

# Type-Based Distributed Access Control

Tom Chothia

Dept of Computer Science  
Stevens Institute of Technology  
Hoboken, NJ, USA.

Dominic Duggan

Dept of Computer Science  
Stevens Institute of Technology  
Hoboken, NJ, USA.

Jan Vitek

Dept of Computer Sciences  
Purdue University  
West Lafayette, IN, USA.

## Abstract

*A type system is presented that combines a weak form of information flow control, termed distributed access control in the paper, with typed cryptographic operations. The motivation is to have a type system that ensures access control while giving the application the responsibility to secure network communications, and to do this safely. The notion of declassification certificates is introduced to support the declassification of encrypted data.*

## 1. Introduction

The notion of *distributed access control* defined in this paper is a weak form of information flow control that associates access restrictions to data

Like information flow control, it is important that distributed access control be enforced statically, at compile-time, though for performance reasons in the case of distributed access control<sup>1</sup>. The motivation for distributed access control is to enforce *accountability*, tracking which principals are responsible for allowing access to particular data.

A challenge with static checking of access control is how to reconcile type-based access control specifications with application-based network security. In current systems that do this static enforcement, network communications are secured by hiding the network inside the trusted computing base (TCB), using remote method invocation (RMI) or some similar mechanism to make the network transparent. In general this is unsatisfactory for the following three reasons: (i) The TCB is greatly enlarged by pushing the middleware into it, increasing the risk of security flaws due to bugs in the TCB code. (ii) In Internet programming, there is some skepticism

<sup>1</sup>Distributed access control is completely impractical, as explained below, unless some large part of it can be done statically, with dynamic checking left to the point where data is unmarshalled.

that the network can be made transparent, and many systems (e.g. some banking systems) that are required to be secure and fault-tolerant are implemented using lower-level mechanisms than RMI because of this. (iii) Relying on the TCB to secure all communications violates the end-to-end principle in systems design. In many cases it will be more appropriate to have the application ensure network security rather than relying on the communication system.

The contribution of this paper is to show how network security can be moved out of the TCB into the application, in a system where distributed access control is enforced through the type system. All current approaches to statically checked distributed access control and information flow control require the network to be secured in the TCB. In short we want to:

1. Use the type system to enforce access control, including after we have handed off data to others, in a distributed environment.
2. Allow applications to secure network communication themselves using cryptographic techniques.
3. Ensure that the requirements of the access control specifications are not violated by the application-based network security (without considering covert channels).

**Distributed Access Control.** The traditional approach to access control is to provide the data on a database on the network, restrict access using access control lists (ACLs), and use authentication to verify the allowability of access requests from remote principals. An example of this is the Taos operating system [9, 3, 8]. In recent years a trend has been to replace or supplement ACLs with delegation certificates, which do not require all accessing principals to be named in an ACL, but puts some limit on the extent to which access rights can be delegated. An example of this is the Simple Public Key Infrastructure (SPKI) [17], as well as various systems for

decentralized trust management (PolicyMaker, Keynote, etc) [12, 11].

We refer to these approaches to access control as *local access control*. It is local in the sense that the data being protected resides on a server somewhere on the network, and a reference monitor mediates access requests to that data based on an ACL and credentials authenticating the accessor, or based on authorization certificates delegated to the accessor. Essentially the accessor has a secure pointer (a capability) pointing to where the data is located on the network. The accessor can retrieve the data over the network using the secure pointer, and access control is enforced locally at the server.

We advocate an alternative approach to access control in a distributed system, *distributed access control*. With this approach, data may not reside centrally on some server. Rather the data itself is copied through the network, transmitted by applications through communication channels. Access control is distributed because it is not enforced at one point in the network where the data is accessed, but may be enforced at myriad points in the network where a copy of the data is accessed. Although there are several possible motivations for this approach, the particular motivation we exploit is *accountability*. The data is stored with certain access rights associated with it. If the data should be leaked, there should be a means for tracking backwards to the source of the leak. If the data is simply provided to the accessor and then forgotten about, there is no way to perform this forensic investigation. With distributed access control, on the other hand, the data is copied to the accessor along with a full specification of the access restrictions. If the accessor provides the data, or some result based on the data, to another party, then the justification for this is recorded for subsequent diagnosis.

**Assumptions.** With enough assumptions about the network environment, distributed access control and accountability are trivial. On the other hand, such assumptions may entail requirements on the operating environment that are so onerous to be for all practical purposes unimplementable. In this work we make the following assumptions:

1. *Intermittent connectivity.* Communications may be sufficiently erratic that the application itself may need to handle the task of transferring data across administrative boundaries. Access decisions may need to be performed at a different network location from where the data was originally stored, *e.g.* data cached on a disconnected machine should remain protected by access control policies.

2. *Trusted hosts.* We assume that there is sufficient control over the local operating environments to enforce access restrictions. In an implementation we could for example assume that the applications run on virtual machines that have not been tampered with.

3. *Insecure network communication.* We allow arbitrary hostile hosts to interfere with traffic between secure hosts. We expose the insecurity of the network to the application and require the application to deal with it. The environment to prevent secure data from being accidentally leaked on insecure communication channels. We do not consider covert communication channels, our motivation is purely to build an audit trail for finding leaks due to improper declassification.

4. *No PKI assumptions.* There is no a single “one-size fits all” approach to PKI. While commercial platforms offer systems such as X509v3 and SPKI, other infrastructures are sure to emerge. We concentrate instead on providing a typed framework in which applications are responsible for safe key distribution. The operations for safe key distribution are a central part of the contribution of this paper.

## 2. Informal Motivation for kDLM

We consider a core language for secure distributed systems equipped with primitives for encryption and digital signing. These primitives are needed to ensure the secrecy of data, including keys, communicated over the network and for ensuring integrity of messages. We provide a type system that expresses precisely the secrecy and integrity properties that these operations are intended to ensure.

There is already a proposal for a type system that enforces secrecy and integrity in the type system [24]. We call this the approach the “Decentralized Label Model” (DLM) and use it as our starting point. Our proposal extends DLM by adding cryptographic operations to the type system. We refer to the extended system as keyed-DLM or simply kDLM.

In DML labelled type in our system has the form  $[T]^{r^w}$ , where  $T$  is an unlabelled type,  $r$  is the label restricting read access to the data (who can consume it), and  $w$  is the label restricting right access to the data (who can produce it). Each label is a set of access control lists (ACLs), each one “owned” or “controlled” by a principal. If the label  $r$  has the form

$$P_1 : \overline{P_1}, \dots, P_m : \overline{P_m}$$

then a principal  $P$  has read access to the data only if

$$P \in \bigcap_{i=1}^m \overline{P_i}.$$

Similarly if the integrity label  $w$  has the above form, then a principal  $P$  has write access to the data only if the above membership condition is true.

One way of viewing the DLM is as a decentralized form of information flow control, where the type system no longer relies on the centralized definition of labels (“high,” “low,” and so on). Decentralized labels also support a controlled form of *declassification*: the owner of an ACL may extend that ACL, so that with the cooperation of all of the owners of the policies in the label, a principal may be added to those allowed to read or write the protected data.

The central idea in our proposal is that add *key names* in the type system. A key name has a “type” that, similarly to an ACL in a label in DLM, identifies an owner principal and a set of reader or writer principals. A key name may be either for encryption or for signing. Each encryption key name has an associated cryptographic key. Here key names are entities in the type system, while cryptographic keys are actually used for public-key encryption. The “type” of the key name constrains the labels in the labelled types of these keys. Both the encryption and decryption keys have the owner principal in the set of possible producers allowed by their integrity keys. The decryption key’s secrecy label cannot allow access to any principals outside those listed in the key name’s ACL.

