E_{class}, E_{keys}); (I_{class} \mid I_{ws} \mid \text{new } (k:\text{Un}); \llbracket b \rrbracket_k^p)$$

We claim that the ways an opponent O can interfere with the behaviour of $Sys(b, p, k)$ correspond to the ways in which an actual opponent lurking on a network could interfere with SOAP-level messages being routed between web servers. The names $c.\ell$ of methods are hidden, so O cannot interfere with calls to local methods. The keys K_{pq} are also hidden, so O cannot decrypt or fake SOAP-level encryption. On the other hand, the names w on which $Sys(b, p, k)$ sends and receives our model of SOAP envelopes are public, and so O is free to intercept, replay, or modify such envelopes.

Our main result is that an opponent cannot disrupt the authenticity properties embedded in our translation. The proof is by showing the translation preserves types.

Theorem 1. If $\emptyset \vdash b : B$ and $p \in \text{Prin}$ and $k \notin \text{dom}(E_0)$ then the system $\text{Sys}(b, p, k)$ is robustly safe.

Proof. See Appendix D.1. □

5. A SOAP-Level Implementation

We have implemented the security abstraction introduced in Section 2 and formalized in Sections 3 and 4 on top of the Microsoft Visual Studio .NET implementation of web services, as a library that web service developers and clients can use. A web service developer adds security attributes to the web methods of the service. The developer also needs to provide a web method to supply a nonce to the client. On the client side, the client writer is provided with a modified proxy class that encapsulates the implementation of the security abstraction and takes into account the security level of the corresponding web service methods. Hence, from a client's point of view, there is no fundamental difference between accessing a web service with security annotations and one without.

Consider an implementation of our running example of a banking service. Here is what (an extract of) the class implementing the web service looks like:

```
class BankingServiceClass : System.Web.Services.WebService
{
    ...
    [WebMethod]
    public int RequestNonce () { ... }

    public DSHeader header;

    [WebMethod]
    [SecurityLevel(Level=SecLevel.Auth)]
    [SoapHeader("header", Direction=Direction.InOut,Required=true)]
    public int Balance (int account) { ... }
}
```

This is the code we currently have, and it is close to the idealized interface we gave in Section 2. The differences are due to implementation restrictions imposed by the development environment. The extract shows that the web service implements the `RequestNonce` method required by the authentication protocol. The `Balance` method is annotated as an authenticated method, and is also annotated to indicate that the headers of the SOAP messages used during a call will be available through the `header` field of the interface. (The class `DSHeader` has fields corresponding to the headers of the SOAP message.) As we shall see shortly, SOAP headers are used to carry the authentication information. Specifically, the authenticated identity of the caller is available in a web method through `header.callerid`.

To implement the security abstraction on the web service side, we use a feature of Visual Studio .NET called SOAP Extensions. Roughly speaking, a SOAP Extension acts like a programmable “filter”. It can be installed on either (or both) of a client or a web service. It gets invoked on every incoming and outgoing SOAP message, and can be used to examine and modify the content of the message before forwarding it to its destination. In our case, the extension will behave differently according to whether the message is incoming or outgoing, and depending on the security level specified. For an outgoing message, if the security level is `None`, the SOAP message is unchanged. If the security level is `Auth`, messages are signed as specified by the protocol: a cryptographic hash of the SOAP body and the appropriate nonce is stored in a custom header of the messages. If the security level is `AuthEnc`, messages are encrypted as specified by the protocol, before being forwarded. For incoming messages, the messages are checked and decrypted, if required. If the security level is `Auth`, the signature of the message checked. If the security level is `AuthEnc`, the message is decrypted before being forwarded. Our implementation uses the SHA1 hash function for signatures, and the RC2 algorithm for symmetric encryption.

To implement the security abstraction on the client side, we provide the client with a new proxy class. The new proxy class provides methods `None`, `Auth`, and `AuthEnc`, that are called by the proxy methods to

initiate the appropriate protocol. The method `None` simply sets up the headers of the SOAP message to include the identity of the caller and the callee. `Auth` and `AuthEnc` do the same, but also make a call to the web service to get a nonce and add it (along with a newly created nonce) to the headers. The actual signature and encryption of the SOAP message is again performed using SOAP Extensions, just as on the web service side.

Our implementation uses a custom SOAP header `DSHeader` to carry information such as nonces, identities, and signatures. It provides the following elements:

<code>callerid</code>	identity of the client
<code>calleeid</code>	identity of the web service provider
<code>np</code>	client nonce
<code>nq</code>	web service nonce
<code>signature</code>	cryptographic signature of the message

Not all of those elements are meaningful for all messages. In addition to these headers, in the cases where the message is encrypted, the SOAP body is replaced by the encrypted body. Appendix A gives actual SOAP messages exchanged between the client and web service during an authenticated call to `Balance`, and an authenticated and encrypted call to `Statement`.

6. A Semantics Using Asymmetric Cryptography

The security abstraction we describe in Section 2 relies on shared keys between principals. This is hardly a reasonable setup in modern systems. In this section, we show that our approach can easily accommodate public-key infrastructures.

6.1. Authenticated Web Methods

We start by describing the protocol and implementation for authenticated web methods. Hence, for now, we assume that all the exported methods of a web service are annotated with `Auth`.

Consider a simple public-key infrastructure for digital signatures. Each principal p has a signing key SKp and a verification key VKp . The signing key is kept private, while the verification key is public. To bind the name of a principal with their verification key, we assume a *certification authority* CA (itself with a signing key $SKCA$ and verification key $VKCA$) that can sign certificates $CertVKp$ of the form $\{p, VKp\}_{SKCA}$. (The notation $\{\cdot\}_K$ is used to represent both asymmetric encryption and signature, differentiating it from symmetric encryption. In the case where $\{M\}_K$ represent a signature, this is simply notation for M along with a token representing the signature of M with asymmetric key K .)

Here is a protocol that uses digital signatures to authenticate messages, for p making a web service call $w:\ell(u_1, \dots, u_n)$ to service w owned by q , including the names of continuation channels used at the spi level. Again, we assume that in addition to the methods of $class(w)$, each web service also supports a method `getnonce`, which we implement specially.

$$\begin{aligned}
 & p \rightarrow q \text{ on } w : CertVKp, n_p, req(getnonce()), k_1 \\
 & q \rightarrow p \text{ on } k_1 : CertVKq, res(getnonce(n_q)) \\
 & p \rightarrow q \text{ on } w : p, \{req(w, \ell(u_1, \dots, u_n), t, q, n_q)\}_{SKp}, k_2 \\
 & q \rightarrow p \text{ on } k_2 : q, \{res(w, \ell(r), t, p, n_p)\}_{SKq}
 \end{aligned}$$

Type of Signing Keys:

$$\begin{aligned}
 AuthMsg(p) &\triangleq \\
 & \text{Union}(req(w:Un, a:Un, t:Un, q:Un, n_q:Public \text{ Response } [end \ req(p, q, w, a, t)]), \\
 & \quad \quad \quad res(w:Un, r:Un, t:Un, q:Un, n_q:Public \text{ Response } [end \ res(q, p, w, r, t)])) \\
 AuthKeys(p) &\triangleq \text{KeyPair}(AuthMsg(p)) \\
 AuthCert &\triangleq (p : Un, \text{Decrypt Key}(AuthMsg(p))) \\
 AuthCertKeys &\triangleq \text{KeyPair}(AuthCert)
 \end{aligned}$$

We will represent the key pair of a signing key and verification key for principal p by a pair DSp , of type

$\text{AuthKeys}(p)$. The key pair for the certification authority will be represented by a pair $DSCA$. We use the following abbreviations:

Key and Certificates Abbreviations:

$SKp \triangleq \text{Encrypt } (DSp)$	p 's signing key
$VKp \triangleq \text{Decrypt } (DSp)$	p 's verification key
$\text{CertVKp} \triangleq \{p, VKp\}_{SKCA}$	p 's certificate

With that in mind, we can amend the translation of Section 4 to accommodate the new protocol. First, we give a new translation for a web method call $w:\ell(u_1, \dots, u_n)$:

New Translation of Web Method Call:

```

 $\llbracket w:\ell(u_1, \dots, u_n) \rrbracket_k^p \triangleq$ 
  new  $(k_1:\text{Un}, k_2:\text{Un}, t:\text{Un}, n_p:\text{Public Challenge } [])$ ;
  begin  $req(p, q, w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t)$ ;
  out  $w$  ( $\text{CertVKp}, n_p, req(getnonce()), k_1$ );
  inp  $k_1$  ( $c:\text{Un}, res(getnonce(n_q:\text{Un}))$ );
  decrypt  $c$  is  $\{cert:(q':\text{Un}, \text{Decrypt Key}(\text{AuthMsg}(q')))\}_{VKCA^{-1}}$ ;
  match  $cert$  is  $(q, vkq:\text{Decrypt Key}(\text{AuthMsg}(q)))$ ;
  cast  $n_q$  is  $(n'_q:\text{Public Response } [end\ req(p, q, w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t)])$ ;
  out  $w$  ( $p, \{req(w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t, q, n'_q)\}_{SKp}, k_2$ );
  inp  $k_2$  ( $q'':\text{Un}, bdy:\text{Un}$ );
  decrypt  $bdy$  is  $\{res(plain:(w':\text{Un}, r:\text{Un}, t':\text{Un}, p':\text{Un},$ 
     $\text{Public Response } [end\ res(p', q, w', r, t')]))\}_{vkq^{-1}}$ ;
  match  $plain$  is  $(w, rest:(r:\text{Res}(w), t':\text{Un}, p':\text{Un},$ 
     $\text{Public Response } [end\ res(p', q, w, r, t')]))$ ;
  split  $rest$  is  $(r:\text{Res}(w), rest':(t':\text{Un}, p':\text{Un},$ 
     $\text{Public Response } [end\ res(p', q, w, r, t')]))$ ;
  match  $rest'$  is  $(t, rest'':(p':\text{Un}, \text{Public Response } [end\ res(p', q, w, r, t')]))$ ;
  match  $rest''$  is  $(p, n'_p:\text{Public Response } [end\ res(p, q, w, r, t)])$ ;
  check  $n_p$  is  $n'_p$ ;
  end  $res(p, q, w, r, t)$ ;
  case  $r$  is  $\ell(x)$ ; out  $k$   $x$ 
  where  $q = owner(w)$ 

```

We also need to give a new implementation for web services, again to take into account the different messages being exchanged:

New Web Service Implementation:

```

 $I_{ws}(w) \triangleq$ 
  repeat inp  $w$  ( $c:\text{Un}, n_p:\text{Un}, bdy:\text{Un}, k_1:\text{Un}$ );
  case  $bdy$  is  $req(getnonce())$ ;
  decrypt  $c$  is  $\{p:\text{Un}, vkp:\text{Decrypt Key}(\text{AuthMsg}(p))\}_{VKCA^{-1}}$ ;
  new  $(n_q:\text{Public Challenge } [])$ ;
  out  $k_1$  ( $\text{CertVKq}, res(getnonce(n_q))$ );
  inp  $w$  ( $p':\text{Un}, cipher:\text{Un}, k_2:\text{Un}$ );
  if  $p = p'$  then
  decrypt  $cipher$  is  $\{req(plain:(w:\text{Un}, a:\text{Un}, t:\text{Un}, q':\text{Un},$ 
     $\text{Public Response } [end\ req(p, q', w, a, t)]))\}_{vkp^{-1}}$ ;
  match  $plain$  is  $(w, rest:(a:\text{Req}(w), t:\text{Un}, q':\text{Un},$ 
     $\text{Public Response } [end\ req(p, q', w, a, t)]))$ ;

```

```

split rest is (a:Req(w),
               t:Un, rest':(q':Un, Public Response [end req(p, q', w, a, t)]));
match rest' is (q, n'_q:Public Response [end req(p, q, w, a, t)]);
check n_q is n'_q;
end req(p, q, w, a, t);
let r:Res(w)=call_w(p, a);
begin res(p, q, w, r, t);
cast n_p is (n'_p:Public Response [end res(p, q, w, r, t)]);
out k_2 (q, {res(w, r, t, p, n'_p)}_{SK_q})
where q = owner(w)

```

Finally, we need to change the top-level environment to account for the new keys, and to add a channel through which we will publish the public keys.

Top-Level Environments:

```

E_class ≜ (c:ℓ:Un) (c,ℓ)∈CIMeth
E_keys ≜ DSCA:AuthCertKeys, (Dsp:AuthKeys(p)) p∈Prin
E_ws ≜ (w:Un) w∈WebService
E_prin ≜ p_1:Prin, ..., p_n:Prin           where Prin = {p_1, ..., p_n}
E_net ≜ net:Un
E_0 ≜ E_ws, E_prin, E_net, E_class, E_keys

```

Publishing can be achieved by simply sending the public keys on a public channel, here *net*:

Public Keys Publishing:

```

I_net ≜ out net (VKCA, (VKp) p∈Prin)

```

We can now establish that the resulting system is robustly safe:

Theorem 2. If $\emptyset \vdash a : A$ and $p \in Prin$ and $k \notin \text{dom}(E_0)$ then the system

```
new (E_class, E_keys); (I_net | I_class | I_ws | new (k:Un); [a]_k^p)
```

is robustly safe.

Proof. See Appendix D.2. □

The protocol we give above to provide authentication has some undesirable properties. Specifically, it requires the server to remember the certificate *CertVKp* and nonce n_p at the time when a nonce is requested. Since anyone can request a nonce, and no authentication is performed at that stage of the protocol, this makes the server severely vulnerable to denial-of-service attacks. The following variation on the protocol achieves the same guarantees, but pushes the exchange of certificates and nonces to later messages, basically just when they are needed.

```

p → q on w : req(getnonce()), k_1
q → p on k_1 : res(getnonce(n_q))
p → q on w : p, CertVKp, n_p, {req(w, ℓ(u_1, ..., u_n), t, q, n_q)}_{SK_p}, k_2
q → p on k_2 : q, CertVKq, {res(w, ℓ(r), t, p, n_p)}_{SK_q}

```

6.2. Authenticated and Encrypted Web Methods

We now describe a protocol and implementation for authenticated and encrypted web methods. Hence, for now, we assume that all the exported methods of a web service are annotated with **AuthEnc**.

The public-key infrastructure we consider for this case is similar to the one for authenticated web methods, except that now we have encryption and decryption keys, as opposed to signing and verification keys. Each principal p has an encryption key *EKp* and a decryption key *DKp*. The decryption key is kept private, while

the encryption key is public. To bind the name of a principal with their encryption key, we again assume a *certification authority* CA (with a signing key $SKCA$ and verification key $VKCA$) that can sign certificates $CertEKp$ of the form $\{p, EKp\}_{SKCA}$.

