Code-Carrying Authorization

Sergio Maffeis[2,3], Martín Abadi[1,2], Cédric Fournet[1], and Andrew D. Gordon[1]

[1] Microsoft Research
[2] University of California, Santa Cruz
[3] Imperial College London

Abstract. In authorization, there is often a wish to shift the burden of proof to those making requests, since they may have more resources and more specific knowledge to construct the required proofs. We introduce an extreme instance of this approach, which we call Code-Carrying Authorization (CCA). With CCA, access-control decisions can partly be delegated to untrusted code obtained at run-time. The dynamic verification of this code ensures the safety of authorization decisions. We define and study this approach in the setting of a higher-order spi calculus. The type system of this calculus provides the needed support for static and dynamic verification.

1 Introduction

The generation, transmission, and checking of evidence plays a central role in authorization. The evidence may include, for instance, certificates of memberships in groups, delegation assertions, and bindings of keys to principals. Typically, the checking is done dynamically, that is, at run-time, in reference monitors. When a reference monitor considers a request from a principal, it evaluates the evidence supplied by the principal in the context of a local policy and other information. It is also possible—and indeed attractive—to perform some of the checking statically, at the time of definition of a system. This static checking may rely on logical reasoning or on type systems, and may guarantee that enforcement of a policy is done thoroughly and correctly.

A growing body of research explores the idea that the evidence may include or may be organized as a logical proof [16,4,15,9,19]. For instance, in the special case of proof-carrying code (PCC), the proofs guarantee code safety, and the requests are typically for running a piece of code [16]. In another example, the clients of a web server may present proofs that their requests should be granted [5]. This idea provides a principled approach to authorization. It also provides an approach to auditing in which the proofs that motivate access-control decisions can be logged and analyzed [19]. While the burden of proof generation shifts to the principal that makes a request, the proof need not be trusted, so the reference monitor still needs to verify the proof. Dynamic proof verification may fail; accordingly, any static checking needs to accommodate this possibility.

Thus arises the question of how to reconcile static checking with proof-carrying and dynamic verification. As an interesting specific instance of this question, one may wonder how to incorporate dynamic verification in the existing typed spi calculus for authorization of Fournet et al. [12]. In that calculus, a static type system guarantees the safe enforcement of an authorization policy. It does not include proofs as first-class objects, nor the possibility of dynamic verification. One might think about adding proofs and proof-checking as primitives to this calculus, in order to support dynamic verification and authorization. While that idea may seem "natural", to our surprise we discovered that a more general idea is both

technically cleaner and more powerful in supporting interesting authorization scenarios. With "Proof-Carrying Authorization" (PCA) [4] in mind, we call this idea "Code-Carrying Authorization" (CCA).

CCA consists in passing not proofs but pieces of code that perform run-time verification. These pieces of code are essentially fragments of a reference monitor. They are themselves checked dynamically, since in general they are not trusted. Analogously, the Open Verifier project [8] has started to explore a generalization of PCC in which mobile code is accompanied by untrusted verifiers.

Following the Curry-Howard isomorphism, one may view proofs as programs. Still, with PCA [4], those programs are only checked, not executed. With CCA, programs are executed as well, though in a controlled way. No additional language for proofs is needed; we can use arbitrary code, subject to dynamic typing. Thus, in comparison with PCA, CCA allows a more open-ended, flexible notion of evidence without requiring the introduction of special syntax.

In the present paper, we explore dynamic verification and authorization in the context of a typed spi calculus. Technically, this calculus is a higher-order spi calculus [3] with dynamic typing. Both the higher-order features and the dynamic typing rely on fairly standard constructs [18,2], though with some new technical complications and new applications. In particular, the dynamic typing can require theorem proving. The calculus includes only shared-key cryptography; further cryptographic primitives might be added as in later work by Fournet et al. [13]. Optionally, the calculus also includes first-class proof hints, which can alleviate or eliminate the theorem-proving task at the reference monitor. We prove results that establish the safety of authorization decisions with respect to policies. (Appendix C contains detailed proofs.)