Key names are purely compile-time entities used for static checking and are stripped along with all type information before a program runs. The “type” of a key name is a *kind*, a form of type for type-level expressions. An encryption key name  $K$  that is generated by the principal  $P$  and is accessible to principals  $P_1 \dots P_m$  has the kind

$$K : \text{EKey}_F(P : P_1 \dots P_m)$$

The flag  $F$  indicates whether this is a *virtual* or *actual* key name. Virtual keys are needed for technical reasons in the static access checks. Actual key names are represented by cryptographic keys written  $a$ . We assume that the set of primitive values is partitioned into at least five sets: the set of channel names (for message passing), the set of encryption keys, the set of decryption keys, the set of signing keys and the set of authentication keys. Furthermore we assume that for every encryption key there is exactly one other key, a decryption key, that is its inverse with respect to cryptographic encryption/decryption. Similarly for every signing key there is exactly one authentication key that is its inverse. We denote the public and private parts of such a key pair by  $a^+$

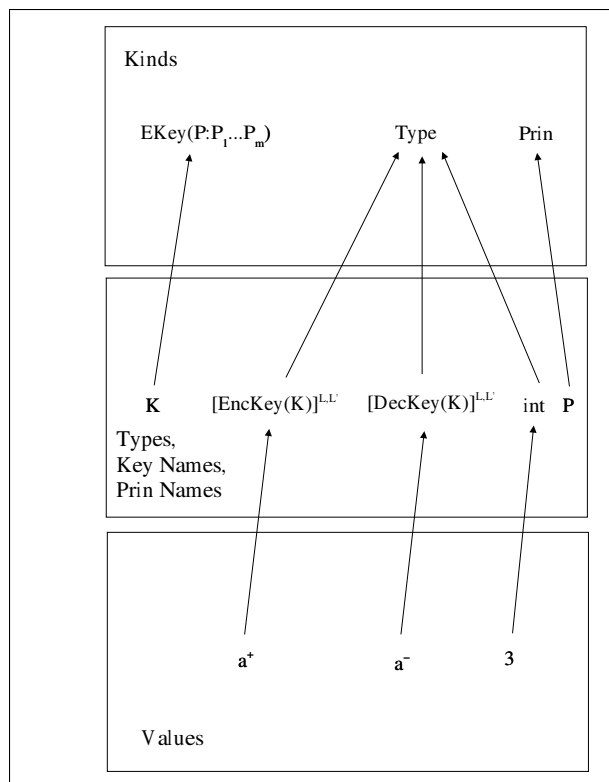


Figure 1. Values, types and kinds

and  $a^-$ , respectively. Then for the encryption key name  $K$  above, we have the typings:

$$a^+ : [\text{EncKey}(K)]^{r_1 w_1} \quad a^- : [\text{DecKey}(K)]^{r_2 w_2}$$

The kind of the key name  $K$  enforces the restriction that the secrecy label  $r_2$  of the private key  $a^-$  cannot allow any principal outside of  $P_1 \dots P_m$  to access the key. A diagram illustrating the relationships between values, types and kinds is provided in Fig. 1. We add integers for comparison. We emphasize that key names  $K$  and principal names  $P$  are purely compile-time entities, and only the actual public and private keys,  $a^+$  and  $a^-$  respectively, exist at run-time. This stratification into values, types and kinds enforces this “phase distinction.”

Purely static access control systems are too constraining for many practical applications. Declassification is a way to circumvent the rigidity of static type checking at the cost of potential security breaches. In our system we choose to impose the following restriction on declassification. The constraints of key names prevent the declassification of a private key to any principal outside those originally allowed in the key name’s kind. However we allow declassification by allowing the “owner” of the key name to issue a *declassification certificate*, of type

$$[K' \text{ reclassifies } K]^{r w}$$

While the private key cannot be accessed by any principal outside the original set  $P_1, \dots, P_m$ , the owner principal  $P$  may issue a declassification certificate that is accessible by some other principal. We describe below how these declassification certificates are then used to declassify data.

Declassification may appear to be an unnecessary complication, but in fact it is the most fundamental operation in our language, and leads to the most profound design decisions in the type system. When we combine labelled types with cryptographic operations, then we must consider the following question<sup>2</sup>:

Should it be possible to declassify encrypted data, making it available to principals that did not have access to it beforehand, without having to decrypt the data, declassify it, and then re-encrypt under another key?

If the answer is no, then we can give our language a simple type system in which the principals mentioned in the decryption key name always are a subset of the ones allowed to encrypt the data. We take the view that this is too restrictive. For example a log server may archive an large number of messages encrypted under different keys depending on the sender, with the server not even necessarily a part of the applications that generated the messages. As a result, the log server will certainly have no part in decrypting messages. Yet if a new principal joins the system and wants to retrieve a particular message, some principal authorized to access that message will have to retrieve it from the server, decrypt it, and forward it re-encrypted under a new key. Without key declassification, the access control system is thus placing unreasonable constraints on the architecture and communication patterns of such distributed systems.

So we allow key declassification, via the mechanism of allowing principals to issue declassification certificates of the form  $K'$  reclassifies  $K$ , signifying that a principal with access to the decryption key for  $K'$  can decrypt values encrypted with the encrypted key for  $K$ . Provided, of course, that the principal has access to the certificate. Furthermore the decrypted data is implicitly declassified at the point of decryption.

Now we must ensure that the combination of encryption, key declassification, and then decryption cannot be used to perform an illegal declassification. Consider the following scenario:

1. There is a value of type  $[T]^{rw}$ . The principal  $P$  is not on the list of principals allowed to access this data (according to the secrecy label  $r$ ).

<sup>2</sup>Since everything that is said about secrecy and encryption can be “dualized” to integrity and digital signing, we only consider the former in our discussion.

2. There is an encryption key  $a^+$  of (unlabelled) type  $\text{EncKey}(K)$ , accessible to all. The kind of  $K$  is compatible with the secrecy label  $r$ , so  $a^+$  can be used to encrypt the data.
3. There is a declassification certificate declassifying the key name  $K$  to some key name  $K'$ . This declassification certificate, and therefore the decryption key for  $K'$ , are accessible to the principal  $P$ .
4. Putting this together,  $P$  can now declassify the data by encrypting it with the encryption key for  $K$ , then decrypting it with the declassification certificate.

So the declassification certificate allows *any* data whose secrecy label is compatible with the encrypting key (i.e., where the value’s secrecy label gives access to any principal that has access to the corresponding decrypting key) to be declassified.

We identify the problem with the DLM as being that it is based on *structural equivalence* of labels and policies, that is, labels  $L$  and  $L'$  are equivalent if they are similar structurally, up to some obvious rearranging of terms. The analogy is with languages such as ML and Modula-3 where type equivalence is structural. The alternative is *name equivalence*. Name equivalence for types means that two types must have the same name; differently named types with the same structure are not considered equivalent. Examples include almost all popular languages, including C and Java.

Declassification of keys creates a connection between key names  $K$  and  $K'$ , delegating decryption rights for  $K$  to  $K'$ . In the DLM, there must then be a connection between the encryption key  $a^+$  for  $K$  and secrecy labels for values that can be encrypted with  $a^+$ , and a connection between the decryption key  $b^-$  for  $K$  and secrecy labels for values that result from decrypting with  $a^-$ . This latter connection is necessary to allow the principal decrypting with the declassification certificate to access the data. The first connection is based on name inclusion (of key names) while the latter two connections are based on (a semantic form of) structural inclusion. It is the latter structural inclusions that allow declassification certificates to be used to perform unintended declassification.