Here is a protocol for p making a web service call $w:\ell(u_1, \dots, u_n)$ to service w owned by q , including the names of continuation channels used at the spi level. Again, we assume that in addition to the methods of $class(w)$, each web service also supports a method *getnonce*, which we implement specially.

$$\begin{aligned} p &\rightarrow q \text{ on } w : CertEKp, req(getnonce()), k_1 \\ q &\rightarrow p \text{ on } k_1 : CertEKq, \{msg_2(q, n_K)\}_{EKp}, res(getnonce(n_q)) \\ p &\rightarrow q \text{ on } w : \{msg_3(w, p, K, n_K)\}_{EKq}, n_p, \{req(\ell(u_1, \dots, u_n), t, n_q)\}_K, k_2 \\ q &\rightarrow p \text{ on } k_2 : \{res(\ell(r), t, n_p)\}_K \end{aligned}$$

This protocol is similar to that for authenticated web methods, except that public key encryption is used to exchange a session-specific shared key K used to encrypt the actual method call. Specifically, in the third message, p chooses a session-specific shared key K , and sends it to q encrypted with q 's public key EKq ; this session key K is used to encrypt the web method call. The result of the web method call is also encrypted with this shared key. To prevent replay attacks, the shared key is bound to a nonce n_K sent by q in the second message.

Type of Keys:

$SKey(p, q, w) \triangleq$ SharedKey(Union($req(a:Un, t:Un, n_q:Public\ Response\ [end\ req(p, q, w, a, t)])$), $res(r:Un, t:Un, n_p:Public\ Response\ [end\ res(p, q, w, r, t)])$))
$AuthEncMsg(p) \triangleq$ Union($msg_2(q:Un, n_K:Private\ Challenge\ [])$, $msg_3(w:Un, q:Un, K:Top,$ $n_K:Private\ Response\ [trust\ K:SKey(p, q, w)])$)
$AuthEncKeys(p) \triangleq KeyPair(AuthEncMsg(p))$
$AuthEncCert \triangleq (p:Un, Encrypt\ Key(AuthEncMsg(p)))$
$AuthEncCertKeys \triangleq KeyPair(AuthEncCert)$

We will represent the key pair of an encryption key and decryption key for principal p by a pair PKp , of type $AuthEncKeys(p)$. The signing key pair for the certification authority will be represented by a pair $DSCA$. We use the following abbreviations:

Key and Certificates Abbreviations:

$EKp \triangleq Encrypt\ (PKp)$	p 's encryption key
$DKp \triangleq Decrypt\ (PKp)$	p 's decryption key
$CertEKp \triangleq \{p, EKp\}_{SKCA}$	p 's certificate

Again, we can amend the translation of Section 4 to accommodate the new protocol. First, we give a new translation for a web method call $w:\ell(u_1, \dots, u_n)$:

New Translation of Web Method Call:

$\llbracket w:\ell(u_1, \dots, u_n) \rrbracket_k^p \triangleq$ new ($k_1:Un, k_2:Un, t:Un, n_p:Public\ Challenge\ []$); begin $req(p, q, w, \ell(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket), t)$; out $w\ (CertEKp, req(getnonce()), k_1)$; inp $k_1\ (c:Un, cipher:Un, res(getnonce(n_q:Un)))$;

```

decrypt c is {cert:(q':Un, Encrypt Key(AuthEncMsg(q')))} VKCA-1;
match cert is (q, ekq:Encrypt Key(AuthEncMsg(q)));
decrypt cipher is {msg2(q':Un, nK:Un)} DKp-1;
if q = q' then
cast nq is (n'q:Public Response [end req(p, q, w, ℓ([u1], ..., [un]), t)]);
new (K:SKey(p, q, w));
witness K:SKey(p, q, w);
cast nK is (n'K:Private Response [trust K:SKey(p, q, w)]);
out w ({msg3(w, p, t, K, n'K)} ekq, np, {req(w, ℓ([u1], ..., [un]), t, n'q)} K, k2);
inp k2 (bdy:Un);
decrypt bdy is {res(plain:(r:Res(w), t':Un,
Public Response [end res(p, q, w, r, t')]))} K;
match plain is (r:Res(w), rest:(t':Un, Public Response [end res(p, q, w, r, t')]));
match rest is (t, n'p:Public Response [end res(p, q, w, r, t)]);
check np is n'p;
end res(p, q, w, r, t);
case r is ℓ(x); out k x
where q = owner(w)

```

We also need to give a new implementation for web services, again to take into account the different messages being exchanged:

New Web Service Implementation:

```

Iws(w) ≜
repeat inp w (c:Un, bdy:Un, k1:Un);
case bdy is req(getnonce());
decrypt c is {p:Un, ekp:Encrypt Key(AuthEncMsg(p))} VKCA-1;
new (nq:Public Challenge []);
new (nK:Private Challenge []);
out k1 (CertEKq, {msg2(q, nK)} ekp, res(getnonce(nq)));
inp w (cipher1:Un, np:Un, cipher2:Un, k2:Un);
decrypt cipher1
is {msg3(plain1:(w:Un, p':Un, K:Top,
Private Response [trust K:SKey(p', q, w)]))} DKq-1;
match plain1 is (w, rest:(p':Un, K:Top,
Private Response [trust K:SKey(p', q, w)]));
match rest is (p, rest':(K:Top, Private Response [trust K:SKey(p, q, w)]));
split rest' is (K:Top, n'K:Private Response [trust K:SKey(p, q, w)]);
check nK is n'K;
trust K is (K':SKey(p, q, w));
decrypt cipher2 is {req(plain2:(a:Req(w), t:Un,
Public Response [end req(p, q, w, a, t)]))} K';
split plain2 is (a:Req(w), t:Un, n'q:Public Response [end req(p, q, w, a, t)]);
check nq is n'q;
end req(p, q, w, a, t);
let r:Res(w)=callw(p, a);
begin res(p, q, w, r, t);
cast np is (n'p:Public Response [end res(p, q, w, r, t)]);
out k2 {res(r, t, n'p)} K'
where q = owner(w)

```

Finally, we need to change the top-level environment to account for the new keys, and to add a channel through which we will publish the public keys.

Top-Level Environments:

$$\begin{aligned}
E_{class} &\triangleq (c:\ell:\text{Un}) \quad (c,\ell) \in \text{CIMeth} \\
E_{keys} &\triangleq \text{DSCA:AuthEncCertKeys}, (PKp:\text{AuthEncKeys}(p)) \quad p \in \text{Prin} \\
E_{ws} &\triangleq (w:\text{Un}) \quad w \in \text{WebService} \\
E_{prin} &\triangleq p_1:\text{Prin}, \dots, p_n:\text{Prin} \quad \text{where } \text{Prin} = \{p_1, \dots, p_n\} \\
E_{net} &\triangleq \text{net}:\text{Un} \\
E_0 &\triangleq E_{ws}, E_{prin}, E_{net}, E_{class}, E_{keys}
\end{aligned}$$

Publishing can be achieved by simply sending the public keys on a public channel, here net :

Public Keys Publishing:

$$I_{net} \triangleq \text{out } \text{net} \ (VKCA, (EKp) \quad p \in \text{Prin})$$

We can now establish that the resulting system is robustly safe:

Theorem 3. If $\emptyset \vdash a : A$ and $p \in \text{Prin}$ and $k \notin \text{dom}(E_0)$ then the system

$$\text{new } (E_{class}, E_{keys}); (I_{net} \mid I_{class} \mid I_{ws} \mid \text{new } (k:\text{Un}); \llbracket a \rrbracket_k^p)$$

is robustly safe.

Proof. See Appendix D.3. □

We can note some further possibilities, with respect to the protocols implemented in this section:

- The protocol implementing authenticated and encrypted invocation uses certificates to essentially negotiate a symmetric key with which to actually perform the encryption. It is straightforward to apply the same idea to the authenticated-only case, negotiating a symmetric key with which to hash the content of the method call (instead of relying on public-key signatures).
- In the above protocol, a new symmetric key is negotiated at every method invocation. A more efficient variation would be to re-use a negotiated symmetric key over multiple web method calls. Once a symmetric key has been negotiated, it can effectively act as a shared key between the two principals, which is the case we investigated in the body of this paper. We can therefore use the above protocol for the first web method call between a principal and a particular service, and the shared-key protocol for subsequent web method calls.

7. Related Work

There has been work for almost twenty years on secure RPC mechanisms, going back to Birrell [10]. More recently, secure RPC has been studied in the context of distributed object systems. As we mentioned, our work was inspired by the work of van Doorn *et al.* [35], itself inspired by [30, 36]. These techniques (or similar ones) have been applied to CORBA [31], DCOM [11], and Java [7, 19].

In contrast, little work seems to have been done on formalizing secure RPC. Of note is the work of Abadi, Fournet, and Gonthier [2, 3], who show how to compile the standard join-calculus into the sjoin-calculus, and show that the compilation is fully abstract. In a subsequent paper [4], they treat similarly and more simply a join-calculus with authentication primitives: each message contains its source address, there is a way to extract the principal owning a channel from the channel, and any piece of code runs as a particular principal. Their fully abstract translation gives very strong guarantees: it shows that for all intents and purposes, we can reason at the highest level (at the level of the authentication calculus). Although our guarantees are weaker, they are easier to establish.

Duggan [18] formalizes an application-level security abstraction by introducing types for signed and encrypted messages; he presents a fully abstract semantics for the abstraction by translation to a spi-calculus.

Much of the literature on security in distributed systems studies the question of *access control*. Intuitively, access control is the process of determining if the principal calling a particular method has permission to access the objects that the method refers to, according to a particular access control policy. There is a distinction to be made between authentication and access control. Authentication determines whether the

principal calling a method is indeed the principal claiming to be calling the method, while access control can use this authenticated identity to determine whether that principal is allowed access. This distinction is made clear in the work of Balfanz *et al.* [7], where they provide authenticated and encrypted communication over Java RMI (using SSL) and use that infrastructure as a basis for a logic-based access control mechanism. The access control decisions are based on the authenticated caller identity obtained from the layer in charge of authentication. This approach is also possible in our framework, which provides access to an authenticated identity as well. We plan to study access control abstractions in our framework. Various forms of access control mechanisms have been formalized via π -calculi, [26, 33, 27], and other process calculi [13, 16]. An access control language based on temporal logic has been defined by Sizer and Wang [34] specifically for web services. Damiani *et al.* [15] describe an implementation of an access control model for SOAP; unlike our work, and the WS-Security proposal, it relies on an underlying secure channel, such as an SSL connection.

Since this work was completed, a series of specifications for web services security has been published, as laid out in a whitepaper from IBM and Microsoft [28]. In particular, WS-Security [6] defines how to add signatures, to apply encryption, and to add principal identities, such as usernames or certificates, to a SOAP envelope. It would be straightforward, for example, to adapt our implementation to produce WS-Security compliant SOAP envelopes. A recent paper shows how to formalize the authentication goals of protocols based on WS-Security using the applied π -calculus [9].

Despite its enjoyable properties, the formal model we use to study the implementation of our security abstraction suffers from some limitations. For instance, it makes the usual Dolev-Yao assumptions that the adversary can compose messages, replay them, or decipher them if it knows the right key, but cannot otherwise “crack” encrypted messages. A more severe restriction is that we cannot yet model insider attacks: principals with shared keys are assumed well-behaved. Work is in progress to extend the Cryptyc type theory to account for malicious insiders. We have not verified the hash-based protocol of Section 2.

8. Conclusions

Authenticated method calls offer a convenient abstraction for developers of both client and server code. Various authorisation mechanisms may be layered on top of this abstraction. This paper proposes such an abstraction for web services, presents a theoretical model, and describes an implementation using SOAP-level security. By typing our formal semantics, we show no vulnerability exists to attacks representable within the spi-calculus, given certain assumptions. Vulnerabilities may exist outside our model—there are no methods, formal or otherwise, to guarantee security absolutely.

While our approach is restricted to proving properties of protocols that can be established using the Cryptyc type and effect system, it is worth pointing out that it is compatible with alternative methods for protocol verification. For instance, it is possible to analyze the protocols we use to implement secure web method calls for security flaws beyond those that can be uncovered using Cryptyc (for instance, flaws involving malicious insiders).

Our work shows that by exploiting recent advances in authenticity types, we can develop a theoretical model of a security abstraction, and then almost immediately obtain precise guarantees. (As with many formal analyses, these guarantees concern the design of our abstraction, and do not rule out code defects in its actual implementation.)

This study furthermore validates the adequacy of the spi-calculus, and Cryptyc in particular, to formally reason about security properties in a distributed communication setting.

Acknowledgments Cryptyc is an ongoing collaboration between Alan Jeffrey and the first author. Ernie Cohen, Cédric Fournet, and Alan Jeffrey made useful suggestions during the writing of this paper.

A. Sample SOAP Messages

We give some sample SOAP messages exchanged during web service method calls of the web service described in Section 5. One thing that is immediately clear is that we are not using standard XML formats for signing and encrypting messages, such as XML-Encryption and XML-Signature. There is no intrinsic difficulty in adapting our infrastructure to use standard formats. The point is that the validation of the security abstraction does not rely on the exact syntax of the SOAP envelopes.