We exploit this calculus in a range of small but challenging examples. These examples illustrate some of the advantages of dynamic verification and of CCA in particular. For instance, in some of the examples, a server can enforce a rich authorization policy while having only simple, generic code; clients provide more detailed code for run-time access control. Such examples are beyond the scope of previous systems.

Sections 2 and 3 describe our calculus and its type system, respectively. Section 4 contains examples. Section 5 extends the calculus with proof hints. Section 6 concludes.

In addition to the research on PCA and on types for authorization cited above, our work is related to a broad range of applications of process calculi to security. These include, for instance, distributed pi calculi with trust relations and mobile code [17,14]. Interestingly, some of these calculi support remote attestation and dynamic subtyping checks (however, with rather different goals and type structures, and no typecase) [10].

2 A Spi Calculus with Dynamic Verification

In this section we review the calculus for authorization on which we build [12], and discuss our extensions for dynamic verification.

2.1 Authorization Logics

Our approach is parametric in the choice of an authorization logic used as a policy language. The only constraint on the logic is that it be monotonic and closed under substitution (see Appendix A). For example, Datalog [7], Binder [11], and CDD [1] are valid authorization

logics. In the rest of the paper, we use Datalog as an authorization logic, and write $S \models C$ when policy S entails the clause C. Informally, entailment means that access requests that depend on C should be granted according to S.

Our running example is based on an electronic conference reviewing system. The conference server contains a policy that controls the access to the database of paper reviews. This policy expresses authorization facts such as $PCMember(alice)$, which means "Alice has been appointed as a member of the program committee of the conference", or authorization rules such as

$$Review(U,ID,R) :- PCMember(U), Opinion(U,ID,R)$$

which means "if a committee member holds a certain opinion on any paper, she can submit a review for that paper". Capitalized variables such as U, ID, and R are bound logical variables. Lower-case identifiers (such as *alice* above), together with any other values of the process language, are uninterpreted logical atoms.

2.2 Process Syntax and Semantics

The core language consists of an asynchronous spi calculus where parallel processes can send messages to each other on named channels. For example, we may write:

$$\mathsf{out}\ a(M) \mid \mathsf{in}\ a(x); P \rightarrow P\{M/x\}$$

The symbol $\rightarrow$ represents a computation step. On the left of $\rightarrow$, we have a parallel composition of a process that sends a message (actually M) on the channel a and a process that receives a message (represented by the formal parameter x) on a and then executes P; on the right is the result, in which the formal parameter is replaced with the actual message.

Messages include channel names, cryptographic keys, pairs, and encryptions. We assume that encryption preserves the integrity of the payload. There are operations for decomposing and matching pairs and for decrypting messages. For example,

$$\mathsf{decrypt}\ \{M\}k\ \mathsf{as}\ \{y\}k; Q \rightarrow Q\{M/y\}$$

represents the only way to "open" the encryption $\{M\}k$ to retrieve M.

Two special constructs have no effects on the semantics of programs, but are annotations that connect the authorization policy to the protocol code: statements and expectations. A *statement*, such as $SentOn(a,b)$, should be manually inserted in the code in order to record that, at a particular execution point, the clause $SentOn(a,b)$ is regarded as true. An *expectation*, such as $\mathsf{expect}\ GoodParam(x)$, should label program points where the clause $GoodParam(x)$ must hold for the run-time value of x. For example, the following code is safe with respect to the policy $GoodParam(X) :- SentOn(a,X)$:

$$(\mathsf{out}\ a(b) \mid SentOn(a,b)) \mid \mathsf{in}\ a(x); (\mathsf{expect}\ GoodParam(x) \mid \mathsf{out}\ c(x))$$

To this core language, we add a new kind of message $(x{:}T)P$ that represents the process P parametrized by x of type T, and operations to spawn such processes and to check the type of messages dynamically. The formal syntax of messages and processes is as follows:

Syntax for Messages and Processes:

a,b,c,k,x,y,z	name
$M,N ::=$	message

x	name
$\{M\}N$	authenticated encryption of M with key N
(M,N)	message pair
$(x{:}T)P$	code P parametric in x
ok	token conveying logical effects (see Section 3)
$P,Q,R ::=$	process
out $M(N)$	asynchronous output of N to channel M
in $M(x{:}T);P$	input of x from channel M (x has scope P)
!in $M(x{:}T);P$	replicated input
new $x{:}T;P$	fresh generation of name x (x has scope P)
$P \mid Q$	parallel composition of P and Q
$\mathbf{0}$	null process
decrypt M as $\{y{:}T\}N;P$	bind y to decryption of M with key N (y has scope P)
split M as $(x{:}T,y{:}U);P$	solve $(x,y) = M$ (x has scope U and P; y has scope P)
match M as $(N,y{:}U);P$	solve $(N,y) = M$ (y has scope P)
spawn M with N	spawn M instantiated with N
typecase M of $x{:}T;P$	typecheck M at type T (x has scope P)
C	statement of clause C
expect C	expectation that clause C is derivable

Notations: $(\widetilde{x}{:}\widetilde{T}) \triangleq (x_1{:}T_1,\ldots,x_n{:}T_n)$ and new $\widetilde{x}{:}\widetilde{T};P \triangleq$ new $x_1{:}T_1;\ldots$ new $x_n{:}T_n;P$
Let $S = \{C_1,\ldots,C_n\}$. We write $S \mid P$ for $C_1 \mid \ldots \mid C_n \mid P$.

Type annotations help to understand the type of newly created names and variables, but are strictly necessary only in the syntax for typecase. For notational convenience, we may omit type annotations, especially for Un types.

Both spawn and typecase are standard constructs. However, in combination they turn out to be very useful for our purposes. For example, a verifier process can accept untrusted messages from the network, check that they are well-typed as processes with input of type T, and then send the code out to the network once again on an untrusted channel, wrapped in an encryption meant to signify that the contents are now guaranteed to be type-safe:

$$\text{in } unCode(x);\text{typecase } x \text{ of } y{:}\mathsf{Pr}(T);\text{out } tsCode(\{y\}k)$$

A code user can accept such encrypted code packages, and run the code passing it a parameter M of the correct type T without further checking:

$$\text{in } tsCode(x);\text{decrypt } x \text{ as } \{y\}k;\text{spawn } y \text{ with } M$$

As usual in the pi calculus, we define the formal semantics of the calculus by a set of structural congruence rules (see Appendix A) that describe what terms should be considered syntactically equivalent, and a set of reduction rules (displayed below) that describe how processes evolve. Most of these reduction axioms are standard. Rule (Red Typecase) requires some typing environment E in which the check $E \vdash M : T$ can be performed. In order to define such environments, we parametrize the reduction relation by an initial environment (which can also be chosen as $\varnothing$ if necessary). Rule (Red Res) dynamically adds the names defined by restriction contexts to the current typing environment, and rule (Red Par) adds the new clauses and names ($env(Q)^{\widetilde{x}}$) defined by parallel contexts. The technical reasons for these definitions, which should become apparent in Section 3, are illustrated in the following small example. Consider the reduction step:

$$\text{new } a{:}T;(\text{typecase } a \text{ of } y{:}T;P) \rightarrow_{\varnothing} \text{new } a{:}T;P\{a/y\}$$

Rules for Reduction: $P \rightarrow_E P'$

$\mathsf{out}\ a(M) \mid \mathsf{in}\ a(x{:}T);P \rightarrow_E P\{M/x\}$	(Red Comm)
$\mathsf{out}\ a(M) \mid \mathsf{!in}\ a(x{:}T);P \rightarrow_E P\{M/x\} \mid \mathsf{!in}\ a(x{:}T);P$	(Red !Comm)
$\mathsf{decrypt}\ \{M\}k\ \mathsf{as}\ \{y{:}T\}k;P \rightarrow_E P\{M/y\}$	(Red Decrypt)
$\mathsf{split}\ (M,N)\ \mathsf{as}\ (x{:}T,y{:}U);P \rightarrow_E P\{M,N/x,y\}$	(Red Split)
$\mathsf{match}\ (M,N)\ \mathsf{as}\ (M,y{:}U);P \rightarrow_E P\{N/y\}$	(Red Match)
$\mathsf{spawn}\ (x)P\ \mathsf{with}\ M \rightarrow_E P\{M/x\}$	(Red Spawn)
$E \vdash M : T \Rightarrow \mathsf{typecase}\ M\ \mathsf{of}\ y{:}T;P \rightarrow_E P\{M/y\}$	(Red Typecase)
$P \rightarrow_{E,env(Q)^{\tilde{x}}} P' \Rightarrow P \mid Q \rightarrow_E P' \mid Q \quad (\text{where } \{\tilde{x}\} \cap fn(P,Q) = \varnothing)$	(Red Par)
$P \rightarrow_{E,x:T} P' \Rightarrow \mathsf{new}\ x{:}T;P \rightarrow_E \mathsf{new}\ x{:}T;P'$	(Red Res)
$P \equiv Q, Q \rightarrow_E Q', Q' \equiv P' \Rightarrow P \rightarrow_E P'$	(Red Struct)