Our solution is to remove some of the structural equivalence and inclusion between labels, but not all of it. To see the need for some form of structural equivalence, consider a hypothetical code fragment:

```
(e) ? foo() : bar()
```

Suppose `foo` and `bar` are procedures provided by different principals, with secrecy labels  $r_f$  and  $r_b$  respectively for their return types. If label inclusion is strictly

name-based, then the result of this conditional must have a secrecy label  $r$  that includes both  $r_f$  and  $r_b$ . If label inclusions are name-based and therefore must be defined, who is responsible for this? Clearly the principal executing the conditional cannot, otherwise they would have the ability to declassify data provided by `foo` and `bar`. If we allow the executing principal to add this inclusion based on some structural containment condition, then we are back where we started, with the DLM. The principals for `foo` and `bar` cannot do this classification, since they know nothing of  $r$  (in general). This motivates keeping labels as *sets*, and the label on the result of the conditional will contain the label  $r_f \cup r_b$ .

We introduce name equivalence, but not name inclusion, on policies. In fact we already have a mechanism for naming policies: key names! Recall that for example for encryption keys, the kind of a key name identifies the principal that generated the key, and also identifies the set of principals that can access the decryption key. Our type system will have two forms of key names, actual and virtual. The former have associated encryption and decryption keys, the latter do not. The latter can be used for static access checking, while the latter are used when data must be encrypted (or of course signed, for integrity) for transmission over the network.

In general a type has the form  $[T]^{r^w}$  where  $r$  is a set of encryption key names and  $w$  a set of signing key names. If a secrecy label  $r$  has the form  $K_1 \dots K_m$  where we have the kindings

$$K_1 : \text{EKey}_F(P_1 : \overline{P_1}) \dots K_m : \text{EKey}_F(P_m : \overline{P_m})$$

then a process executing for principal  $P$  cannot access this value unless  $P \in \bigcap_{i=1}^m P_i$ . This is similar to the read access restriction under the DLM, except that our approach adds the extra level of indirection through the key names and their kinds.

To encrypt a value with this secrecy label, one must have public keys for all of the key names in  $r$  (or  $w$ ), of types

$$[\text{EncKey}(K_1)]^{r_1 w_1}, \dots, [\text{EncKey}(K_m)]^{r_m w_m}.$$

To decrypt the resulting value, a process executing for principal  $P'$  must have access to declassification certificates, declassifying the key names  $\overline{K_m}$  under which the data is encrypted to key names  $\overline{K'_m}$ , so  $P'$  must have declassification certificates of types

$$[K'_1 \text{ reclassifies } K_1]^{r'_1 w'_1} \dots [K'_m \text{ reclassifies } K_m]^{r'_m w'_m}$$

Decrypting the value with these certificates produces a value with secrecy label  $K'_1 \dots K'_m$ , and the process for  $P'$  has read access to this value if  $P' \in \bigcap_{i=1}^m \overline{P'_i}$ , where each  $K'_i$  has kind  $\text{EKey}_{\text{Actual}}(P'_i : \overline{P'_i})$ .

The type system requires that every secret key be secured by the secrecy label on its type, so that it cannot be accidentally leaked to any principals outside of those specified in the kind of its key name. Every secrecy label for a secret key must therefore contain one or more other secrecy key names, that must have already been defined before the secret key's own key name was introduced. There is the danger of an infinite regress here: how does one define the first key? This is part of the motivation for virtual key names: since such key names do not have any associated public-private key pairs, they can be introduced without the need for any other secrecy key names to protect their private keys. The other motivation for virtual key names is convenience and performance: we should not have to endure the expense of generating new public-private key pairs every time we introduce new key names that will only be used for statically checked access control.

Declassification certificates have two applications: The first of which is of course decrypting values that were encrypted under the declassified key. The other application is to declassify a value that has not been encrypted. Our language includes a declassify operation that requires a declassification certificate. When a value is so declassified, the result is annotated with the declassification certificate. This has both practical and theoretical motivation. It supports our original motivation for distributed access control, of building an audit trail; and it is necessary in order to ensure that types are preserved during computation.

Is this three-level system necessary? An alternative approach might be to collapse the three levels into a two-level system of values and types, with a dependent type system (types depend on values; our current system includes dependent kinds but not dependent types). However the problem with this approach is that key names and principal names become values, and now we have to deal unnecessarily with what the access restrictions should be for them. The three-level system avoids this complication. If we wish to extend the system with dynamic principal and key names (the former is available in the JIF implementation [24] but not its meta-theory), this can be done by adding special tag types indexed by the principal name or key name. Indeed public-private key pairs are essentially values for such tag types for key names. Virtual key names avoid the unnecessary introduction of such tags where they are not needed. As we have seen, this is crucial to breaking the cycle in introducing keys with access protected by other keys. It is not clear how this cycle would be broken in a two-level dependent type system.

$A \in \text{Arity, Kind}$	$::=$	Prin	Principal
		EKey <sub>F</sub> ( $P : \bar{P}$ )	Encryption key
		SKey <sub>F</sub> ( $P : \bar{P}$ )	Signing key
		Type	Type
$F \in \text{Key Flag}$	$::=$	Actual	Actual key
		Virtual	Virtual key
$K, P, T \in \text{Key, Prin, Type}$	$::=$	$k, p, t$	Variable, name
		DecKey( $K$ )	Decryption key type
		EncKey( $K$ )	Encryption key type
		AuthKey( $K$ )	Authentication key type
		SignKey( $K$ )	Signing key type
		$K_1$ reclassifies $K_2$	Reclassification cert
		$\mathcal{E}\{LT\}$	Actually encrypted
		$\mathcal{S}\{LT\}$	Actually signed
		Chan( $LT$ )	Channel type
		$\langle t : A \rangle LT$	Package
$L \in \text{Label}$	$::=$	$\{K_1, \dots, K_m\}$	
$LT \in \text{Labelled type}$	$::=$	$[T]^{L_1 L_2}$	

**Figure 2. Syntax of Kinds and Types**

### 3. Type System: Formal Description

We now provide a more detailed description of the type system. The syntax of types and kinds is provided in Fig. 2. The system of arities or kinds organizes the well-formed types, principal names and key names. Besides the kind of types (Type) and principal names (Prin), there are two kinds for key names: encryption (secrecy) key names EKey<sub>F</sub>( $P : \bar{P}$ ) and signing (integrity) key names SKey<sub>F</sub>( $P : \bar{P}$ ). This kind identifies the creator of the key ( $P$ ) and the principals who may have access to the private part of the public-private key pair for that key ( $\{\bar{P}\}$ ). Every key name kind also has a flag indicating if it is a virtual key (static checking only) or an actual key (dynamic checking may also be performed, using cryptographic operations).

The types then consist of key types, indexed by key names, the types of reclassification certificates, and the types of actually (i.e., cryptographically) encrypted or signed values,  $\mathcal{E}\{LT\}$  and  $\mathcal{S}\{LT\}$ . The types are rounded out by the other (non-cryptographic) types of the base language, in this case channel types and package types. The latter are also known as *existential types* [21]; they are practically necessary for transmitting prin-

cipal and key information when principal and key names are static second-class entities, at the same level as types. For example to send an encryption key to a process that requires an encryption key created by the principal *root*, no matter the underlying key name, we send it a value of type<sup>3</sup>

$$\langle k : \text{EKey}_F(\text{root} : (\text{root})) \rangle [\text{EncKey}(k)]^{L_1 L_2}.$$

Kinds and types depend on each other, resulting in a dependent kind system. The formation rules for kinds and types must therefore be formulated in a mutually recursive fashion, and these formation rules intertwined with the rules for environment formation in the type system. An environment is a sequence of pairs, binding (variables or names) to (kinds or types). We have two forms of environments:

$$\begin{aligned} TE \in \text{Type Env} & ::= \varepsilon \mid (t : A) \mid TE_1, TE_2 \\ VE \in \text{Value Env} & ::= \varepsilon \mid (x : LT) \mid (a : LT) \mid VE_1, VE_2 \end{aligned}$$

<sup>3</sup>Even this trivial example suggests a need for some notion of parameterized types (generics) for parameterizing values by key names and principal names, and even sets of principal names because of key name kinds. Presumably some notion of bounded type quantification would also be useful. However we do not consider this avenue any further in this paper.