A.1. An Authenticated Call

We describe an authenticated call to the `Balance` method. The messages exchanged to obtain the nonce are standard SOAP messages. The following message is the request from Alice to the web service to execute the `Balance` method on argument 12345. Notice the `DSHeader` element holding the identity of the principals involved, as well as the nonces and the cryptographic signature.

```
<?xml version="1.0" encoding="utf-8"?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <soap:Header>
    <DSHeader xmlns="http://tempuri.org/">
      <callerid>Alice</callerid>
      <calleeid>Bob</calleeid>
      <np>13</np>
      <nq>42</nq>
      <signature>
        3E:67:75:28:3B:AD:DF:32:E7:6C:D3:66:2A:CF:E7:8A:3F:0A:A6:0D
      </signature>
    </DSHeader>
  </soap:Header>
  <soap:Body>
    <Balance xmlns="http://tempuri.org/">
      <account>12345</account>
    </Balance>
  </soap:Body>
</soap:Envelope>
```

The response from the web service has a similar form:

```
<?xml version="1.0" encoding="utf-8"?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <soap:Header>
    <DSHeader xmlns="http://tempuri.org/">
      <callerid>Alice</callerid>
      <calleeid>Bob</calleeid>
      <np>13</np>
      <nq>42</nq>
      <signature>
        8D:31:52:6E:08:F0:89:7B:1E:12:3F:5E:63:EE:B0:D2:63:89:CA:73
      </signature>
    </DSHeader>
  </soap:Header>
  <soap:Body>
    <BalanceResponse xmlns="http://tempuri.org/">
      <BalanceResult>100</BalanceResult>
    </BalanceResponse>
  </soap:Body>
</soap:Envelope>
```

A.2. Authenticated and Encrypted Call

We describe an authenticated and encrypted call, this time to the `Statement` method. Again, the messages exchanged to obtain the nonce are standard SOAP messages. The following message is the request from

Alice to the web service to execute the `Statement` method on argument 12345. As in the authenticated call above, the `DSHeader` element holds identity information. The body of the message itself is encrypted. Note that the nonce `nq` must be encrypted according to the protocol, so its encrypted value is included in the encrypted data, and its element is reset to a dummy value (here, -1). Similarly, the signature is unused and set to a dummy value.

```
<?xml version="1.0" encoding="utf-8"?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <soap:Header>
    <DSHeader xmlns="http://tempuri.org/">
      <callerid>Alice</callerid>
      <calleeid>Bob</calleeid>
      <np>13</np>
      <nq>-1</nq>
      <signature>4E:00:6F:00</signature>
    </DSHeader>
  </soap:Header>
  <soap:Body>
    9D:8F:95:2B:BC:60:B1:73:A7:C4:82:F5:39:20:97:F7:69:71:66:
    D3:A3:A0:90:B9:9B:FE:71:0A:65:C1:EF:EE:99:CB:4D:8A:40:37:
    CA:1E:D0:03:50:34:76:8C:E3:F3:30:DD:C9:34:19:D4:04:CB:39:
    7D:1A:84:2F:CA:30:DA:68:7E:E1:CB:07:9C:EB:79:F9:E9:4B:47:
    5B:94:56:D7:22:0E:02:CD:AA:F5:D3:40:C1:EC:13:FB:B9:E6:4F:
    13:CD:70:FD:BA:18:80:FC:50:F3:75:F2:2F:95:50:5D:41:7E:C8:
    8B:BB:AB:76:C9:59:BA:E2:3B:E5:4D:79:71:E4:AD:18:5A:4B:EA:
    29:17:30:90:66:08:27:ED:B4:BD:2E:89:06:6D:0B:56:40:43:35:
    A1:77:AE:12:7E:4B:19:26:B5:24:1A:D9:67:3D:A0:91
  </soap:Body>
</soap:Envelope>
```

The response is similarly encoded. Notice that this time the nonce `np` must be encrypted, so its value is again included in the encrypted data, and its element is reset to a dummy value.

```
<?xml version="1.0" encoding="utf-8"?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <soap:Header>
    <DSHeader xmlns="http://tempuri.org/">
      <callerid>Alice</callerid>
      <calleeid>Bob</calleeid>
      <np>-1</np>
      <nq>-1</nq>
      <signature>4E:00:6F:00</signature>
    </DSHeader>
  </soap:Header>
  <soap:Body>
    98:FD:6A:5B:38:0A:82:95:3F:01:EC:D3:55:F9:AA:35:4D:18:DB:
    1B:7D:9D:FE:3F:78:52:29:99:C9:41:84:EE:B1:42:12:B2:02:AC:
    63:F5:0C:92:9B:DB:75:FB:6C:8B:65:EB:3C:42:6B:79:70:AF:61:
    2A:C2:7B:ED:96:E1:D6:7A:F6:D2:0C:DF:BC:2A:4C:93:B3:D0:7B:
    7D:2D:83:18:60:D2:D8:05:EB:73:74:2D:75:A2:B2:57:C9:04:B4:
    C1:E6:66:54:BA:42:86:AF:22:72:3D:B7:90:CF:03:22:E5:C4:47:
    03:F0:77:A0:30:01:C9:FE:78:A1:AB:FA:B1:CB:EE:E2:0B:F2:79:
    17:1B:8E:82:E2:13:F4:66:52:76:6D:BA:1B:E9:8E:75:15:90:37:
```

```

0A:64:ED:F3:9C:18:94:EC:4F:CF:61:92:38:EF:A9:46:E8:4E:E9:
4A:E6:8A:C9:5E:ED:A7:34:72:3E:72:A2:BE:0D:DC:07:22:45:B0:
E6:79:33:8F:CD:90:B8:97:DB:BA:3B:B2:8B:38:38:B6:5B:F1:11:
FB:DD:88:CE:9A:3E:B4:E6:31:13:CB:1C:F3:B5:17:D8:9B:CF:2E:
65:23:4D:BA:ED:72:6D:F4:53:97:B8:7A:D2:9C:2C:10:58:A3:0E:
FE:48:A2:2A:2A:57:AE:6D:69:4D:97:90:EF:9F:C6:7E:9B
</soap:Body>
</soap:Envelope>

```

B. Semantics of the Object Calculus

In this appendix, we give a formal description of the operational semantics and typing rules of the object calculus. We first describe some encodings showing the expressiveness of the calculus.

B.1. Encoding Arithmetic

The calculus is simple enough that questions about whether or not it is sufficiently expressive to be of interest arise. This is especially likely since there are no recursive functions in the calculus, and it is not clear that it is even Turing complete. That the calculus indeed is Turing complete is a consequence of the fact that we can write recursive classes and methods, and that we have a null object. The following example shows an encoding of natural numbers as a class *Num*, with the typical recursive definition of addition:

```

class Num
  Num pred
  Num succ()
  new Num(this)
  Num add(Num x)
  if x.pred = null then
    this
  else this.add(x.pred).succ()

```

We define *zero* as *new Num(null)*, *one* as *zero.succ()*, and so on.

B.2. Formalization of proxy objects

We mentioned in the text that we can easily express proxy objects within the calculus. For completeness, here is a detailed formalization of such proxy objects. First, we assume a map $proxy \in WebService \rightarrow Class$, assigning to every web service $w \in WebService$ a proxy class $proxy(w)$. We further assume that for each $w \in WebService$,

- $dom(methods(class(w))) \cup \{Id\} = dom(methods(proxy(w)))$,
- $fields(proxy(w)) = \emptyset$,
- $methods(proxy(w))(Id) = (Id(), owner(w))$, and
- for all $\ell \in dom(methods(class(w)))$,

$$methods(proxy(w))(\ell) = (B(A_1 x_1, \dots, A_n x_n), w:\ell(x_1, \dots, x_n)),$$

$$\text{where } methods(class(w))(\ell) = (B(A_1 x_1, \dots, A_n x_n), b).$$

B.3. Operational Semantics

The operational semantics is defined by a transition relation, written $a \rightarrow^p a'$, where a and a' are method bodies, and p is the principal evaluating the body a .

To specify the semantics, we need to keep track of which principal is currently running a method body.

We add a new method body form to our object calculus, $p[a]$, meaning p running body a . This form does not appear in code written by the user, but only arises through the transitions of the semantics.

Extended Method Bodies:

$a, b \in \text{Body} ::=$	method body
\dots	as in Section 3
$p[a]$	body a running as p

Transitions:

(Red Let 1)	(Red Let 2)
$\frac{a \rightarrow^p a'}{\text{let } x=a \text{ in } b \rightarrow^p \text{let } x=a' \text{ in } b}$	$\frac{}{\text{let } x=v \text{ in } b \rightarrow^p b\{x \leftarrow v\}}$
(Red If)	
$\frac{}{\text{if } u = v \text{ then } a_{\text{true}} \text{ else } a_{\text{false}} \rightarrow^p a_{u=v}}$	
(Red Field)	
$\frac{\text{fields}(c) = f_i \mapsto A_i \quad i \in 1..n \quad j \in 1..n}{(\text{new } c(v_1, \dots, v_n)).f_j \rightarrow^p v_j}$	
(Red Invoke)(where $v = \text{new } c(v_1, \dots, v_n)$)	
$\frac{\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i) \quad i \in 1..n \quad j \in 1..n \quad \text{sig}_j = B(A_1 x_1, \dots, A_m x_m)}{v.\ell_j(u_1, \dots, u_m) \rightarrow^p b_j\{\text{this} \leftarrow v, x_k \leftarrow u_k \quad k \in 1..m\}}$	
(Red Remote)	
$\frac{\text{owner}(w) = q \quad \text{class}(w) = c}{w:\ell(u_1, \dots, u_n) \rightarrow^p q[\text{new } c(p).\ell(u_1, \dots, u_n)]}$	
(Red Prin 1)	(Red Prin 2)
$\frac{a \rightarrow^q a'}{q[a] \rightarrow^p q[a']}$	$\frac{}{q[v] \rightarrow^p v}$

B.4. Type System

The judgments of our type system all depend on an *environment* E , that defines the types of all variables in scope. An environment takes the form $x_1:A_1, \dots, x_n:A_n$ and defines the type A_i for each variable x_i . The domain $\text{dom}(E)$ of an environment E is the set of variables whose types it defines.

Environments:

$D, E ::=$	environment
\emptyset	empty
$E, x:A$	entry
$\text{dom}(x_1:A_1, \dots, x_n:A_n) \triangleq \{x_1, \dots, x_n\}$	domain of an environment

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

Judgments $E \vdash \mathcal{J}$:

$E \vdash \diamond$	good environment
$E \vdash a : A$	good expression a of type A

We write $E \vdash \mathcal{J}$ when we want to talk about both kinds of judgments, where \mathcal{J} stands for either \diamond or $a : A$.

The following rules define an environment $x_1:A_1, \dots, x_n:A_n$ to be well-formed if each of the names x_1, \dots, x_n are distinct.

Rules for Environments:

$$\frac{}{\emptyset \vdash \diamond} \quad \frac{(\text{Env } \emptyset) \quad (\text{Env } x) \text{ (where } x \notin \text{dom}(E))}{E \vdash \diamond}$$

We present the rules for deriving the judgment $E \vdash a : A$ that assigns a type A to a value or method body a . These rules are split into two tables, one for values, and one for method bodies.

Rules for Typing Values:

$$\frac{(\text{Val } x) \quad E = E_1, x:A, E_2 \quad E \vdash \diamond}{E \vdash x : A} \quad \frac{(\text{Val } \text{null}) \quad E \vdash \diamond}{E \vdash \text{null} : c}$$

$$\frac{(\text{Val Object}) \quad \text{fields}(c) = f_i \mapsto A_i^{i \in 1..n} \quad E \vdash v_i : A_i \quad \forall i \in 1..n}{E \vdash \text{new } c(v_1, \dots, v_n) : c} \quad \frac{(\text{Val Princ}) \quad E \vdash \diamond}{E \vdash p : \text{Id}}$$

Rules for Typing Method Bodies:

$$\frac{(\text{Body Let}) \quad E \vdash a : A \quad E, x:A \vdash b : B}{E \vdash \text{let } x=a \text{ in } b : B}$$

$$\frac{(\text{Body If}) \quad E \vdash u : A \quad E \vdash v : A \quad E \vdash a : B \quad E \vdash b : B}{E \vdash \text{if } u = v \text{ then } a \text{ else } b : B}$$

$$\frac{(\text{Body Field}) \quad E \vdash v : c \quad \text{fields}(c) = f_i \mapsto A_i^{i \in 1..n} \quad j \in 1..n}{E \vdash v.f_j : A_j}$$

$$\frac{(\text{Body Invoke}) \quad E \vdash v : c \quad \text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)^{i \in 1..n} \quad j \in 1..n \quad \text{sig}_j = B(A_1 x_1, \dots, A_m x_m) \quad E \vdash u_k : A_k \quad \forall k \in 1..m}{E \vdash v.\ell_j(u_1, \dots, u_m) : B}$$

$$\frac{(\text{Body Remote}) \quad \text{class}(w) = c \quad \text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)^{i \in 1..n} \quad j \in 1..n \quad \text{sig}_j = B(A_1 x_1, \dots, A_m x_m) \quad E \vdash u_i : A_i \quad \forall i \in 1..m}{E \vdash w.\ell_j(u_1, \dots, u_m) : B} \quad \frac{(\text{Body Princ}) \quad E \vdash a : A}{E \vdash p[a] : A}$$

We make the following assumption on the execution environment.

Assumptions on the Execution Environment:

- (1) For each $w \in \text{WebService}$, $\text{fields}(\text{class}(w)) = \text{CallerId} : \text{Id}$.
- (2) No tagged expression $p[a]$ occurs within the body of any method; such expressions occur only at runtime, to track the call stack of principals.
- (3) for each $c \in \text{Class}$ and each $\ell \in \text{dom}(\text{methods}(c))$, if $\text{methods}(c)(\ell) = (B(A_1 x_1, \dots, A_n x_n), b)$, then $\text{this}:c, x_1:A_1, \dots, x_n:A_n \vdash b : B$.

It is straightforward to show that our type system is sound, that is, that the type system ensures that methods that typecheck do not get stuck when evaluating. Some care is needed to make this precise, since evaluation can block if one attempts to access a field of a null object, or to invoke a method on a null object. (We could introduce an *error token* in the semantics and propagate that error token when such a case is encountered, but this would needlessly complicate the semantics, at least for our purposes.) Soundness can be derived as usual via Preservation and Progress theorems. To establish these, we first need the following lemmas:

Lemma 1. The following properties of judgments hold:

- (**Exchange**) if $E, x:A, y:B, E' \vdash \mathcal{J}$, then $E, y:B, x:A, E' \vdash \mathcal{J}$;
- (**Weakening**) if $E \vdash \mathcal{J}$ and $x \notin \text{dom}(E)$, then $E, x:A \vdash \mathcal{J}$;
- (**Strengthening**) if $E, x:B \vdash a : A$ and $x \notin \text{fv}(a)$, then $E \vdash a : A$.

Proof. Straightforward. □

We often use the above properties silently in the course of proofs.

Lemma 2 (Substitution). If $E, x:B \vdash a : A$ and $E \vdash v : B$, then $E \vdash a\{x \leftarrow v\} : A$

Proof This is a straightforward proof by induction on the height of the typing derivation for $E \vdash a : A$. We proceed by case analysis on the form of a .