Notation: $P \rightarrow_E^{*\equiv} P'$ is $P \equiv P'$ or $P \rightarrow_E^* P'$.

By (Red Res), this reduction takes place if $\mathsf{typecase}\ a\ \mathsf{of}\ y{:}T;P \rightarrow_{a:T} P\{a/y\}$, and this is a valid instance of (Red Typecase) since the typing environment is now $a{:}T$, and $a{:}T \vdash a : T$ is clearly a valid typing judgment.

These rules allow a typecase process $\mathsf{typecase}\ M\ \mathsf{of}\ y{:}T;P$ to reduce provided the message M can be typechecked in an environment E that collects clauses and names defined in any parallel context. In an implementation, it may be impractical to collect the full environment, because, for example, E takes the form E', E'' where the clauses and names of E' are local, while those in E'' are distributed across remote machines. Still, it is fine for an implementation to typecheck the message in the local environment E', because, by a standard weakening lemma, if $E' \vdash M : T$ then also $E', E'' \vdash M : T$. Such an implementation would not admit reduction steps that depend on implicit knowledge of remote clauses and names. This is not a problem in our theory, as we are concerned with safety properties; in practice, we can convey knowledge of remote clauses and names by explicit use of cryptography, as in the examples in later sections.

For brevity, we use derived notations for tuples and pattern-matching, and omit type annotations when they are not necessary(see Appendix A for a formal definition). The tuple $(M_1, M_2, \ldots, M_n)$ abbreviates the nested pairs $(M_1, (M_2, \ldots, M_n))$. We write $\mathsf{tuple}\ M\ \mathsf{as}\ (\underline{N}_1, \ldots, \underline{N}_n);P$ to pattern-match a tuple, where M is a tuple, and each $\underline{N}_i$ is an atomic pattern (either a variable pattern x, or a constant pattern $=M$, where M is a message to be matched). For each variable, we introduce a split, and for each constant a match. For example, for a fresh z we have

$$\mathsf{tuple}\ (a,b,c)\ \mathsf{as}\ (x,{=}b,y);P \stackrel{\triangle}{=} \mathsf{split}\ (a,(b,c))\ \mathsf{as}\ (x,z);\mathsf{match}\ z\ \mathsf{as}\ (b,z);\mathsf{split}\ (z,z)\ \mathsf{as}\ (y,z);P$$

We also allow pattern-matching in conjunction with input and decryption processes.

2.3 Safety

Relying on the operational semantics, we give a formal definition of safety (much as in [12]). This notion makes precise the intuitive relation between assumptions, expectations, and program execution. The idea is that a process is safe if whenever during an execution the statement $\mathsf{expect}\ C$ is reached (i.e., it appears at the top level, possibly inside some nested name restrictions) the environment has accumulated enough rules and facts to entail C.

Safety:

A process P is *safe for* E if and only if whenever $P \rightarrow_E^{*\equiv} \mathsf{new}\ \tilde{x}{:}\tilde{T};(\mathsf{expect}\ C \mid P')$,
we have $P' \equiv \mathsf{new}\ \tilde{y}{:}\tilde{U};(S \mid P'')$ and $S \cup \mathit{clauses}(E) \models C$ with $(\{\tilde{y}\} \cap \mathit{fn}(C)) = \varnothing = (\{\tilde{x},\tilde{y}\} \cap \mathit{dom}(E))$.