$e \in \text{Expression}$	$::=$	$w, x, y, z$	Variable
		$a, b, c, n$	Name (Channel, Key)
		$\text{newKey}\langle k : A \rangle\{e\}$	New virtual key
		$\text{newKey}\langle k : A \rangle(a^+ : LT_1, a^- : LT_2)\{e\}$	New actual key
		$\text{reclassifyCert}_{K_1, K_2}()$	Reclassify cert (virtual key)
		$\text{reclassifyCert}_{K_1, K_2}(e)$	Reclassify cert (actual key)
		$\text{chain}_{K_1, K_2, K_3}(e_1, e_2)$	Chain certs
		$\text{reclassify}_{K_1, K_2}(e_1, e_2)$	Reclassify value
		$\text{encrypt}_{\bar{K}}(\bar{e}_1, e_2)$	Encrypt
		$\text{decrypt}_{\bar{K}_1, \bar{K}_2}(\bar{e}_1, e_2)$	Decrypt
		$\text{sign}_{\bar{K}_1, \bar{K}_2}(\bar{e}_1, e_2)$	Sign
		$\text{auth}_{\bar{K}}(\bar{e}_1, e_2)$	Authenticate
		$\text{new}(n : LT)\{e\}$	New channel
		$\text{fork}\{e\}$	New process
		$\text{send}(e_1, e_2)$	Message send
		$\text{receive}(a)$	Message receive
		$\text{pack}_{\langle t : A \rangle LT}(K, e)$	Package (key, data)
		$\text{unpack } e_1 \text{ to } \langle k : A \rangle(x : LT)\{e_2\}$	Open package
$R \in \text{Process}$	$::=$	$e$	Sequential process
		$\text{new}(k : A)\{R\}$	Key name
		$\text{new}(a : LT)\{R\}$	Channel, key
		$(R_1 \mid R_2)$	Parallel composition

**Figure 3. Syntax of Values, Expressions and Processes**

The sequence concatenation operation ( $\cdot$ ) is assumed to be associative with  $\epsilon$  (empty sequence) as its unit.

The type environment  $TE$  is used in the definition of a metafunction that, given a label (set of key names), computes the set of principals that is defined by that label in that type environment:

$$\begin{aligned}
 PRINS_{TE}(\{\}) &= \{\} \\
 PRINS_{TE}(L_1 \cup L_2) &= PRINS_{TE}(L_1) \cap PRINS_{TE}(L_2) \\
 PRINS_{TE}(\{k\}) &= \{\bar{P}\} \text{ if } \exists TE_1, TE_2. \\
 &\quad TE = (TE_1, k : EKey_F(P : \bar{P}), TE_2) \\
 PRINS_{TE}(\{k\}) &= \{\bar{P}\} \text{ if } \exists TE_1, TE_2. \\
 &\quad TE = (TE_1, k : SKey_F(P : \bar{P}), TE_2)
 \end{aligned}$$

We make use of the fact that the only forms of key names in our system are atomic names. If some extension of the system were to consider more structured key names,

then this metafunction would have to be defined mutually recursively (though the recursion would be well-founded) with the type formation rules.

The formation rules for (labelled) types check, e.g., that in an encryption key  $\text{EncKey}(K)$ , the argument  $K$  denotes a key name with an encryption key name kind. These rules also check label constraints that are placed on public and private key types by the kinds of the corresponding key names. For example for encryption and decryption keys we have these type formation rules:

$$\frac{TE \vdash L_1, L_2 \text{ label} \quad TE \vdash K : EKey_F(P : \bar{P}) \quad P \in PRINS_{TE}(L_2) \quad PRINS_{TE}(L_1) \subseteq \{\bar{P}\}}{TE \vdash [\text{DecKey}(K)]^{L_1 L_2} : \text{Type}} \quad (\text{LCON DECKEY})$$

$$\frac{TE \vdash L_1, L_2 \text{ label} \quad TE \vdash K : \text{EKey}_F(P : \bar{P}) \quad \frac{P \in \text{PRINS}_{TE}(L_2)}{TE \vdash [\text{EncKey}(K)]^{L_1 L_2} : \text{Type}}}{(L\text{CON ENCKEY})}$$

The first rule checks that the decryption key has the owning principal  $P$  in its integrity label, and that the set of principals allowed access to the decryption key by its secrecy label is contained in the principals allowed by the corresponding key name's kind. The second rule checks that the encryption key has the owning principal  $P$  in its integrity label. Analogous (dual) rules hold for signing and authentication keys.

The kinds of key names also place a constraint on the labels of declassification certificates. For example for declassification certificates for encryption keys, we have the following:

$$\frac{TE \vdash L_1, L_2 \text{ label} \quad TE \vdash K_1 : \text{EKey}_{F_1}(P_1 : \bar{P}_1) \quad TE \vdash K_2 : \text{EKey}_{F_2}(P_2 : \bar{P}_2) \quad \frac{P_2 \in \text{PRINS}_{TE}(L_2) \quad \{P_1\} \subseteq \text{PRINS}_{TE}(L_1)}{TE \vdash [K_1 \text{ reclassifies } K_2]^{L_1 L_2} : \text{Type}}}{(L\text{CON RECL DECR})}$$

This rule requires that the integrity label for the certificate include the principal name owning the key name  $K_2$  that is declassified; only this principal could have declassified the key name. The secrecy label for the certificate must define a set of principals that is contained in the set of principals that are allowed to access the private decryption key for  $K_1$ , the key name to which  $K_2$  is declassified. Because of this latter condition, only principals that are already allowed to decrypt values encrypted with the public key for  $K_1$  (using the latter's private key) will also be allowed to decrypt values encrypted with the public key for  $K_2$  (using the declassification certificate). So reclassification certificates allow keys to be shared between principals, even in unanticipated ways, while certificates are the basis for building an audit trail of such key sharing.

The syntax of values, expressions and processes is provided in Fig. 3, where values  $V$  denote a subset of the expressions  $e$  where no further evaluation is necessary. For type preservation purposes, we also introduce a notion of *annotated values*  $\langle\langle V \rangle\rangle_{P,C}$ . The motivation is that for example an expression may have secrecy label  $\{K\}$  but then be dynamically declassified to have secrecy label  $\{K'\}$ . With no relationship between the two key names, we cannot argue that the result of declassification has the same type as the original value, and therefore evaluation does not appear to preserve types. In our language, declassification is only possible when it is authorized by a declassification certificate. We restore the desired invariant of type preservation under evaluation by assuming that values are annotated by chains of

such certificates. A declassification operation then adds its certificate chain to the annotation of the value it is declassifying. Annotations on values are not purely a fiction for preserving type preservation; they are consistent with the original motivation for distributed access control, *viz* building a fine-grain audit trail of accesses and declassifications.

The type system for expressions in general uses judgements of the form

$$TE; VE \vdash e \in^P [T]^{L_1 L_2}$$

to check that the expression (code)  $e$  is well-formed with annotated type  $[T]^{L_1 L_2}$ , under the assumption that it will be evaluated (executed) under the authority of the principal  $P$ , in the corresponding type and value environments  $TE$  and  $VE$ , respectively. We modularize the type rules for expressions with three judgement forms:  $TE; VE \vdash^I e \in^P [T]^{L_1 L_2}$ ,  $TE; VE \vdash^O e \in^P [T]^{L_1 L_2}$  and  $TE; VE \vdash e \in^P [T]^{L_1 L_2}$ . The first of these denotes a type judgement where we have checked that the principal  $P$  is allowed to access the result of the expression  $e$ , according to the secrecy label  $L_1$ . Typically such a judgement will be used for the premises for an expression type rule. The second judgement form is used as a conclusion for an expression type rule, denoting that the integrity condition on the conclusion has not yet been checked, i.e., that the principal  $P$  is allowed to compute the result of the expression  $e$ , according to the integrity label  $L_2$ . So the general form of a type derivation is

$$\frac{\begin{array}{c} \vdots \\ TE; VE \vdash e \in^P [T]^{L_1 L_2} \\ \hline TE; VE \vdash^I e \in^P [T]^{L_1 L_2} \end{array} \quad \dots \quad \frac{\begin{array}{c} \vdots \\ TE; VE \vdash e \in^P [T]^{L_1 L_2} \\ \hline TE; VE \vdash^I e \in^P [T]^{L_1 L_2} \end{array}}{TE; VE \vdash^O e \in^P [T]^{L_1 L_2}}}{TE; VE \vdash e \in^P [T]^{L_1 L_2}}$$

In some cases the input and output checking may be circumvented, for example for variables, or for example where we sign a value with a private key and do not expect the access restrictions on the key to propagate to the resulting ciphertext (since we are ignoring covert channels and only considering access control, as far as we are concerned the key cannot be recovered from the ciphertext).