- Case $a = x$: Since $E, x:B \vdash x : A$, we must have $A = B$. Since $a\{x \leftarrow v\} = v$ and $E \vdash v : B$, we have $E \vdash v : A$, as required.
- Case $a = y$, where $y \neq x$: Since x is not free in y , $E, x:B \vdash y : A$ implies $E \vdash y : A$, by the Strengthening Lemma, as required.
- Case $a = \text{null}$: Since x is not free in null , $E, x:B \vdash \text{null} : A$ implies $E \vdash \text{null} : A$, as required.
- Case $a = p$: Since x is not free in p , $E, x:B \vdash p : A$ implies $E \vdash p : A$, as required.
- Case $a = \text{newc}(v_1, \dots, v_n)$: We have the equation $\text{newc}(v_1, \dots, v_n)\{x \leftarrow v\} = \text{newc}(v_1\{x \leftarrow v\}, \dots, v_n\{x \leftarrow v\})$. We have $E, x:B \vdash \text{newc}(v_1, \dots, v_n) : A$, hence $E, x:B \vdash v_i : A_i$ if $\text{fields}(c) = f_i \mapsto A_i$ $i \in 1..n$. By the induction hypothesis, we know $E \vdash v_i\{x \leftarrow v\} : A_i$ for all $i \in 1..n$. Hence, we can derive $E \vdash \text{newc}(v_1\{x \leftarrow v\}, \dots, v_n\{x \leftarrow v\}) : A$, as required.
- Case $a = \text{let } y=a_0 \text{ in } b$: Without loss of generality, we can take $y \neq x$, since y is bound in b . Note that $(\text{let } y=a_0 \text{ in } b)\{x \leftarrow v\} = \text{let } y=a_0\{x \leftarrow v\} \text{ in } b\{x \leftarrow v\}$. We have $E, x:B \vdash \text{let } y=a_0 \text{ in } b : A$, hence $E, x:B \vdash a_0 : A_0$ for some A_0 , and $E, y:A_0, x:B \vdash b : A$. By the induction hypothesis, $E \vdash a_0\{x \leftarrow v\} : A_0$ and $E, y:A_0 \vdash b\{x \leftarrow v\} : A$, and hence $E \vdash \text{let } x=a_0\{x \leftarrow v\} \text{ in } b\{x \leftarrow v\} : A$, as required.
- Case $a = \text{if } u_0 = u_1 \text{ then } a_0 \text{ else } a_1$: We have $(\text{if } u_0 = u_1 \text{ then } a_0 \text{ else } a_1)\{x \leftarrow v\} = \text{if } u_0\{x \leftarrow v\} = u_1\{x \leftarrow v\} \text{ then } a_0\{x \leftarrow v\} \text{ else } a_1\{x \leftarrow v\}$. We have $E, x:B \vdash \text{if } u_0 = u_1 \text{ then } a_0 \text{ else } a_1$, hence $E, x:B \vdash u_0 : A'$, $E, x:B \vdash u_1 : A'$, $E, x:B \vdash a_0 : A$ and $E, x:B \vdash a_1 : A$. Applying the induction hypothesis to these judgments, we can derive

$$E \vdash \text{if } u_0\{x \leftarrow v\} = u_1\{x \leftarrow v\} \text{ then } a_0\{x \leftarrow v\} \text{ else } a_1\{x \leftarrow v\} : A$$

as required.

The remaining cases are similar, upon noting that:

- $(u.f_j)\{x \leftarrow v\} = u\{x \leftarrow v\}.f_j$,
- $(u.\ell(u_1, \dots, u_m))\{x \leftarrow v\} = u\{x \leftarrow v\}(\ell(u_1\{x \leftarrow v\}, \dots, u_m\{x \leftarrow v\}))$,
- $(w:\ell(u_1, \dots, u_n))\{x \leftarrow v\} = w:(u_1\{x \leftarrow v\}, \dots, u_n\{x \leftarrow v\})$, and
- $(p[a])\{x \leftarrow v\} = p[a\{x \leftarrow v\}]$. □

Theorem 4 (Preservation). If $E \vdash a : A$ and $a \rightarrow^p a'$ then $E \vdash a' : A$.

Proof We proceed by induction on the height of the typing derivation for $E \vdash a : A$. Since $a \rightarrow^p a'$, a cannot be a value v .

- Case $a = \text{let } x=v \text{ in } b$: Since $E \vdash a : A$, we have $E \vdash v : B$ and $E, x:B \vdash b : A$. We must have $a' = b\{x \leftarrow v\}$. Applying the Substitution Lemma, we have $E \vdash b\{x \leftarrow v\} : A$, as required.

- Case $a = \text{let } x=a_0 \text{ in } b$, where a_0 is not a value: We have $E \vdash a_0 : B$, and $E, x:B \vdash b : A$. Since $a \rightarrow^p a'$, we must have $a_0 \rightarrow^p a'_0$. By induction hypothesis, $E \vdash a'_0 : B$, and hence $E \vdash \text{let } x=a'_0 \text{ in } b : A$, as required.
- Case $a = \text{if } u = v \text{ then } a_0 \text{ else } a_1$: Note that either $a \rightarrow^p a_0$ or $a \rightarrow^p a_1$. In both cases, since $E \vdash \text{if } u = v \text{ then } a_0 \text{ else } a_1 : A$, we have $E \vdash a_0 : A$ and $E \vdash a_1 : A$, as required.
- Case $a = (\text{new } c(v_1, \dots, v_n)).f_j$: We have $\text{fields}(c) = f_i \mapsto A_i^{i \in 1..n}$. The type derivation for a is as follows:

$$\frac{\frac{E \vdash v_i : A_i^{i \in 1..n}}{E \vdash \text{new } c(v_1, \dots, v_n) : c}}{E \vdash (\text{new } c(v_1, \dots, v_n)).f_j : A_j}$$

Since $a' = v_j$, we have $E \vdash v_j : A_j$, as required.

- Case $a = (\text{new } c(v_1, \dots, v_n)).\ell_j(u_1, \dots, u_m)$: Let $v = \text{new } c(v_1, \dots, v_n)$. We have $\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)^{i \in 1..n}$, where $\text{sig}_j = B(A_1 x_1, \dots, A_m x_m)$. By the typing derivation for $E \vdash a : B$, we have $E \vdash u_k : A_k$ for all $k \in 1..m$, and $E \vdash v : c$. By assumption on the execution environment, we know $\text{this}:c, x_1:A_1, \dots, x_m:A_m \vdash b : B$. Applying the Substitution and the Weakening Lemmas, we get $E \vdash b\{\text{this} \leftarrow v, x_k \leftarrow u_k^{k \in 1..m}\} : B$, as required.
- Case $a = w:\ell_j(u_1, \dots, u_m)$: We have $\text{class}(w) = c$, $\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)^{i \in 1..n}$ where $\text{sig}_j = B(A_1 x_1, \dots, A_m x_m)$. By the typing derivation for $E \vdash a : B$, we have $E \vdash u_i : A_i$ for all $i \in 1..m$. We can therefore derive the required type for $a' = q[\text{new } c(p).\ell(u_1, \dots, u_m)]$:

$$\frac{\frac{E \vdash \text{new } c(p) : c \quad E \vdash u_i : A_i \quad \forall i \in 1..m}{E \vdash \text{new } c(p).\ell(u_1, \dots, u_m) : B}}{E \vdash q[\text{new } c(p).\ell(u_1, \dots, u_m)] : B}$$

- Case $a = q[v]$: Since $E \vdash q[v] : A$, we have $E \vdash v : A$, and $q[v] \rightarrow^p v$, as required.
- Case $a = q[a_0]$, where a_0 is not a value: Since $E \vdash q[a_0] : A$, we have $E \vdash a_0 : A$, and since $a \rightarrow^p a'$, we must have $a_0 \rightarrow^q a'_0$. By induction hypothesis, $E \vdash a'_0 : A$, and hence $E \vdash q[a'_0] : A$, as required. \square

To state the Progress Theorem, we need to recognize programs that are blocked because of a *null* in object position. We say a method body a is *null-blocked* if, essentially, it is stuck trying to access a field of a null object, or invoke a method on a null object. Formally, a is null-blocked if it is of the form $\text{null}.f_j$, $\text{null}.\ell(u_1, \dots, u_n)$, $\text{let } x=a \text{ in } b$ (where a is null-blocked), or $q[a]$ (where a is null-blocked).

Theorem 5 (Progress). If $\emptyset \vdash a : A$ and a is not a value and is not null-blocked, and $p \in \text{Prin}$, then $a \rightarrow^p a'$ for some a' .

Proof Again, we proceed by induction on the height of the typing derivation for $\emptyset \vdash a : A$. We assume a is not a value, and a is not null-blocked.

- Case $a = \text{let } x=a_0 \text{ in } b$: We consider two subcases, depending on whether a_0 is a value or not.
 - Case a_0 is a value v : We have $a \rightarrow^p b\{x \leftarrow v\}$.
 - Case a_0 is not a value: Since $\emptyset \vdash a : A$, we have $\emptyset \vdash a_0 : B$ for some B , a_0 not a value. Since a is not null-blocked, a_0 is not null-blocked. Hence, by induction hypothesis, we have $a_0 \rightarrow^p a'_0$. Hence, we have $a \rightarrow^p \text{let } x=a'_0 \text{ in } b$.
- Case $a = \text{if } u = v \text{ then } a_0 \text{ else } a_1$: We have $a \rightarrow^p a_0$ or $a \rightarrow^p a_1$ depending on the result of $u = v$.
- Case $a = v.f_j$: Since $\emptyset \vdash a : A$ and a is not null-blocked, we must have $v = \text{new } c(u_1, \dots, u_n)$, and $\text{fields}(c) = f_i \mapsto A_i^{i \in 1..n}$. Therefore, we have $v.f_j \rightarrow^p u_j$.
- Case $a = v.\ell_j(u_1, \dots, u_m)$: Since $\emptyset \vdash a : A$ and a is not null-blocked, we must have $v = \text{new } c(u_1, \dots, u_n)$, $\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)$, and $\text{sig}_j = B(A_1 x_1, \dots, A_m x_m)$. Therefore, we have $v.\ell_j(u_1, \dots, u_m) \rightarrow^p b_j\{\text{this} \leftarrow v, x_k \leftarrow u_k^{k \in 1..m}\}$.
- Case $a = w:\ell(u_1, \dots, u_m)$: The following transition rule $w:\ell(u_1, \dots, u_m) \rightarrow^p q[\text{new } c(p).\ell(u_1, \dots, u_m)]$ applies, with $\text{owner}(w) = q$ and $\text{class}(w) = c$.
- Case $a = q[a_0]$: We consider two subcases, depending on whether a_0 is a value or not.

- Case a_0 is a value v : We have $q[v] \rightarrow^p v$.
- Case a_0 is not a value: Since $\emptyset \vdash q[a_0] : A$, we have $\emptyset \vdash a_0 : A$, a_0 not a value. Since a is not null-blocked, a_0 is not null-blocked. Hence, by induction hypothesis, we have $a_0 \rightarrow^q a'_0$, and $q[a_0] \rightarrow^p q[a'_0]$. \square

We can now state soundness formally. We say a method body a is stuck if a is not a value, a is not null-blocked, and there is no a' and p such that $a \rightarrow^p a'$. We write $a \rightarrow^* a'$ to mean that there exists a sequence a_1, \dots, a_n and principals p_1, \dots, p_{n+1} such that $a \rightarrow^{p_1} a_1 \rightarrow^{p_2} \dots \rightarrow^{p_n} a_n \rightarrow^{p_{n+1}} a'$. (Hence, \rightarrow^* is a kind of transitive closure of \rightarrow^p .)

Theorem 6 (Soundness). If $\emptyset \vdash a : A$, and $a \rightarrow^* a'$, then a' is not stuck.

Proof. A straightforward induction on the number of transitions in $a \rightarrow^* a'$. \square

C. The Spi-Calculus in More Detail

We give an overview of the language and type system on which our analysis of web services depends. We give the syntax in detail, but for the sake of brevity give only an informal account of the operational semantics and type system. Full details are in a technical report [21], from which some of the following explanations are drawn. Some constructs primitive here are actually derived forms in the original calculus.

Names, Messages:

$k ::= \text{Encrypt} \mid \text{Decrypt}$	key attribute
m, n, x, y, z	name: nonce, key, key-pair
$L, M, N ::=$	message
x	name
(M_1, \dots, M_n)	record, $n \geq 0$
$t_i(M)$	tagged union
$\{M\}_N$	symmetric encryption
$\{M\}_N$	asymmetric encryption
$k(M)$	key-pair component

The message x is a name, representing a channel, nonce, symmetric key, or asymmetric key-pair. We do not differentiate in the syntax or operational semantics between key-pairs used for public key cryptography and those used for digital signatures.

The message (M_1, \dots, M_n) is a record with n fields, M_1, \dots, M_n .

The message $t_i(M)$ is message M tagged with tag t_i . The message $\{M\}_N$ is the ciphertext obtained by encrypting the plaintext M with the symmetric key N .

The message $\{M\}_N$ is the ciphertext obtained by encrypting the plaintext M with the asymmetric encryption key N .

The message $\text{Decrypt}(M)$ is the decryption key (or signing key) component of the key-pair M , and $\text{Encrypt}(M)$ is the encryption key (or verification key) component of the key-pair M .

Types and Effects:

$\ell ::= \text{Public} \mid \text{Private}$	nonce attribute
$S, T, U ::=$	type
Un	data known to the opponent
$(x_1:T_1, \dots, x_n:T_n)$	dependent record, $n \geq 0$
$\text{Union}(t_1(T_1), \dots, t_n(T_n))$	tagged union
Top	top
$\text{SharedKey}(T)$	shared-key type
$\text{KeyPair}(T)$	asymmetric key-pair
$k \text{ Key}(T)$	encryption or decryption part
$\ell \text{ Challenge } es$	challenge type
$\ell \text{ Response } fs$	response type
$e, f ::=$	atomic effect

end L	end-event labelled L
check ℓN	name-check for a nonce N
trust $M:T$	trust that $M:T$
$\mathcal{e}s, \mathcal{f}s ::=$ [e_1, \dots, e_n]	effect multiset of atomic effects

The type Un describes messages that may flow to or from the opponent, which we model as an arbitrary process of the calculus. We say that a type is *public* if messages of the type may flow to the opponent. Dually, we say a type is *tainted* if messages from the opponent may flow into the type. The type Un is both public and tainted.

The type $(x_1:T_1, \dots, x_n:T_n)$ describes a record (M_1, \dots, M_n) where each $M_i : T_i$. The scope of each variable x_i consists of the types T_{i+1}, \dots, T_n . Type $(x_1:T_1, \dots, x_n:T_n)$ is public just if all of the types T_i are public, and tainted just if all of the types T_i are tainted.