There are two key generation operations, one for virtual keys and one for actual keys. The type rule for virtual key generation is as follows:

$$\frac{A = \text{EKey}_{\text{Virtual}}(P : \bar{P}) \quad TE \vdash A \text{ kind} \quad (TE, k : A); VE \vdash e \in^P [T]^{LL'} \quad k \notin \text{tv}([T]^{L_1 L_2})}{TE; VE \vdash^O \text{newKey}(k : A)\{e\} \in^P [T]^{LL'}} (VAL \text{ NEWKEY VIRT})$$



$TE \vdash \mathbf{tenv}$	Well-formed type environment
$TE \vdash VE \mathbf{venv}$	Well-formed value environment
$TE \vdash A \mathbf{kind}$	Well-formed kind
$TE \vdash LT : A$	Well-formed type
$TE \vdash L \mathbf{label}$	Well-formed label
$TE; VE \vdash^I e \in^P [T]^{L_1 L_2}$	Type judgement (inputs checked)
$TE; VE \vdash^O e \in^P [T]^{L_1 L_2}$	Type judgement (outputs to be checked)
$TE; VE \vdash e \in^P [T]^{L_1 L_2}$	Type judgement (outputs checked)

**Figure 4. Judgements for Type System**

This rule introduces a new encryption key name  $k$ , owned by the principal  $P$  that will evaluate this expression. This key name is local to the block  $e$ , but can escape this block if it is bundled up in a package. When used in a secrecy label, such a key name restricts accesses to some subset of the principals  $\{\bar{P}\}$ . The key is only useful for this form of static access checking.

The type rule for actual key generation is as follows:

$$\begin{array}{c}
A = \text{EKey}_{\text{Actual}}(P : \bar{P}) \quad TE \vdash A \mathbf{kind} \\
LT_1 = [\text{EncKey}(k)]^{L_1 L'_1} \quad (TE, k : A) \vdash LT_1 : \text{Type} \\
LT_2 = [\text{DecKey}(k)]^{L_2 L'_2} \quad (TE, k : A) \vdash LT_2 : \text{Type} \\
(TE, k : A); (VE, a^+ : LT_1, a^- : LT_2) \vdash e \in^P [T]^{L L'} \\
\quad k \notin \text{tv}([T]^{L_1 L_2}) \\
\hline
TE; VE \vdash^O \text{newKey}(k : A)(a^+ : LT_1, a^- : LT_2)\{e\} \in^P [T]^{L L'} \\
\text{(VAL NEWKEY ACT)}
\end{array}$$

Actual key generation introduces not only a new key name, but also a public-private key pair for that key name. This key pair is denoted by the pair of dual names  $(a^+, a^-)$ . The encryption key  $a^+$  has an encryption key type indexed by the new key name, and the compatibility restrictions between the kind  $A$  of this key name and the labels of the key type are enforced before the key and its type are added to the value environment. Similarly for the decryption key  $a^-$ .

There are two forms of declassification for keys: declassifying a virtual key and declassifying an actual key. For the former we have the type rule:

$$\begin{array}{c}
TE \vdash K_1 : \text{EKey}_F(P_1 : \bar{P}_1) \quad TE \vdash K_2 : \text{EKey}_{\text{Virtual}}(P_2 : \bar{P}_2) \\
TE \vdash [K_1 \text{ reclassifies } K_2]^{L L'} : \text{Type} \\
TE \vdash VE \mathbf{venv} \quad TE \vdash P : \text{Prin} \\
\hline
TE; VE \vdash^O \text{reclassifyCert}_{K_1, K_2}() \in^P [K_1 \text{ reclassifies } K_2]^{L L'} \\
\text{(VAL RECL VIRT)}
\end{array}$$

There is a little subtlety here. Suppose the principal  $P$  producing the certificate is not the same as the princi-

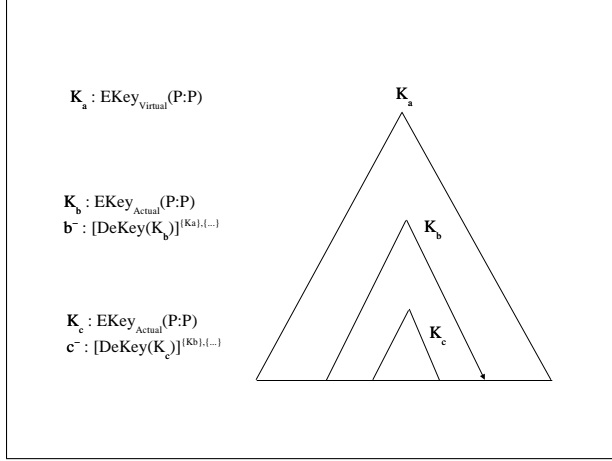
pal  $P_2$  that owns the key being declassified? In that case we will have  $P$  included as one of the possible principals that created the certificate (after we check the integrity condition on the conclusion). Since this means that the certificate may have come from a principal other than the owner of the key being declassified, the declassification certificate will be effectively useless. So an attacker can always “spooft” a declassification certificate, but they must subvert the integrity checking to have this certificate taken seriously.

When declassifying an actual key name  $K_2$ , we must provide the associated decryption key. This is so that the principal receiving the declassification certificate can subsequently use it to declassify data encrypted under the encryption key for  $K_2$ . In this case we must also check that the decryption key could only have come from the principal that created the key name ( $P_2$ ).

$$\begin{array}{c}
TE \vdash K_1 : \text{EKey}_F(P_1 : \bar{P}_1) \quad TE \vdash K_2 : \text{EKey}_{\text{Actual}}(P_2 : \bar{P}_2) \\
TE \vdash [K_1 \text{ reclassifies } K_2]^{L L'} : \text{Type} \\
TE; VE \vdash^I e \in^P [\text{DecKey}(K_2)]^{L L'} \quad \text{PRINS}_{TE}(L') = \{P_2\} \\
\hline
TE; VE \vdash^O \text{reclassifyCert}_{K_1, K_2}(e) \in^P [K_1 \text{ reclassifies } K_2]^{L L'} \\
\text{(VAL RECL ACT)}
\end{array}$$

Now there is a subtlety with the type rule for declassifying actual keys. The certificate that results from declassification has the same secrecy label  $L$  as the original decryption key, therefore the declassification certificate can only be accessed by the same set of principals  $\{\bar{P}_2\}$  that already have access to the private key. Therefore this declassification certificate must itself be declassified.

We illustrate this point with an example, shown in Fig. 5. A principal  $P$  has created three key names,  $K_a$ ,  $K_b$  and  $K_c$ . The private key for  $K_c$  is protected by  $K_b$ , while the private key for  $K_b$  is protected by  $K_a$ . The hierarchy stops here because  $K_a$  is a virtual key. Now the principal  $P$  wants to declassify  $K_c$  to the key  $K'_c$  of another principal  $P'$ . So  $P$  generates declassification cer-



**Figure 5. Key hierarchies and declassification**

certificates of types:

$$\begin{aligned}
 e_1 &\in [K'_c \text{ reclassifies } K_c] \{K_b\} \{\dots\} \\
 e_2 &\in [K'_b \text{ reclassifies } K_b] \{K_a\} \{\dots\} \\
 e_3 &\in [K'_a \text{ reclassifies } K_a] \{K'_c\} \{\dots\}
 \end{aligned}$$

The last certificate is allowed by the rule for declassifying virtual key names, above, that places no restrictions on the secrecy label for the certificate. Assuming that  $K'_a$ ,  $K'_b$  and  $K'_c$  are keys of  $P'$  analogous to the keys  $K_a$ ,  $K_b$  and  $K_c$  of  $P$ , then  $P'$  can use the last certificate to declassify the second certificate, and then use the second certificate to declassify the first certificate.

An obvious operation for declassification certificates is certificate chaining. This is provided by the  $\text{chain}_{K_1, K_2, K_3}(e_1, e_2)$  construct. In all of this, there is no problem with declassifying virtual key names to actual key names or vice versa. The only place where these certificates have computational import is in decryption, to which we turn below.

There are two applications of declassification certificates, as noted. The first application is in declassifying data, analogous to declassification in the DLM model:

$$\begin{array}{c}
 TE; VE \vdash^J e_1 \in^P [K' \text{ reclassifies } K]^{L_1 L'_1} \\
 TE; VE \vdash^J e_2 \in^P [T]^{L_2 \uplus \{K\}} L'_2 \\
 TE \vdash K : EKey_F(P_0 : \bar{P}) \quad PRINS_{TE}(L'_1) = \{P_0\} \\
 TE \vdash [T]^{L_2 \uplus \{K'\}} L'_2 : \text{Type} \\
 \hline
 TE; VE \vdash^O \text{reclassify}_{K', K}(e_1, e_2) \in^P [T]^{L_2 \uplus \{K'\}} L'_2 \\
 \text{(VAL RECLASSIFY)}
 \end{array}$$

In this rule the operation  $\uplus$  denotes disjoint union. So declassification replaces the key name  $K$  in the secrecy label with the key name  $K'$ . The access restrictions

on the declassification certificate do not propagate to the declassified value, since there is no construct in our language for extracting certificates from annotated values. We assume that all code within a single process is typed according to the type system described here, and that all communication over the network is encrypted in order to avoid leaks due to low-level attackers that do not respect the type system and can intercept messages. We return to this point in the conclusions in Sect. 6.

The second application of declassification certificates is in declassifying data that has already been encrypted. In fact we fold this operation into the decryption operation itself. Rather than taking a sequence of decryption keys, corresponding to the keys under which the data is encrypted, this operation takes a sequence of declassification certificates. If a key name is not being declassified as part of decryption, then the corresponding certificate can simply declassify the key name to itself. The rules for encryption and decryption are as follows:

$$\begin{array}{c}
 TE; VE \vdash^J \bar{e} \in^P [\overline{\text{EncKey}(K)}]^{L \bar{L}} \\
 TE \vdash \bar{K} : EKey_{\text{Actual}}(P : \bar{P}') \quad \overline{PRINS_{TE}(L')} = \{\bar{P}\} \\
 TE; VE \vdash^J e_0 \in^P [T]^{L_0 L'_0} \quad L_0 = \{\bar{K}\} \quad TE \vdash L''_0 \text{ label} \\
 \hline
 TE; VE \vdash^O \text{encrypt}_{\bar{K}}(\bar{e}, e_0) \in^P [\mathcal{E}\{T\}]^{L_0 L'_0} \\
 \text{(VAL ENCRYPT)}
 \end{array}$$

$$\begin{array}{c}
 TE; VE \vdash^J \bar{e} \in^P [\bar{K}' \text{ reclassifies } \bar{K}]^{L \bar{L}} \\
 TE \vdash \bar{K} : EKey_{\text{Actual}}(P : \bar{P}') \quad \overline{PRINS_{TE}(L')} = \{\bar{P}\} \\
 TE; VE \vdash^J e_0 \in^P [\mathcal{E}\{T\}]^{L_0 L'_0} \quad L'_0 = \{\bar{K}'\} \\
 \hline
 TE; VE \vdash^O \text{decrypt}_{\bar{K}, \bar{K}'}(\bar{e}, e_0) \in^P [T]^{L_0 \cup L'_0 L'_0} \\
 \text{(VAL DECRYPT)}
 \end{array}$$

The encryption rule checks that the encryption keys that are provided do in fact come from the principals that own the corresponding key names. These encryption keys must of course match the key names in the secrecy label for the data being encrypted. The secrecy label on the result is arbitrary (it might be empty). In a real sense encryption is a way of circumventing the type system, but in a “safe” fashion since the resulting value cannot be reintroduced without a dynamic check. This check is by the decryption operation. It again checks that the declassification certificates do indeed come from the principals that own the key names for the decryption keys being declassified by the certificates. The decryption operation asserts that the encrypted value was encrypted with the keys  $\{\bar{K}\}$ , but this can only be ascertained by a run-time check. If we implement encryption and decryption as iterated single-key operations, then we must assume some canonical ordering of the keys, so a sequence of keys is always applied in that canonical order to encrypt data under several keys. Additional assumptions may be necessary. The result of decryption has a secrecy label that includes the keys  $\bar{K}'$  plus also any restrictions that prop-

agate from the expression that produces the encrypted data.

There are analogous but essentially dual rules for integrity key names. We omit the obvious details in order to conserve space.

## 4. Operational Semantics

In this section we consider an operational semantics for our language. This semantics (perhaps surprisingly) does run-time access checking. We verify that types are preserved under evaluation. We use this result to verify progress: a process' evaluation can only block on a failure to decrypt, or a failure to authenticate a signature (or, when we consider message-passing operations, on a blocking receive). This verifies that the run-time access checking is unnecessary: the type system successfully prevents any access violations.

To describe and reason about the semantics, we extend the syntax of expressions as described in Fig. 6. The class of values  $V$  describes the result of evaluation, where no further evaluation steps are necessary. The language contains as values names (channels and keys), certificates, encrypted and signed data, and packages. We then annotate these values with a principal and a certificate chain. The principal and certificate chain annotations are used to ensure type preservation during evaluation. The value may have been computed by a principal different from the current principal, with data that are not accessible to the current principal. The certificate chain is extended every time the value is declassified, to justify giving it the declassified type it has even though the original value does not. The certificate is a sequence of (principal, declassification certificate) pairs, that is extended incrementally every time the annotated value is declassified.

The syntax of expressions is now extended with such annotated values  $v$ , and annotating expressions  $\langle e \rangle$ , and also configurations of the form  $P[e]$ . The latter denotes an expression (thread)  $e$  that is executing for the principal  $P$ . We allow such configurations to be nested; this happens when code executing for one principal does a method or procedure call to code supplied by another principal. Some of the computation rules are provided in Fig. 7.

We have omitted the type rules for processes, but they are completely standard and straightforward.

**Theorem 1 (Type Preservation)** *The following propositions are true:*

1. Suppose  $TE; VE \vdash R_1 \mathbf{proc}$  and  $R_1 \xrightarrow{TE; VE} R_2$ , then  $TE; VE \vdash R_2 \mathbf{proc}$ .

2. Suppose  $TE; VE \vdash e \in^P [T]^{LL'}$  and  $e \xrightarrow{TE; VE; P} e'$ , then  $TE; VE \vdash e' \in^P [T]^{LL'}$ .