The type $\text{Union}(t_1(T_1), \dots, t_n(T_n))$ describes a tagged message $t_i(M)$ where $i \in 1..n$ and $M : T_i$. Type $\text{Union}(t_1(T_1), \dots, t_n(T_n))$ is public just if all of the types T_i are public, and tainted just if all of the types T_i are tainted.

The type Top describes all well-typed messages; it is tainted but not public.

The type $\text{SharedKey}(T)$ describes symmetric keys for encrypting messages of type T ; it is public or tainted just if T is both public and tainted.

The type $\text{KeyPair}(T)$ describes asymmetric key-pairs for encrypting or signing messages of type T ; it is public or tainted just if T is both public and tainted. The key-pair can be used for public-key cryptography just if T is tainted, and for digital signatures just if T is public.

The type $\text{Encrypt Key}(T)$ describes an encryption or signing key for messages of type T ; it is public just if T is tainted, and it is tainted just if T is public.

The type $\text{Decrypt Key}(T)$ describes a decryption or verification key for messages of type T ; it is public just if T is public, and it is tainted just if T is tainted.

The types ℓ **Challenge** $\mathcal{e}s$ and ℓ **Response** $\mathcal{f}s$ describe nonce challenges and responses, respectively. The effects $\mathcal{e}s$ and $\mathcal{f}s$ embedded in these types represent certain events. An outgoing challenge of some type ℓ **Challenge** $\mathcal{e}s$ can be cast into a response of type ℓ **Response** $\mathcal{f}s$ and then returned, provided the events in the effect $\mathcal{e}s + \mathcal{f}s$ have been justified, as explained below. Therefore, if we have created a fresh challenge at type ℓ **Challenge** $\mathcal{e}s$, and check that it equals an incoming response of type ℓ **Response** $\mathcal{f}s$, we can conclude that the events in $\mathcal{e}s + \mathcal{f}s$ may safely be performed. The attribute ℓ is either **Public** or **Private**; the former means the nonce may eventually be public, while the latter means the nonce is never made public. Type **Public Challenge** $\mathcal{e}s$ is public, or tainted, just if $\mathcal{e}s = []$. Type **Public Response** $\mathcal{f}s$ is always public, but tainted just if $\mathcal{e}s = []$. Neither **Private Challenge** $\mathcal{e}s$ nor **Private Response** $\mathcal{f}s$ is public; both are tainted.

An effect $\mathcal{e}s$ is a multiset, that is, an unordered list of atomic effects, e or f . Effects embedded in challenge or response types signify that certain actions are justified, that is, may safely be performed. An atomic effect end L justifies a single subsequent end-event labelled L , and is justified by a distinct, preceding begin-event labelled L . An atomic effect check ℓN justifies a single subsequent check that an ℓ response equals an ℓ challenge named N , where ℓ is **Public** or **Private**, and is justified by freshly creating the challenge N . An atomic effect trust $M:T$ justifies casting message M to type T , and is justified by showing that M indeed has type T .

Processes:

$O, P, Q, R ::=$	process
out $M N$	output
inp $M (x:T); P$	input (x bound in P)
repeat inp $M (x:T); P$	replicated input (x bound in P)
split M is $(x_1:T_1, \dots, x_n:T_n); P$	record splitting
match M is $(N, y:T); P$	pair matching (y bound in P)
case M is $t_i(x_i:T_i); P_i$ $i \in 1..n$	tagged union case (t_i distinct)
if $M = N$ then P else Q	conditional (new)
new $(x:T); P$	name generation (x bound in P)
$P \mid Q$	composition
stop	inactivity

decrypt M is $\{x:T\}_N;P$	symmetric decrypt (x bound in P)
decrypt M is $\{\{x:T\}\}_{N^{-1}};P$	asymmetric decrypt (x bound in P)
check M is $N;P$	nonce-checking
begin $L;P$	begin-assertion
end $L;P$	end-assertion
cast M is $(x:T);P$	cast to nonce type
witness $M:T;P$	witness testimony
trust M is $(x:T);P$	trusted cast

The processes $\text{out } M \ N$ and $\text{inp } M \ (x:T); P$ are output and input, respectively, along an asynchronous, unordered channel M . If an output $\text{out } x \ N$ runs in parallel with an input $\text{inp } x \ (y); P$, the two can interact to leave the residual process $P\{y \leftarrow N\}$, the outcome of substituting N for each free occurrence of y in P . We write $\text{out } x \ (M); P$ as a simple shorthand for $\text{out } x \ M \mid P$.

The process $\text{repeat inp } M \ (x:T); P$ is replicated input, which behaves like input, except that each time an input of N is performed, the residual process $P\{y \leftarrow N\}$ is spawned off to run concurrently with the original process $\text{repeat inp } M \ (x:T); P$.

The process $\text{split } M \text{ is } (x_1:T_1, \dots, x_n:T_n); P$ splits the record M into its n components. If M is (M_1, \dots, M_n) , the process behaves as $P\{x_1 \leftarrow M_1\} \dots \{x_n \leftarrow M_n\}$. Otherwise, it deadlocks, that is, does nothing.

The process $\text{match } M \text{ is } (N, y:U); P$ splits the pair (binary record) M into its two components, and checks that the first one is N . If M is (N, L) , the process behaves as $P\{y \leftarrow L\}$. Otherwise, it deadlocks.

The process $\text{case } M \text{ is } t_i(x_i:T_i); P_i \text{ }^{i \in 1..n}$ checks the tagged union M . If M is $t_j(L)$ for some $j \in 1..n$, the process behaves as $P\{x_i \leftarrow L\}$. Otherwise, it deadlocks.

The process $\text{if } M = N \ \text{then } P \ \text{else } Q$ behaves as P if M and N are the same message, and otherwise as Q . (This process is not present in the original calculus [21] but is a trivial and useful addition.)

The process $\text{new } (x:T); P$ generates a new name x , whose scope is P , and then runs P . This abstractly represents nonce or key generation.

The process $P \mid Q$ runs processes P and Q in parallel.

The process stop is deadlocked.

The process $\text{decrypt } M \text{ is } \{x:T\}_N; P$ decrypts M using symmetric 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.

The process $\text{decrypt } M \text{ is } \{\{x:T\}\}_{N^{-1}}; P$ decrypts M using asymmetric key N . If M is $\{L\}_{\text{Encrypt}(K)}$ and N is $\text{Decrypt}(K)$, then the process behaves as $P\{x \leftarrow L\}$. Otherwise, it deadlocks.

The process $\text{check } M \text{ is } N; P$ checks the messages M and N are the same name before executing P . If the equality test fails, the process deadlocks.

The process $\text{begin } L; P$ autonomously performs a begin-event labelled L , and then behaves as P .

The process $\text{end } L; P$ autonomously performs an end-event labelled L , and then behaves as P .

The process $\text{cast } M \text{ is } (x:T); P$ binds the message M to the variable x of type T , and then runs P . In well-typed programs, M is a challenge of type ℓ **Challenge** *es*, and T is a response type ℓ **Challenge** *fs*. This is the only way to populate a response type.

The process $\text{witness } M:T; P$ simply runs P , but is well-typed only if M has the type T . This is the only way to justify a $\text{trust } M:T$ effect.

The process $\text{trust } M \text{ is } (x:T); P$ binds the message M to the variable x of type T , and then runs P . In well-typed programs, this cast is justified by a previous run of a $\text{witness } M:T; Q$ process.

Next, we recall the notions of process safety, opponents, and robust safety introduced in Section 4. The notion of a run of a process can be formalized by an operational semantics.

Safety:

A process P is *safe* if and only if
 for every run of the process and for every L ,
 there is a distinct $\text{begin } L$ event for every $\text{end } L$ event.

Opponents and Robust Safety:

A process P is *assertion-free* if and only if
 it contains no begin- or end-assertions.

A process P is *untyped* if and only if
the only type occurring in P is Un .
An *opponent* O is an assertion-free untyped process.
A process P is *robustly safe* if and only if
 $P \mid O$ is safe for every opponent O .

Our problem, then, is to show that processes representing protocols are robustly safe. We appeal to a type and effect system to establish robust safety (but not to define it). The system involves the following type judgments.

Judgments $E \vdash \mathcal{J}$:

$E \vdash \diamond$	good environment
$E \vdash \epsilon s$	good effect ϵs
$E \vdash T$	good type T
$E \vdash M : T$	good message M of type T
$E \vdash P : \epsilon s$	good process P with effect ϵs

We omit the rules defining these judgments, which can be found in [21]; our previous informal explanation of types should give some intuitions.

We made two additions to the language as defined in [21], namely the empty record type $()$ (and corresponding empty record message $()$), and the conditional form $\text{if } M = N \text{ then } P \text{ else } Q$. The empty record type can be handled by simply extending the typing rules for records to the case where there are no elements. The main consequence of this is that the type $()$ will be isomorphic to the type Un , by the extended subtyping rules. The extension of spi to handle the conditional is similarly straightforward, except that we need to actually add a transition rule to the operational semantics, and a new typing rule to propagate the effects. For completeness, we describe the additions here, with the understanding that they rely on terminology defined and explained in [21]:

Extensions to Spi for the Conditional:

$[\text{if } M = N \text{ then } P_{\text{true}} \text{ else } P_{\text{false}}] + \mathcal{A}s \rightarrow [P_{M=N}] + \mathcal{A}s$	transition rule
(Proc If) $\frac{E \vdash M : \text{Top} \quad E \vdash N : \text{Top} \quad E \vdash P : \epsilon s \quad E \vdash Q : fs}{E \vdash \text{if } M = N \text{ then } P \text{ else } Q : \epsilon s \vee fs}$	typing rule

The type and effect system can guarantee the robust safety of a process, according to the following theorem [21]:

Theorem 7 (Robust Safety). If $x_1:\text{Un}, \dots, x_n:\text{Un} \vdash P : []$ then P is robustly safe.

D. Proofs

D.1. Proof of Theorem 1

A consequence of the types translation for our calculus is that $\llbracket A \rrbracket$ is isomorphic to Un for all types A . Formally,

Lemma 3. $\llbracket A \rrbracket <:\!> \text{Un} \vdash$ for all types A .

In practice, this means that we can replace $\llbracket A \rrbracket$ by Un in type derivations, and vice versa.

Some general remarks on typing are in order. A consequence of Lemma 3, as well as our general use of types, reveals that we rely on typing exclusively to show security properties, not to establish standard safety results. For instance, we do not use types to ensure that the type of the arguments supplied at method invocation match the type of the parameters to the method. Indeed, the only channel type in our translation has itself type Un .

In order to prove Theorem 1, we first establish some lemmas.

Lemma 4.

- (1) If $E \vdash v : A$ then $E_{prin}, \llbracket E \rrbracket \vdash \llbracket v \rrbracket : \llbracket A \rrbracket$.
(2) If $E \vdash a : A$ and $E_0, \llbracket E \rrbracket \vdash p : \text{Prin}$ and $k \notin \text{dom}(E_0, \llbracket E \rrbracket)$ then:

$$E_0, \llbracket E \rrbracket, k:\text{Un} \vdash \llbracket a \rrbracket_k^p : []$$

- (3) If $c \in \text{Class}$ and $\ell \in \text{dom}(\text{methods}(c))$ then $E_0 \vdash I_{class}(c, \ell) : []$.
(4) If $w \in \text{WebService}$ then $E_0 \vdash I_{ws}(w) : []$.

Proof. (1) We prove this by induction on the height of the type derivation for $E \vdash v : A$:

- Case $v = x$: Since $E \vdash x : A$, we must have $x:A \in E$. By definition of the translation for environment, $x:\llbracket A \rrbracket \in \llbracket E \rrbracket$, hence $E_{prin}, \llbracket E \rrbracket \vdash x : \llbracket A \rrbracket$, as required.
- Case $v = \text{null}$: We have $E \vdash \text{null} : c$. Since $\llbracket c \rrbracket = \text{Union}(\text{null}(), c(\text{Un}))$ and $\llbracket \text{null} \rrbracket = \text{null}()$, we have $E_{prin}, \llbracket E \rrbracket \vdash \text{null}() : \text{Union}(\text{null}(\text{Un}), c(\text{Un}))$, as required.
- Case $v = \text{new } c(v_1, \dots, v_n)$: Since $E \vdash v : A$, where $A = c$, we have $\text{fields}(c) = f_i \mapsto A_i^{i \in 1..n}$, and $E \vdash v_i : A_i$ for all $i \in 1..n$. Let $E' = E_{prin}, \llbracket E \rrbracket$. By induction hypothesis, $E' \vdash \llbracket v_i \rrbracket : \llbracket A_i \rrbracket$ for all $i \in 1..n$. We can now derive:

$$\frac{\frac{\frac{E' \vdash \llbracket v_i \rrbracket : \llbracket A_i \rrbracket \quad \forall i \in 1..n}{E' \vdash (\llbracket v_1 \rrbracket, \dots, \llbracket v_n \rrbracket) : (\llbracket A_1 \rrbracket, \dots, \llbracket A_n \rrbracket)}}{E' \vdash \llbracket v_1 \rrbracket, \dots, \llbracket v_n \rrbracket : (\text{Un}, \dots, \text{Un})}}{E' \vdash \llbracket v_1 \rrbracket, \dots, \llbracket v_n \rrbracket : \text{Un}}}{E' \vdash c(\llbracket v_1 \rrbracket, \dots, \llbracket v_n \rrbracket) : \text{Union}(\text{null}(\text{Un}), c(\text{Un}))}$$

- Case $v = p$: Since $E \vdash p : A$ (with $A = \text{Id}$), we have $p \in \text{Prin}$, hence $p:\text{Prin} \in E_{prin}$. Since $\llbracket \text{Id} \rrbracket = \text{Prin}$, we have $E_{prin}, \llbracket E \rrbracket \vdash p : \text{Prin}$, as required.

(2) Again, we proceed by induction on the height of the type derivation for $E \vdash a : A$.