Define an *evaluation context*  $\mathbf{C}[\cdot]$  in the normal fashion. Define a *process context* by:

$$\begin{aligned} \mathbf{P}[\cdot] ::= & \mathbf{C}[\cdot] \mid (\mathbf{P}[\cdot] \mid R) \mid (R \mid \mathbf{P}[\cdot]) \\ & \mid \mathbf{new}(k : A)\{\mathbf{P}[\cdot]\} \mid \mathbf{new}(a : LT)\{\mathbf{P}[\cdot]\} \end{aligned}$$

This is used to identify places in a process where there should be a redex, and if not then computation is blocked. The proof of the following uses the fact that the dynamic access checks are performed as type checks, and also of course using the previous theorem.

**Theorem 2 (Progress)** *The following proposition are true:*

1. Suppose  $TE; VE \vdash R \mathbf{proc}$  and  $R = \mathbf{P}[e]$  where  $e$  is neither a value nor a redex. Then  $e$  is either an application of decryption, or an application of authentication.
2. Suppose  $TE; VE \vdash e \in^P [T]^{LL'}$  and  $e = \mathbf{P}[e']$  where  $e'$  is neither a value nor a redex. Then  $e'$  is either an application of decryption, or an application of authentication.

## 5. Related Work

The motivation for this work has been the need for proper programming abstractions for developing applications that must manage some or all of the task of securing their communication in a network environment. Almost all of the work in the area of wide-area languages has focused on security, for example, providing abstractions of secure channels [6, 5, 4], controlling key distribution [14, 13], reasoning about security protocols [7, 1], tracking untrustworthy hosts in the system [20, 29], etc.

Abadi [1] considers a type system for ensuring that secrecy is preserved in security protocols implemented in that type system. For securing communication over untrusted networks he includes a “universal” type  $\mathbf{Un}$  inhabited by encrypted values. His type system prevents “secrets” from being leaked to untrusted parties, but allows encrypted values to be divulged. In an analogous way, encrypted and signed values in our type system provide a way to temporarily subvert the access controls in the type system, with the secrecy and integrity properties enforced by labels reasserted when the ciphertext is decrypted/authenticated.

Gordon and Jeffrey [18, 19] have developed a type-based approach to verifying authentication protocols.

$V \in \text{Value}$	$::=$	$a, b, c, n$	Name (Channel, Key)
		$\text{reclassifyCert}_{K_1, K_2}()$	Reclassify cert (virtual key)
		$\text{reclassifyCert}_{K_1, K_2}(v)$	Reclassify cert (actual key)
		$\text{encrypt}_{\bar{K}}(v_1, v_2)$	Encrypted data
		$\text{sign}_{\bar{K}_1; \bar{K}_2}(v_1, v_2)$	Signed data
		$\text{pack}_{(t:A)LT}(K, v)$	Package (key, data)
$v \in \text{Annotated Value}$	$::=$	$\langle\langle V \rangle\rangle_{P,C}$	
$C \in \text{Certificate chain}$	$::=$	$((P_1, v_1), \dots, (P_m, v_m))$	
$e \in \text{Expression}$	$::=$	$\langle e \rangle$	Add annotation
		$v$	Annotated value
		$P[e]$	Configuration

**Figure 6. Additional Syntax for Semantics**

They develop a dependent type system, but in their case the intention is to maintain a connection in the type system between messages and nonces, while nonce types and an effect type system use correspondence assertions to ensure the freshness of communications.

Abadi and Blanchet [2, 10] pursue an approach to analyzing security protocols, initially for secrecy properties but later generalizing it to integrity properties. They have a notion of a type of “secret,” and a type system that ensures that secret items are never put on channels shared with untrusted parties. They can translate types in their system into logic programs that can then be used to check protocols for correctness. The emphasis of this work is somewhat different, since Bruno and Blanchet work in a more “black and white” environment where there are trusted parties and untrusted parties. In contrast our interest is in a more refined type system where we allow certain parties to access certain data. Nevertheless for obvious reasons of space and exposition, we have for now avoided the issue of prevented data from leaking on untrusted channels. There may be scope for synergy with the work of Abadi and Blanchet in the future.

All of these works are focused on verifying secrecy and integrity properties of security protocols. As such the type systems that they use are far more sophisticated than the average programmer will use, while at the same time they give very strong guarantees of secrecy and integrity. The focus of our work is not protocol verification, but building accountable systems: engineering a system where accesses can be logged, but doing it in such a way that the performance of the system is not killed by the demands of credentials checking. So for example we make no attempt to cope with replay attacks. Our type system is compatible with a type system

that has been designed and implemented for Java [24]. It remains to be seen whether either it or our type system are useful in practice.

Sumii and Pierce [32] describe the cryptographic lambda calculus, an extension of the typed lambda calculus with symmetric cryptographic operations along the lines of the spi-calculus. They then use reasoning techniques for the polymorphic lambda calculus (in particular, relational parametricity) to verify properties of programs implemented using cryptographic techniques.

Other work on security in programming languages has focused on ensuring safety properties of untrusted code [26, 25, 22] and preventing unwanted security flows in programs [15, 23, 33, 28]. Sabelfeld and Myers [30] provide an excellent overview of work in language-based information-flow security. Our security concerns have largely been with access control, but there is a clear and obvious relationship to the decentralized label model of JIF. This relationship exists because we initially based our type system on decentralized labels, as explained in Sect. 2.

Pottier and Conchon [28] have developed an interesting approach to encoding information flows in the lambda-calculus, allowing non-interference to be checked in an operational manner, and Pottier [27] has extended this to the pi-calculus. Because of the operational nature of their work, it appears plausible that it could be useful in proving some form of non-interference for our language. But in the presence of declassification, it is a well-known problem to define what safety guarantees are provided by information flow control [30, 34]. In any case, this has not been the focus of our concern, as indicated by Sect. 1.

$$\begin{array}{c}
\frac{TE; VE \vdash v \in^P [T]^{LL'}}{\langle V \rangle \xrightarrow{TE; VE; P} \langle\langle V \rangle\rangle_{P, ()}} \quad (\text{RED PRIN ANN}) \\
\\
\frac{TE; VE \vdash^I v_1 \in^P [T]^{L_1 L_1'} \quad TE; VE \vdash^I v_2 \in^P [T]^{L_3 L_3'}}{\text{reclassify}_{K_1, K_2}(v_1, v_2) \xrightarrow{TE; VE; P} \text{CEXTEND}(v_2, ((P, v_1)))} \quad (\text{RED PRIN RECL}) \\
\\
\frac{v_1 = \langle\langle \text{chain}_{K_1, K_0, K_2}(v_2, v_3) \rangle\rangle_{P_1, C_1} \quad TE; VE \vdash^I v_1 \in^P [T]^{L_0 L_0'} \quad v'_2 = \text{CEXTEND}(v_2, C_1) \quad v'_3 = \text{CEXTEND}(v_3, C_1)}{\text{decrypt}_{\overline{K_1}, \overline{K_1}, \overline{K_2}, \overline{K_2}}((\overline{v_1}, v_1), v) \xrightarrow{TE; VE; P} \text{reclassify}_{K_1, K_0}(v'_2, \text{decrypt}_{\overline{K_1}, \overline{K_0}, \overline{K_2}, \overline{K_2}}((\overline{v_1}, v'_3), v))} \quad (\text{RED PRIN DECCH}) \\
\\
\frac{v_1 = \langle\langle \text{reclassifyCert}_{K_1, K_2}(v_0) \rangle\rangle_{P', C} \quad TE; VE \vdash^I v_1 \in^P [T]^{L_2 L_2'} \quad v_0 = \text{CEXTEND}(v'_0, C)}{\text{decrypt}_{\overline{K_1}, \overline{K_1}, \overline{K_2}, \overline{K_2}}((\overline{v_1}, v_1), v) \xrightarrow{TE; VE; P} \text{reclassify}_{K_1, K_0}(v_1, \text{decrypt}_{\overline{K_1}, \overline{K_0}, \overline{K_2}, \overline{K_2}}((\overline{v_1}, v'_0), v))} \quad (\text{RED PRIN DECRT}) \\
\\
\frac{\overline{v} = \langle\langle a^- \rangle\rangle_{P_1, C_1} \quad v' = \langle\langle \text{encrypt}_{\overline{K}}(\langle\langle a^+ \rangle\rangle_{P_2, C_2}, v'') \rangle\rangle_{P', C} \quad TE; VE \vdash^I \overline{v} \in^P [\overline{T}]^{\overline{L_1} \overline{L_1}'} \quad TE; VE \vdash^I v' \in^P [T]^{L'' L'''}}{\text{decrypt}_{\overline{K_1}, \overline{K_2}}(\overline{v}, v') \xrightarrow{TE; VE; P} \text{CEXTEND}(v'', C)} \quad (\text{RED PRIN DECACT})
\end{array}$$