- Case $a = v$: We can apply the result of part (1). Since $E \vdash v : A$, then $E_{prin}, \llbracket E \rrbracket \vdash \llbracket v \rrbracket : \llbracket A \rrbracket$. We can derive:

$$\frac{\frac{E_0, \llbracket E \rrbracket \vdash \llbracket v \rrbracket : \llbracket A \rrbracket}{E_0, \llbracket E \rrbracket, k:\text{Un} \vdash k : \text{Un} \quad E_0, \llbracket E \rrbracket \vdash \llbracket v \rrbracket : \text{Un}}}{E_0, \llbracket E \rrbracket, k:\text{Un} \vdash \text{out } k \llbracket v \rrbracket : []}$$

- Case $a = \text{let } x=a_0 \text{ in } b$: We have $E \vdash a_0 : B$ for some B , and $E, x:B \vdash b : A$. Applying the induction hypothesis, we derive $E_0, \llbracket E \rrbracket, k':\text{Un} \vdash \llbracket a_0 \rrbracket_{k'}^p : []$ and $E_0, \llbracket E \rrbracket, x:\llbracket B \rrbracket, k:\text{Un} \vdash \llbracket b \rrbracket_k^p : []$. Let $E' = E_0, \llbracket E \rrbracket, k:\text{Un}$. We can now derive:

$$\frac{\frac{\frac{E', k':\text{Un} \vdash \llbracket a \rrbracket_{k'}^p : [] \quad \frac{E', k':\text{Un} \vdash k' : \text{Un} \quad \frac{E', k':\text{Un}, x:\llbracket B \rrbracket \vdash \llbracket b \rrbracket_k^p : []}{E', k':\text{Un}, x:\text{Un} \vdash \llbracket b \rrbracket_k^p : []}}{E', k':\text{Un} \vdash \text{inp } k' (x:\text{Un}); \llbracket b \rrbracket_k^p : []}}{E', k':\text{Un} \vdash \llbracket a \rrbracket_{k'}^p \mid \text{inp } k' (x:\text{Un}); \llbracket b \rrbracket_k^p : []}}{E' \vdash \text{new } (k':\text{Un}); (\llbracket a \rrbracket_{k'}^p \mid \text{inp } k' (x:\text{Un}); \llbracket b \rrbracket_k^p) : []}$$

- Case $a = \text{if } u = v \text{ then } a_0 \text{ else } a_1$: We have $E \vdash u : B$, $E \vdash v : B$, $E \vdash a_0 : A$, and $E \vdash a_1 : A$. Applying the induction hypothesis, we derive $E_0, \llbracket E \rrbracket, k:\text{Un} \vdash \llbracket a_0 \rrbracket_k^p : []$ and $E_0, \llbracket E \rrbracket, k:\text{Un} \vdash \llbracket a_1 \rrbracket_k^p : []$. By (1), we also have $E_0, \llbracket E \rrbracket \vdash \llbracket u \rrbracket : \llbracket B \rrbracket$ and $E_0, \llbracket E \rrbracket \vdash \llbracket v \rrbracket : \llbracket B \rrbracket$. This gives us $E_0, \llbracket E \rrbracket, k:\text{Un} \vdash \text{if } \llbracket u \rrbracket = \llbracket v \rrbracket \text{ then } \llbracket a_0 \rrbracket_k^p \text{ else } \llbracket a_1 \rrbracket_k^p : []$, as required.
- Case $a = v.f_j$: We have $E \vdash v.f_j : A_j$, where $E \vdash v : c$ and $\text{fields}(c) = f_i \mapsto A_i^{i \in 1..n}$. By (1), $E_0, \llbracket E \rrbracket \vdash \llbracket v \rrbracket : \llbracket c \rrbracket$. Let $E' = E_0, \llbracket E \rrbracket, k:\text{Un}$. First, let us derive that $E', y : \text{Un} \vdash \text{split } y \text{ is } (x_1:\llbracket A_1 \rrbracket, \dots, x_n:\llbracket A_n \rrbracket)$;

out $k x_j : []$. Let $E'' = x_1 : \llbracket A_1 \rrbracket, \dots, x_n : \llbracket A_n \rrbracket$. (We trim environments where possible to reduce clutter.)

$$\frac{\frac{\frac{E', y : \text{Un}, E'' \vdash x_j : \llbracket A_j \rrbracket}{E' \vdash k : \text{Un}}{E', y : \text{Un}, E'' \vdash x_j : \text{Un}}}{E', y : \text{Un} \vdash y : \text{Un}} \quad \frac{E', y : \text{Un}, E'' \vdash \text{out } k x_j : []}{E', y : \text{Un}, E'' \vdash \text{out } k x_j : []}}{E', y : \text{Un} \vdash \text{split } y \text{ is } (x_1 : \llbracket A_1 \rrbracket, \dots, x_n : \llbracket A_n \rrbracket); \text{out } k x_j : []}$$

We can now derive:

$$\frac{\frac{\frac{E' \vdash \llbracket v \rrbracket : \text{Union}(\text{null}(\text{Un}), c(\text{Un}))}{E', y : \text{Un} \vdash \text{stop} : []}}{E', y : \text{Un} \vdash \text{split } y \text{ is } (x_1 : \llbracket A_1 \rrbracket, \dots, x_n : \llbracket A_n \rrbracket); \text{out } k x_j : []}}{E' \vdash \text{case } \llbracket v \rrbracket \text{ is } \text{null}(y : \text{Un}); \text{stop}} \quad \frac{}{\text{is } c(y); \text{split } y \text{ is } (x_1 : \llbracket A_1 \rrbracket, \dots, x_n : \llbracket A_n \rrbracket); \text{out } k x_j : []}}$$

- Case $a = v.\ell_j(u_1, \dots, u_m)$: We have $E \vdash v.\ell_j(u_1, \dots, u_m) : B$, where $E \vdash v : c$, $\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)_{i \in 1..n}$, $\text{sig}_j = B(A_1 x_1, \dots, A_m x_m)$, and $E \vdash u_k : A_k$ for all $k \in 1..m$. By (1), $E_0, \llbracket E \rrbracket \vdash \llbracket u_k \rrbracket : \llbracket A_k \rrbracket$ for all $k \in 1..m$. Let $E' = E_0, \llbracket E \rrbracket, k : \text{Un}$. First, let us derive that $E', y : \text{Un} \vdash \text{out } c.\ell(p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k) : []$.

$$\frac{E', y : \text{Un} \vdash c.\ell : \text{Un} \quad E', y : \text{Un} \vdash (p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k) : \text{Un}}{E', y : \text{Un} \vdash \text{out } c.\ell(p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k) : []}$$

We can derive:

$$\frac{\frac{\frac{E' \vdash \llbracket v \rrbracket : \text{Union}(\text{null}(\text{Un}), c(\text{Un}))}{E', y : \text{Un} \vdash \text{stop} : []}}{E', y : \text{Un} \vdash \text{out } c.\ell(p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k) : []}}{E' \vdash \text{case } \llbracket v \rrbracket \text{ is } \text{null}(y : \text{Un}); \text{stop}} \quad \frac{}{\text{is } c(y); \text{out } c.\ell(p, \llbracket v \rrbracket, \llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket, k) : []}}$$

- Case $a = w.\ell_j(u_1, \dots, u_m)$: We have $E \vdash w.\ell_j(u_1, \dots, u_m) : B$, where $\text{class}(w) = c$, $\text{owner}(w) = q$, $\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i)_{i \in 1..n}$, $\text{sig}_j = B(A_1 x_1, \dots, A_m x_m)$, and $E \vdash u_k : A_k$ for all $k \in 1..m$. By (1), $E_0, \llbracket E \rrbracket \vdash \llbracket u_k \rrbracket : \llbracket A_k \rrbracket$ for all $k \in 1..m$. Rather than giving the full type derivation for the translation of a web service call, we outline the derivation of effects:

```

new (k1:Un, k2:Un, t:Un, np:Public Challenge []);
// Effect: [check Public np]
begin req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t);
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
out w (req(getnonce()), k1);
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
inp k1 (res(getnonce(nq:Un)));
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
cast nq is (n'q:Public Response [end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]);
// Effect: [check Public np]
out w (p, {req(w, ℓ(⟦u1⟧, ..., ⟦un⟧), t, n'q)}Kpq, np, k2);
// Effect: [check Public np]
inp k2 (q':Un, bdy:Un);
// Effect: [check Public np]
decrypt bdy is {res(plain)}Kpq;
// Effect: [check Public np]
match plain is (w, rest:
  (r:Res(w), t':Un, Public Response [end res(p, q, w, r, t')]));
// Effect: [check Public np]
split rest is (r:Res(w), rest':
  (t':Un, n'p:Public Response [end res(p, q, w, r, t')]));
// Effect: [check Public np]
match rest' is (t, n'p:Public Response [end res(p, q, w, r, t)]);
// Effect: [check Public np]
check np is n'p;
// Effect: [end res(p, q, w, r, t)]
end res(p, q, w, r, t);
// Effect: []
case r is ℓ(x); out k x
// Effect: []

```

- (3) Recall that we assume that method bodies are well-typed, that is, we assume for c, ℓ_j with $methods(c) = \ell_i \mapsto (sig_i, b_i)$ and $sig_j = B(A_1 a_1, \dots, A_m a_m)$, that $this:c, x_1:A_1, \dots, x_m:A_m \vdash b_j : B$. By clause (2) above, this means that $E_0, this:\llbracket c \rrbracket, x_1:\llbracket A_1 \rrbracket, \dots, x_m:\llbracket A_m \rrbracket, k:\mathbf{Un} \vdash \llbracket b_j \rrbracket_k^p : []$. Applying Lemma 3, we derive $E_0, this:\mathbf{Un}, x_1:\mathbf{Un}, \dots, x_m:\mathbf{Un}, k:\mathbf{Un} \vdash \llbracket b_j \rrbracket_k^p : []$. We can now easily derive the following:

$$\frac{
\frac{
\frac{
E_0, z:\mathbf{Un} \vdash z : (\text{Prin}, \mathbf{Un}, \dots, \mathbf{Un})
}{E_0, z:\mathbf{Un}, p:\text{Prin}, this:\mathbf{Un}, x_1:\mathbf{Un}, \dots, x_n:\mathbf{Un}, k:\mathbf{Un} \vdash \llbracket b_j \rrbracket_k^p : []}
}{E_0 \vdash c.\ell : \mathbf{Un} \quad E_0, z:\mathbf{Un} \vdash \text{split } z \text{ is } (p:\text{Prin}, this:\mathbf{Un}, x_1:\mathbf{Un}, \dots, x_n:\mathbf{Un}, k:\mathbf{Un}); \llbracket b_j \rrbracket_k^p : []}
}{E_0 \vdash \text{repeat } \text{inp } c.\ell (z); \text{split } z \text{ is } (p:\text{Prin}, this:\mathbf{Un}, x_1:\mathbf{Un}, \dots, x_n:\mathbf{Un}, k:\mathbf{Un}); \llbracket b_j \rrbracket_k^p : []}
}{E_0 \vdash I_{\text{class}}(c, \ell) : []}$$

- (4) Let $w \in \text{WebService}$, with $owner(w) = q$. First, note that the following derivation is admissible:

$$\frac{
E_0, E \vdash p : \text{Prin} \quad E_0, E \vdash a : \text{Req}(w) \quad E_0, E, r:\text{Res}(w) \vdash P : es
}{E_0, E \vdash \text{let } r:\text{Res}(w)=\text{call}_w(p, a); P : es}$$

(The proof is a straightforward, if longish, type derivation.) Rather than giving the full type derivation for the implementation of web service w , we outline the derivation of effects:

```

repeat inp  $w$  ( $bdy:Un, k_1:Un$ );
// Effect: []
case  $bdy$  is  $req(getnonce())$ ;
// Effect: []
new ( $n_q:Public$  Challenge []);
// Effect: [check Public  $n_q$ ]
out  $k_1$  ( $res(getnonce(n_q))$ );
// Effect: [check Public  $n_q$ ]
inp  $w$  ( $p':Un, cipher:Un, n_p:Un, k_2:Un$ );
// Effect: [check Public  $n_q$ ]
 $\prod_{p \in Prin}$  if  $p = p'$  then
// Effect: [check Public  $n_q$ ]
decrypt  $cipher$  is  $\{req(plain)\}_{K_{pq}}$ ;
// Effect: [check Public  $n_q$ ]
match  $plain$  is ( $w, rest:(a:Req(w), t:Un,$ 
Public Response [end  $req(p, q, w, a, t)]$ );
// Effect: [check Public  $n_q$ ]
split  $rest$  is ( $a:Req(w), t:Un, n'_q:Public$  Response [end  $req(p, q, w, a, t)]$ );
// Effect: [check Public  $n_q$ ]
check  $n_q$  is  $n'_q$ ;
// Effect: [end  $req(p, q, w, a, t)$ ]
end  $req(p, q, w, a, t)$ ;
// Effect: []
let  $r:Res(w)=call_w(p, a)$ ;
// Effect: []
begin  $res(p, q, w, r, t)$ ;
// Effect: [end  $res(p, q, w, r, t)$ ]
cast  $n_p$  is ( $n'_p:Public$  Response [end  $res(p, q, w, r, t)$ ]);
// Effect: []
out  $k_2$  ( $q, \{res(w, r, t, n'_p)\}_{K_{pq}}$ );
// Effect: []

```

□

Lemma 5. If $\emptyset \vdash a : A$ and $p \in Prin$ and $k \notin dom(E_0)$ then:

$$E_{ws}, E_{prin} \vdash \text{new } (E_{class}, E_{keys}); (I_{class} \mid I_{ws} \mid \text{new } (k:Un); \llbracket a \rrbracket_k^p) : []$$

Proof. This is a corollary of Lemma 4. Specifically, we can derive:

$$\frac{\frac{E_0 \vdash I_{class}(c, \ell) : [] \quad (c, \ell) \in ClMeth}{E_0 \vdash I_{class} : []} \quad \frac{E_0 \vdash I_{ws}(w) : [] \quad w \in WebService}{E_0 \vdash I_{ws} : []} \quad \frac{E_0, k:Un \vdash \llbracket a \rrbracket_k^p : []}{E_0 \vdash \text{new } (k:Un); \llbracket a \rrbracket_k^p : []}}{\frac{E_0 \vdash I_{class} \mid I_{ws} \mid \text{new } (k:Un); \llbracket a \rrbracket_k^p : []}{E_{ws}, E_{prin} \vdash \text{new } (E_{class}, E_{keys}); (I_{class} \mid I_{ws} \mid \text{new } (k:Un); \llbracket a \rrbracket_k^p) : []}}$$

□

We can now prove Theorem 1.

Theorem. If $\emptyset \vdash a : A$ and $p \in Prin$ and $k \notin dom(E_0)$ then the system

$$\text{new } (E_{class}, E_{keys}); (I_{class} \mid I_{ws} \mid \text{new } (k:Un); \llbracket a \rrbracket_k^p)$$

is robustly safe.

Proof. By Lemma 5,

$$E_{ws}, E_{prin} \vdash \text{new } (E_{class}, E_{keys}); (I_{class} \mid I_{ws} \mid \text{new } (k:Un); \llbracket a \rrbracket_k^p) : [].$$

Robust safety of the system follows by Theorem 7.