Figure 7. Selected Computation Rules

## 6. Conclusions

This paper has combined two ideas:

1. First, the notion of type-based cryptographic operations [16], with the intention that some of the secrecy and integrity properties of those operations can be checked statically. This can sometimes avoid the expense of cryptographic operations at run-time. At the very least, it provides a way for specifying the security guarantees of a channel provided from lower layers in the protocol stack to upper layers.
2. Second, the notion of decentralized labels, that combine access control and some form of information flow control [23, 24]. We are strictly interested in access control and so emphasize this aspect of decentralized labels.

The payoff from this combination is that we can now see a way to move network security out of the TCB and into the application, with the type system making sure that the application does not violate the label constraints (modulo declassification). For example a distributed implementation of JIF [35] puts all network security in the TCB and makes the network transparent to the program. It appears implausible that this will scale over the network, or in mission-critical or fault-tolerant systems.

There are obviously many issues to explore further. Our focus due to space has been on parties interacting and preventing data from “flowing” (directly) to unde-

sirable parties. What about untyped attackers on the network? For now we assume that all network communication is encrypted and signed, but there should be ways to interact safely with untrusted parties, declassifying data from the world and reading data from the world. We have some ideas on how to proceed with this.

Cryptographic types [16] were based on the idea that there would be subnets (or just OS buffers) where it would be safe to assume that data was secure without being encrypted and signed. That work did not consider how to delineate where on the network it would be secure to send data intended for certain parties, and where it would be necessary to build trusted channels from untrusted channels. The current work still does not address this issue. We expect to address it in the near future, but back in the context of cryptographic types where the language and the problems are simpler.

It would be good to develop more notions of correctness for our work. For example it appears plausible that some form of robust non-interference could be demonstrated for this language [34], based on simulation relations. However the language that we work with is rich and we do not expect such a result to be easy.

Another direction to consider is accountability: formalizing it, perhaps based on causality types [31], and verifying that our certificate chains for annotated values are in some sense a good measure of accountability. These are all areas for future work.

## References

- [1] M. Abadi. Security by typing in security protocols. In *Theoretical Aspects of Computer Science*, pages 611–638, 1997.
- [2] M. Abadi and B. Blanchet. Analyzing security protocols with secrecy types and logic programs. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 33–44, 2002.
- [3] M. Abadi, M. Burrows, B. Lampson, and G. Plotkin. A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, 1993. Also appeared as SRC Research Report 70.
- [4] M. Abadi, C. Fournet, and G. Gonthier. Secure communications processing for distributed languages. In *IEEE Symposium on Security and Privacy*, 1999.
- [5] M. Abadi, C. Fournet, and G. Gonthier. A top-down look at a secure message. In *Foundations of Software Technology and Theoretical Computer Science*, pages 122–141, 1999.
- [6] M. Abadi, C. Fournet, and G. Gonthier. Authentication primitives and their compilation. In *Proceedings of ACM Symposium on Principles of Programming Languages*, 2000.
- [7] M. Abadi and A. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1):1–70, January 1999.
- [8] M. Abadi, B. Lampson, M. Burrows, and E. Wobber. Authentication in distributed systems: Theory and practice. *ACM Trans. Comput. Syst.*, 10(4):265–310, 1992. Also appeared as SRC Research Report 83.
- [9] M. Abadi, E. Wobber, M. Burrows, and B. Lampson. Authentication in the Taos operating system. *ACM Transactions on Computer Systems*, 12(1):3–32, 1994. Also appeared as SRC Research Report 117.
- [10] B. Blanchet. From secrecy to authenticity in security protocols. In *9th International Static Analysis Symposium (SAS'02)*, pages 242–259, 2002.
- [11] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis. The role of trust management in distributed systems security. In J. Vitek and C. Jensen, editors, *Secure Internet Programming: Security Issues for Distributed and Mobile Objects*, volume 1603 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
- [12] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *IEEE Symposium on Security and Privacy*, 1996.
- [13] L. Cardelli, G. Ghelli, and A. D. Gordon. Secrecy and group creation. In *Concurrency Theory (CONCUR)*. Springer-Verlag, 2000.
- [14] G. Castagna and J. Vitek. A calculus of secure mobile computations. In *Internet Programming Languages*, Lecture Notes in Computer Science. Springer-Verlag, 1999.
- [15] D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Communications of the ACM*, 1977.
- [16] D. Duggan. Cryptographic types. In *Computer Security Foundations Workshop*, Nova Scotia, Canada, 2002. IEEE Press.
- [17] C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen. Simple public key certificate. Internet Draft draft-ietf-spki-cert-structure-05.txt, 1998.
- [18] A. D. Gordon and A. Jeffrey. Authenticity by typing for security protocols. In *IEEE Computer Security Foundations Workshop (CSFW)*, June 2001.
- [19] A. D. Gordon and A. Jeffrey. Types and effects for asymmetric cryptographic protocols. In *IEEE Computer Security Foundations Workshop (CSFW)*, June 2002.
- [20] M. Hennessy and J. Riely. Type-safe execution of mobile agents in anonymous networks. In *Secure Internet Programming: Security Issues for Distributed and Mobile Objects*, Lecture Notes in Computer Science. Springer-Verlag, 1999.
- [21] J. Mitchell and G. Plotkin. Abstract types have existential types. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, 1988.
- [22] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):528–569, 1999.
- [23] A. C. Myers and B. Liskov. Complete, safe information flow with decentralized labels. In *IEEE Symposium on Security and Privacy*, 1998.
- [24] A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4), 2000.
- [25] G. Necula. Proof-carrying code. In *Proceedings of ACM Symposium on Principles of Programming Languages*, 1997.
- [26] G. Necula and P. Lee. Safe kernel extensions without run-time checking. In *Operating Systems Design and Implementation*, 1996.
- [27] F. Pottier. A simple view of type-secure information flow in the pi-calculus. In *Computer Security Foundations Workshop*, 2002.
- [28] F. Pottier and S. Conchon. Information flow inference for free. In *Proceedings of ACM International Conference on Functional Programming*, 2000.
- [29] J. Riely and M. Hennessy. Trust and partial typing in open systems of mobile agents. In *Proceedings of ACM Symposium on Principles of Programming Languages*, 1999.
- [30] A. Sabelfeld and A. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 2002.
- [31] P. Sewell and J. Vitek. Secure composition of untrusted code: Wrappers and causality types. In *Computer Security Foundations Workshop*, 2000.
- [32] E. Sumii and B. C. Pierce. Logical relations for encryption. In *Computer Security Foundations Workshop*, June 2001.
- [33] D. Volpano and G. Smith. A type-based approach to program security. In *Proceedings of the International Joint Conference on Theory and Practice of Software Development*. Springer-Verlag, 1997.

- [34] S. Zdancewic and A. C. Myers. Robust declassification. In *Computer Security Foundations Workshop*, 2001.
- [35] S. Zdancewic, L. Zheng, N. Nystrom, and A. C. Myers. Untrusted hosts and confidentiality: Secure program partitioning. In *Symposium on Operating Systems Principles*, 2001.