□

D.2. Proof of Theorem 2

Theorem. If $\emptyset \vdash a : A$ and $p \in \text{Prin}$ and $k \notin \text{dom}(E_0)$ then the system

$$\text{new } (E_{\text{class}}, E_{\text{keys}}); (I_{\text{net}} \mid I_{\text{class}} \mid I_{\text{ws}} \mid \text{new } (k:\text{Un}); \llbracket a \rrbracket_k^P)$$

is robustly safe.

Proof. Rather than giving a full proof, we point out the parts of the proof of Theorem 1 that need to be updated. Essentially, we need to show that the new semantics for web method invocations is effect-free, and similarly for the new implementation of web services. These occur in the proof of Lemma 4, part (2) and (4).

As we did in Lemma 4, rather than giving the full type derivation for the translation of a web service call, we outline the derivation of effects:

```

new (k1:Un, k2:Un, t:Un, np:Public Challenge []);
// Effect: [check Public np]
begin req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t);
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
out w (CertVKp, np, req(getnonce()), k1);
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
inp k1 (c:Un, res(getnonce(nq:Un)));
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
decrypt c is {cert:(q':Un, Decrypt Key(AuthMsg(q')))}VKCA-1;
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
match cert is (q, vkq:Decrypt Key(AuthMsg(q)));
// Effect: [check Public np, end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]
cast nq is (n'q:Public Response [end req(p, q, w, ℓ(⟦u1⟧, ..., ⟦un⟧), t)]);
// Effect: [check Public np]
out w (p, {req(w, ℓ(⟦u1⟧, ..., ⟦un⟧), t, q, n'q)}SKP, k2);
// Effect: [check Public np]
inp k2 (q'':Un, bdy:Un);
// Effect: [check Public np]
decrypt bdy is {res(plain:(w':Un, r:Un, t':Un, p':Un,
Public Response [end res(p', q, w', r, t')])vkq-1);
// Effect: [check Public np]
match plain is (w, rest:(r:Res(w), t':Un, p':Un,
Public Response [end res(p', q, w, r, t')])];
// Effect: [check Public np]
split rest is (r:Res(w), rest':(t':Un, p':Un,
Public Response [end res(p', q, w, r, t')])];
// Effect: [check Public np]
match rest' is (t, rest'':(p':Un, Public Response [end res(p', q, w, r, t')])];
// Effect: [check Public np]
match rest'' is (p, n'p:Public Response [end res(p, q, w, r, t)]);
// Effect: [check Public np]
check np is n'p;
// Effect: [end res(p, q, w, r, t)]
end res(p, q, w, r, t);
// Effect: []
case r is ℓ(x); out k x
// Effect: []

```

For the new implementation of web service w , rather than giving the full type derivation, we outline the derivation of effects:

```

repeat inp  $w$  ( $c:\text{Un}, n_p:\text{Un}, bdy:\text{Un}, k_1:\text{Un}$ );
// Effect: []
case  $bdy$  is  $req(getnonce())$ ;
// Effect: []
decrypt  $c$  is  $\{p:\text{Un}, vkp:\text{Decrypt Key}(\text{AuthMsg}(p))\}_{VKCA^{-1}}$ ;
// Effect: []
new ( $n_q:\text{Public Challenge []}$ );
// Effect: [check Public  $n_q$ ]
out  $k_1$  ( $\text{CertVK}q, res(getnonce(n_q))$ );
// Effect: [check Public  $n_q$ ]
inp  $w$  ( $p':\text{Un}, cipher:\text{Un}, k_2:\text{Un}$ );
// Effect: [check Public  $n_q$ ]
if  $p = p'$  then
// Effect: [check Public  $n_q$ ]
decrypt  $cipher$  is  $\{req(plain:(w:\text{Un}, a:\text{Un}, t:\text{Un}, q':\text{Un},$ 
Public Response [end  $req(p, q', w, a, t)])\}_{vkp^{-1}}$ ;
// Effect: [check Public  $n_q$ ]
match  $plain$  is ( $w, rest:(a:\text{Req}(w), t:\text{Un}, q':\text{Un},$ 
Public Response [end  $req(p, q', w, a, t)])$ );
// Effect: [check Public  $n_q$ ]
split  $rest$  is ( $a:\text{Req}(w),$ 
 $t:\text{Un}, rest':(q':\text{Un}, \text{Public Response [end } req(p, q', w, a, t)]$ );
// Effect: [check Public  $n_q$ ]
match  $rest'$  is ( $q, n'_q:\text{Public Response [end } req(p, q, w, a, t)]$ );
// Effect: [check Public  $n_q$ ]
check  $n_q$  is  $n'_q$ ;
// Effect: [end  $req(p, q, w, a, t)$ ]
end  $req(p, q, w, a, t)$ ;
// Effect: []
let  $r:\text{Res}(w)=\text{call}_w(p, a)$ ;
// Effect: []
begin  $res(p, q, w, r, t)$ ;
// Effect: [end  $res(p, q, w, r, t)$ ]
cast  $n_p$  is ( $n'_p:\text{Public Response [end } res(p, q, w, r, t)]$ );
// Effect: []
out  $k_2$  ( $q, \{res(w, r, t, p, n'_p)\}_{SKq}$ );
// Effect: []

```

□

D.3. Proof of Theorem 3

Theorem. If $\emptyset \vdash a : A$ and $p \in \text{Prin}$ and $k \notin \text{dom}(E_0)$ then the system

$$\text{new } (E_{\text{class}}, E_{\text{keys}}); (I_{\text{net}} \mid I_{\text{class}} \mid I_{\text{ws}} \mid \text{new } (k:\text{Un}); \llbracket a \rrbracket_k^p)$$

is robustly safe.

Proof. Rather than giving a full proof, we point out the parts of the proof of Theorem 1 that need to be updated. Essentially, we need to show that the new semantics for web method invocations is effect-free, and similarly for the new implementation of web services. These occur in the proof of Lemma 4, part (2) and (4).

As we did in Lemma 4, rather than giving the full type derivation for the translation of a web service call, we outline the derivation of effects:

```

new (k1:Un, k2:Un, t:Un, np:Public Challenge []);
// Effect: [check Public np]
begin req(p, q, w, ℓ([u1], ..., [un]), t);
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)]
out w (CertEKp, req(getnonce()), k1);
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)] inp k1 (c:Un, cipher:Un, res(getnonce(nq:Un)));
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)]
decrypt c is {cert:(q':Un, Encrypt Key(AuthEncMsg(q')))}VKCA-1;
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)]
match cert is (q, ekq:Encrypt Key(AuthEncMsg(q)));
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)]
decrypt cipher is {msg2(q':Un, nK:Un)}DKP-1;
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)]
if q = q' then
// Effect: [check Public np, end req(p, q, w, ℓ([u1], ..., [un]), t)]
cast nq is (n'q:Public Response [end req(p, q, w, ℓ([u1], ..., [un]), t)]);
// Effect: [check Public np]
new (K:SKey(p, q, w));
// Effect: [check Public np]
witness K:SKey(p, q, w);
// Effect: [check Public np, trust K:SKey(p, q, w)]
cast nK is (n'K:Private Response [trust K:SKey(p, q, w)]);
// Effect: [check Public np]
out w ({msg3(w, p, t, K, n'K)}ekq, np, {req(w, ℓ([u1], ..., [un]), t, n'q)K, k2});
// Effect: [check Public np]
inp k2 (bdy:Un);
// Effect: [check Public np]
decrypt bdy is {res(plain:(r:Res(w), t':Un,
Public Response [end res(p, q, w, r, t')])K);
// Effect: [check Public np]
match plain is (r:Res(w), rest:(t':Un, Public Response [end res(p, q, w, r, t')])];
// Effect: [check Public np]
match rest is (t, n'p:Public Response [end res(p, q, w, r, t)]);
// Effect: [check Public np]
check np is n'p;
// Effect: [end res(p, q, w, r, t)]
end res(p, q, w, r, t);
// Effect: []
case r is ℓ(x); out k x
// Effect: []

```

For the new implementation of web service w , rather than giving the full type derivation, we outline the derivation of effects:

```

repeat inp w (c:Un, bdy:Un, k1:Un);
// Effect: []
case bdy is req(getnonce());
// Effect: []
decrypt c is {p:Un, ekp:Encrypt Key(AuthEncMsg(p))}VKCA-1;
// Effect: []
new (nq:Public Challenge []);
// Effect: [check Public nq]
new (nK:Private Challenge []);
// Effect: [check Public nq, check Private nK]
out k1 (CertEKq, {msg2(q, nK)ekp, res(getnonce(nq))});
// Effect: [check Public nq, check Private nK]

```

```

inp w (cipher1:Un, np:Un, cipher2:Un, k2:Un);
// Effect: [check Public nq, check Private nK]
decrypt cipher1
  is {msg3(plain1:(w:Un, p':Un, K:Top,
    Private Response [trust K:SKey(p', q, w)]))}_{DKq-1};
// Effect: [check Public nq, check Private nK]
match plain1 is (w, rest:(p':Un, K:Top,
  Private Response [trust K:SKey(p', q, w)]));
// Effect: [check Public nq, check Private nK]
match rest is (p, rest':(K:Top, n'K:Private Response [trust K:SKey(p, q, w)]));
// Effect: [check Public nq, check Private nK]
split rest' is (K:Top, n'K:Private Response [trust K:SKey(p, q, w)]);
// Effect: [check Public nq, check Private nK]
check nK is n'K;
// Effect: [check Public nq, trust K:SKey(p, q, w)]
trust K is (K':SKey(p, q, w));
// Effect: [check Public nq]
decrypt cipher2 is {req(plain2:(a:Req(w), t:Un,
  Public Response [end req(p, q, w, a, t)]))}_{K'};
// Effect: [check Public nq]
split plain2 is (a:Req(w), t:Un, n'q:Public Response [end req(p, q, w, a, t)]);
// Effect: [check Public nq]
check nq is n'q;
// Effect: [end req(p, q, w, a, t)]
end req(p, q, w, a, t);
// Effect: []
let r:Res(w)=callw(p, a);
// Effect: []
begin res(p, q, w, r, t);
// Effect: [end res(p, q, w, r, t)]
cast np is (n'p:Public Response [end res(p, q, w, r, t)]);
// Effect: []
out k2 {res(r, t, n'p)}_{K'}
// Effect: []

```

□

E. First-Class Web Services

The model of web services captured by our calculus in Section 3 does not consider web services to be values. This reflects the fact that current WSDL does not allow for web services to be passed as requests or results. On the other hand, a web service has a simple representation as a string, namely the URL used to access the web service, and this string *can* be passed as a request or a result. Hence, it is possible, in a sense, to pass web services as values given the current web services infrastructure. In this section, we explore an extension of our object calculus that allows web services as first-class values. The main point here is to show that there is no real difficulty in modelling this aspect of the web services infrastructure. Our main result is type safety. We expect it would be straightforward to translate this extended calculus into the spi-calculus, but we do not describe this in detail.

For the sake of keeping this section essentially self-contained, we give the full syntax and semantics of the extended object calculus.

E.1. Syntax

We assume finite sets *Prin*, *WebService*, *Class*, *Field*, *Meth* of principal, web service, class, field, and method names, respectively.

Classes, Fields, Methods, Principals, Web Services:

$c \in Class$	class name
$f \in Field$	field name
$\ell \in Meth$	method name
$p \in Prin$	principal name
$w \in WebService$	web service name

There are now three kinds of data type: Id is the type of principal identifiers, $c \in Class$ is the type of instances of class c , and $WS(c)$ is the type of web services with implementation class $c \in Class$. A method signature specifies the types of its arguments and result.

Types and Method Signatures:

$A, B \in Type ::=$	type
Id	principal identifier
c	object
$WS(c)$	web service
$sig \in Sig ::= B(A_1 x_1, \dots, A_n x_n)$	method signature (x_i distinct)

As in Section 3, an execution environment defines the services and code available in the distributed system.

Execution Environment: ($fields, methods, owner, class$)

$fields \in Class \rightarrow (Field \xrightarrow{fin} Type)$	fields of a class
$methods \in Class \rightarrow (Meth \xrightarrow{fin} Sig \times Body)$	methods of a class
$owner \in WebService \rightarrow Prin$	service owner
$class \in WebService \rightarrow Class$	service implementation

The owner and implementation class of a web service need not be globally known. We can assume that the representation of a web service w carries representations of its owner and its implementation class, which $class$ and $owner$ simply read off. Since we assume web services are given, and we do not provide for ways to actually create new web services, there is no loss of generality in taking this particular approach.

The syntax of method bodies and values is that of the original object calculus, with the differences that web services are values, and that we do not assume that web service invocations require a fixed web service.

Values and Method Bodies:

x, y, z	name: variable, argument
$u, v \in Value ::=$	value
x	variable
$null$	null
$new\ c(v_1, \dots, v_n)$	object
p	principal identifier
w	web service
$a, b \in Body ::=$	method body
v	value
$let\ x=a\ in\ b$	let-expression
$if\ u = v\ then\ a\ else\ b$	conditional
$v.f$	field lookup
$v.\ell(u_1, \dots, u_n)$	method call
$v:\ell(u_1, \dots, u_n)$	service call
$p[a]$	body a running as p

We again require a method body of the form $p[a]$, meaning p running body a , to keep track of which principal is running a method body in the upcoming operational semantics.

E.2. Operational Semantics

The operational semantics is defined by a transition relation, written $a \rightarrow^p a'$, where a and a' are method bodies, and p is the principal evaluating the body a .

Transitions:

(Red Let 1) $\frac{a \rightarrow^p a'}{\text{let } x=a \text{ in } b \rightarrow^p \text{let } x=a' \text{ in } b}$	(Red Let 2) $\frac{}{\text{let } x=v \text{ in } b \rightarrow^p b\{x \leftarrow v\}}$	(Red If) $\frac{}{\text{if } u = v \text{ then } a_{\text{true}} \text{ else } a_{\text{false}} \rightarrow^p a_{u=v}}$
(Red Field) $\frac{\text{fields}(c) = f_i \mapsto A_i \quad i \in 1..n \quad j \in 1..n}{(\text{new } c(v_1, \dots, v_n)).f_j \rightarrow^p v_j}$		
(Red Invoke)(where $v = \text{new } c(v_1, \dots, v_n)$) $\frac{\text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i) \quad i \in 1..n \quad j \in 1..n \quad \text{sig}_j = B(A_1 x_1, \dots, A_m x_m)}{v.\ell_j(u_1, \dots, u_m) \rightarrow^p b_j\{\text{this} \leftarrow v, x_k \leftarrow u_k \quad k \in 1..m\}}$		
(Red Remote) $\frac{\text{owner}(w) = q \quad \text{class}(w) = c}{w:\ell(u_1, \dots, u_n) \rightarrow^p q[\text{new } c(p).\ell(u_1, \dots, u_n)]}$	(Red Prin 1) $\frac{}{a \rightarrow^q a'}$	(Red Prin 2) $\frac{}{q[v] \rightarrow^p v}$

E.3. Type System

The judgments of our type system all depend on an *environment* E , that defines the types of all variables in scope. An environment takes the form $x_1:A_1, \dots, x_n:A_n$ and defines the type A_i for each variable x_i . The domain $\text{dom}(E)$ of an environment E is the set of variables whose types it defines.

Environments:

$D, E ::=$	environment
\emptyset	empty
$E, x:A$	entry
$\text{dom}(x_1:A_1, \dots, x_n:A_n) \triangleq \{x_1, \dots, x_n\}$	domain of an environment

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

Judgments $E \vdash \mathcal{J}$:

$E \vdash \diamond$	good environment
$E \vdash a : A$	good expression a of type A

We write $E \vdash \mathcal{J}$ when we want to talk about both kinds of judgments, where \mathcal{J} stands for either \diamond or $a : A$.

The following rules define an environment $x_1:A_1, \dots, x_n:A_n$ to be well-formed if each of the names x_1, \dots, x_n are distinct.

Rules for Environments:

(Env \emptyset)	(Env x)(where $x \notin \text{dom}(E)$)
$\emptyset \vdash \diamond$	$E \vdash \diamond$
$\emptyset \vdash \diamond$	$E, x:A \vdash \diamond$

We present the rules for deriving the judgment $E \vdash a : A$ that assigns a type A to a value or method body a . These rules are split into two tables, one for values, and one for method bodies.

Rules for Typing Values:

$\frac{(\text{Val } x) \quad E = E_1, x:A, E_2 \quad E \vdash \diamond}{E \vdash x : A}$	$\frac{(\text{Val null}) \quad E \vdash \diamond}{E \vdash \text{null} : c}$	$\frac{(\text{Val WS}) \quad E \vdash \diamond \quad \text{class}(w) = c}{E \vdash w : WS(c)}$
$\frac{(\text{Val Object}) \quad \text{fields}(c) = f_i \mapsto A_i \quad i \in 1..n \quad E \vdash v_i : A_i \quad \forall i \in 1..n}{E \vdash \text{new } c(v_1, \dots, v_n) : c}$		$\frac{(\text{Val Princ}) \quad E \vdash \diamond}{E \vdash p : Id}$

Rules for Typing Method Bodies:

$\frac{(\text{Body Let}) \quad E \vdash a : A \quad E, x:A \vdash b : B}{E \vdash \text{let } x=a \text{ in } b : B}$	$\frac{(\text{Body If}) \quad E \vdash u : A \quad E \vdash v : A \quad E \vdash a : B \quad E \vdash b : B}{E \vdash \text{if } u = v \text{ then } a \text{ else } b : B}$	
$\frac{(\text{Body Field}) \quad E \vdash v : c \quad \text{fields}(c) = f_i \mapsto A_i \quad i \in 1..n \quad j \in 1..n}{E \vdash v.f_j : A_j}$		
$\frac{(\text{Body Invoke}) \quad E \vdash v : c \quad \text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i) \quad i \in 1..n \quad j \in 1..n \quad \text{sig}_j = B(A_1 x_1, \dots, A_m x_m) \quad E \vdash u_k : A_k \quad \forall k \in 1..m}{E \vdash v.\ell_j(u_1, \dots, u_m) : B}$		
$\frac{(\text{Body Remote}) \quad E \vdash v : WS(c) \quad \text{methods}(c) = \ell_i \mapsto (\text{sig}_i, b_i) \quad i \in 1..n \quad j \in 1..n \quad \text{sig}_j = B(A_1 x_1, \dots, A_m x_m) \quad E \vdash u_i : A_i \quad \forall i \in 1..m}{E \vdash v.\ell_j(u_1, \dots, u_m) : B}$		$\frac{(\text{Body Princ}) \quad E \vdash a : A}{E \vdash p[a] : A}$

We make the following assumption on the execution environment.

Assumptions on the Execution Environment:

- | |
|---|
| <ol style="list-style-type: none"> (1) For each $w \in \text{WebService}$, $\text{fields}(\text{class}(w)) = \text{CallerId} : Id$. (2) No tagged expression $p[a]$ occurs within the body of any method; such expressions occur only at runtime, to track the call stack of principals. (3) for each $c \in \text{Class}$ and each $\ell \in \text{dom}(\text{methods}(c))$, if $\text{methods}(c)(\ell) = (B(A_1 x_1, \dots, A_n x_n), b)$, then $\text{this}:c, x_1:A_1, \dots, x_n:A_n \vdash b : B$. |
|---|

We can establish the soundness of the type system of this extended object calculus by essentially the same way we established the soundness of the type system of the original object calculus. Recall that a method body is null-blocked if it is of the form $\text{null}.f_j$, $\text{null}.\ell(u_1, \dots, u_n)$, $\text{let } x=a \text{ in } b$ (where a is null-blocked), or $q[a]$ (where a is null-blocked). A method body is stuck if a is not a value, a is not null-blocked, and there is no a' and p such that $a \rightarrow^p a'$. We write $a \rightarrow^* a'$ to mean that there exists a sequence a_1, \dots, a_n and principals p_1, \dots, p_{n+1} such that $a \rightarrow^{p_1} a_1 \rightarrow^{p_2} \dots \rightarrow^{p_n} a_n \rightarrow^{p_{n+1}} a'$.

Theorem 8 (Soundness). If $\emptyset \vdash a : A$, and $a \rightarrow^* a'$, then a' is not stuck.

Proof. A straightforward adaptation of the proof of Theorem 6, via corresponding Preservation and Progress theorems. \square

To illustrate the usefulness of first-class web services, consider the following simple example, where the fact that web services can be passed as arguments to methods is quite natural. Suppose, as we did in Section 3, that there are two principals $Alice, Bob \in \text{Prin}$, and a web service $\text{cal} = \text{http://mycalendar.com/CalendarService}$,

where we have $class(cal) = CalendarServiceClass$. The web service cal maintains an appointment calendar for principals. It offers web methods to query a principal's calendar for a free time slot, and to reserve time slots. More precisely, the service has the following interface:

```
class CalendarServiceClass
  Id CallerId
  Bool Available(Id account, Time from, Time to)
    ⟨check if selected time slot is free for account⟩
  Void Reserve(Id account, Time from, Time to)
    ⟨reserve time slot for account⟩
```

(We assume that the classes *Bool*, *Time*, and *Void* are provided in the execution environment. The details of their implementation are irrelevant to our discussion.)

Suppose that Alice has an account on cal , and that she wants to make an appointment with a calendar-enabled banking service—that is, a banking service that offers a web method for scheduling appointments with a bank advisor via a calendar service. Consider a calendar-enabled version of the banking service of Section 3. Let $w = http://bob.com/BankingService$, where we have $owner(w) = Bob$ and $class(w) = BankingServiceClass$. We add a web method *MakeAppt* to *BankingServiceClass* that takes as argument a time period during which the appointment is sought, and a calendar service that the banking service can query to confirm that a common free time slot is available between the client and the bank advisor. The interface of the augmented banking service is as follows:

```
class BankingServiceClass
  Id CallerId
  Num Balance(Num account)
    if account = 12345 then
      if this.CallerId = Alice then 100 else null
    else ...
  Time MakeAppt(Time from, Time to, WS(CalendarService) cs)
    ... cs.Available(CallerId, ...) ...
```

Hence, if Alice wants to make an appointment sometime within the next week, she could issue the web method call $w:MakeAppt(18/11/02:08:00, 23/11/02:17:00, cal)$. (We assume appropriate syntax for constants of type *Time*.) During the evaluation of this web method invocation, the implementation of *MakeAppt* will make calls to $cal:Available$ to find a time slot suitable to Alice, and finally a call to $cal:Reserve$ to reserve a time slot. A principal with an account on a different calendar service c would call $w:MakeAppt$ passing in c as the calendar service.

References

- [1] M. Abadi and L. Cardelli. *A Theory of Objects*. Springer, 1996.
- [2] M. Abadi, C. Fournet, and G. Gonthier. Secure communications implementation of channel abstractions. In *13th IEEE Symposium on Logic in Computer Science (LICS'98)*, pages 105–116, 1998.
- [3] M. Abadi, C. Fournet, and G. Gonthier. Secure communications processing for distributed languages. In *IEEE Computer Society Symposium on Research in Security and Privacy*, pages 74–88, 1999.
- [4] M. Abadi, C. Fournet, and G. Gonthier. Authentication primitives and their compilation. In *27th ACM Symposium on Principles of Programming Languages (POPL'00)*, pages 302–315, 2000.
- [5] M. Abadi and A.D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148:1–70, 1999.
- [6] B. Atkinson, G. Della-Libera, S. Hada, M. Hondo, P. Hallam-Baker, C. Kaler, J. Klein, B. LaMacchia, P. Leach, J. Manfredelli, H. Maruyama, A. Nadalin, N. Nagaratnam, H. Prafullchandra, J. Shewchuk, and D. Simon. Web services security (WS-Security), version 1.0. Available from <http://msdn.microsoft.com/library/en-us/dnglobspec/html/ws-security.asp>, April 2002.
- [7] D. Balfanz, D. Dean, and M. Spreitzer. A security infrastructure for distributed Java applications. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 15–26. IEEE Computer Society Press, 2000.
- [8] T. Barclay, J. Gray, E. Strand, S. Ekblad, and J. Richter. TerraService.NET: An introduction to web services. Technical Report MS-TR-2002-53, Microsoft Research, June 2002.
- [9] K. Bhargavan, C. Fournet, and A. D. Gordon. A semantics for web services authentication. In *31st ACM Symposium on Principles of Programming Languages (POPL'04)*, pages 198–209, 2004. An extended version appears as Microsoft Research Technical Report MSR-TR-2003-83.

- [10] A. D. Birrell. Secure communication using remote procedure calls. *ACM Transactions on Computer Systems*, 3(1):1–14, 1985.
- [11] D. Box. *Essential COM*. Addison Wesley Professional, 1997.
- [12] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. Nielsen, S. Thatte, and D. Winer. Simple object access protocol (SOAP) 1.1. Available from <http://www.w3.org/TR/SOAP>, 2000.
- [13] L. Cardelli and A.D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240:177–213, 2000.
- [14] E. Christensen, F. Curbera, G. Meredith, and S. Weerawarana. Web services description language (WSDL) 1.2. Available from <http://www.w3.org/TR/2002/WD-wsd112-20020709>, 2002.
- [15] E. Damiani, S. De Capitani di Vimercati, S. Paraboschi, and P. Samarati. Securing SOAP e-services. *International Journal of Information Security (IJIS)*, 1(2):100–115, 2002.
- [16] R. De Nicola, G. Ferrari, and R. Pugliese. Types as specifications of access policies. In *Secure Internet Programming 1999*, volume 1603 of *Lecture Notes in Computer Science*, pages 117–146. Springer, 1999.
- [17] D. Dolev and A.C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, IT-29(2):198–208, 1983.
- [18] D. Duggan. Cryptographic types. In *15th IEEE Computer Security Foundations Workshop*, pages 238–252. IEEE Computer Society Press, 2002.
- [19] P. Eronen and P. Nikander. Decentralized Jini security. In *Proceedings of Network and Distributed System Security 2001 (NDSS2001)*, pages 161–172, 2001.
- [20] Google. Google Web APIs (beta). <http://www.google.com/apis>, July 2002.
- [21] A.D. Gordon and A. Jeffrey. Types and effects for asymmetric cryptographic protocols. In *15th IEEE Computer Security Foundations Workshop*, pages 77–91. IEEE Computer Society Press, 2002. An extended version appears as Technical Report MSR-TR-2002-31, Microsoft Research, August 2002.
- [22] A.D. Gordon and A. Jeffrey. Authenticity by typing for security protocols. *Journal of Computer Security*, 11(4):451–521, 2003.
- [23] A.D. Gordon and A. Jeffrey. Typing correspondence assertions for communication protocols. *Theoretical Computer Science*, 300:379–409, 2003.
- [24] A.D. Gordon and R. Pucella. Validating a web service security abstraction by typing. In *2002 ACM Workshop on XML Security*, pages 18–29, 2002.
- [25] 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.
- [26] M. Hennessy and J. Riely. Resource access control in systems of mobile agents. In *Proceedings HLCL'98*, volume 16(3) of *Electronic Notes in Theoretical Computer Science*. Elsevier, 1998.
- [27] D. Hoshina, E. Sumii, and A. Yonezawa. A typed process calculus for fine-grained resource access control in distributed computation. In *Fourth International Symposium on Theoretical Aspects of Computer Software (TACS2001)*, volume 2215 of *Lecture Notes in Computer Science*, pages 64–81. Springer, 2001.
- [28] IBM Corporation and Microsoft Corporation. Security in a web services world: A proposed architecture and roadmap. White paper available from <http://msdn.microsoft.com/library/en-us/dnwssecur/html/securitywhitepaper.asp>, April 2002.
- [29] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *Object Oriented Programming: Systems, Languages and Applications (OOPSLA '99)*, pages 132–146. ACM Press, 1999.
- [30] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, 1992.
- [31] U. Lang and R. Schreiner. *Developing Secure Distributed Systems with CORBA*. Artech House, 2002.
- [32] R. Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999.
- [33] P. Sewell. Global/local subtyping and capability inference for a distributed π -calculus. In *25th International Colloquium on Automata, Languages, and Programming (ICALP'98)*, volume 1443 of *Lecture Notes in Computer Science*, pages 695–706. Springer, 1998.
- [34] E. G. Sirer and K. Wang. An access control language for web services. In *Proceedings of the ACM Symposium on Access Control Models and Technologies*, pages 23–30. ACM Press, 2002.
- [35] L. van Doorn, M. Abadi, M. Burrows, and E. Wobber. Secure network objects. In *IEEE Computer Society Symposium on Research in Security and Privacy*, pages 211–221, 1996.
- [36] T. Wobber, M. Abadi, M. Burrows, and B. Lampson. Authentication in the Taos operating system. *ACM Transactions on Computer Systems*, 12(1):3–32, 1994.
- [37] T.Y.C. Woo and S.S. Lam. A semantic model for authentication protocols. In *IEEE Computer Society Symposium on Research in Security and Privacy*, pages 178–194, 1993.