The side-condition on the alpha-convertible names $\tilde{y}$ prevents confusing them with names in C, and the one on $\tilde{x},\tilde{y}$ avoids clashes with names defined in E.

It is also important to know when a process is safe even if it is executed in parallel with a malicious opponent. Following a common approach, we model the opponent as an arbitrary untyped process, with no statements or expectations.

Opponents and Robust Safety:

A process O is an *opponent* if and only if it contains no statement or expectation, and every type annotation is Un.
A process P is *robustly safe for* E if and only if for any opponent O, $P \mid O$ is safe for $E,\tilde{x}{:}\widetilde{\mathsf{Un}}$, where $\tilde{x}$ are the free names of O not in the domain of E.

For example, the process $P = \mathsf{out}\ b(a) \mid \mathsf{in}\ b(x);\mathsf{expect}\ A(x)$ is safe for $A(a)$, but not robustly safe, as an opponent that replaces a with c can lead to an unsatisfied expectation: $\mathsf{in}\ b(x);\mathsf{out}\ b(c) \mid P \rightarrow_{A(a)}^{*} \mathsf{expect}\ A(c)$.

3 A Type System for Robust Safety

We present a dependent type system that statically guarantees safety and robust safety. We extend the system of [12] with a type constructor $\mathsf{Pr}(T)$ for process code parametric in T, and rules for the spawn and typecase constructs. Most of the rules in this section (including those for new constructs) are largely standard rules adapted to the present context. We are pleased by how much advantageous reuse has been possible.

We prove that typability with respect to an environment E entails safety for E and, if all the types in E are Un ("untrusted"), also robust safety.

3.1 Types and Environments

Type Un is inhabited by any message that may come or go to the opponent, like for example a ciphertext that can be considered untrusted until it is decrypted. Upon decryption, one may reason that the contents were created by a principal that knows the encryption key. Types $\mathsf{Ch}(T)$ and $\mathsf{Key}(T)$ are inhabited by secure channels or secret keys for communicating or encrypting messages of type T. A dependent type $(x{:}T,U)$ is inhabited by the pairs (M,N) where M has type T, and N has type $U\{M/x\}$. Type $\mathsf{Ok}(S)$ is inhabited only by the token ok, and is used to attach effects to the payload of channels and keys. When a variable in the environment has type $\mathsf{Ok}(S)$, it is safe to assume that S holds.

Syntax for Types:

$T,U ::=$	type
Un	public data
$\mathsf{Ch}(T)$	channel for messages of type T
$\mathsf{Key}(T)$	secret key for plaintexts of type T
$(x{:}T,U)$	dependent pair (scope of x is U)
$\mathsf{Pr}(T)$	process code parametric in type T
$\mathsf{Ok}(S)$	ok to assume the clauses S

T is *generative* iff T is of the form Un, $\mathsf{Ch}(U)$, or $\mathsf{Key}(U)$, for some U.
Notation: $(x_1{:}T_1,\ldots,x_n{:}T_n,T_{n+1}) \triangleq (x_1{:}T_1,\ldots,(x_n{:}T_n,T_{n+1}))$

For example, the type declaration

$$kra : \mathsf{Key}(id{:}\mathsf{Un}, r{:}\mathsf{Un}, \mathsf{Ok}(Opinion(alice,id,r)))$$

says that *kra* is a key for encrypting a tuple like (*paper*,*text*,ok) where *paper* and *text* are untrusted values and the ok token indicates that the key conveys the logical effect *Opinion*(*alice*,*paper*,*text*).

Typing environments are lists of name bindings and clauses. We write $dom(E)$ for the set of names defined (i.e., appearing to the left of a binding ":") in environment E. We write $env(P)$ for the top-level clauses of process P, with suitable name bindings for any top-level restrictions, and $clauses(E)$ for the clauses contained at the top level and inside the top-level Ok types of E.

Syntax for Environments, and Functions: $dom(E)$, $env(P)$, $clauses(E)$

$E ::=$	environment
$\varnothing$	empty
$E,x{:}T$	x has type T
E,C	C is a valid clause

Notation: $E(x) = T$ if $E = E',x{:}T,E''$

$$dom(E,C) = dom(E) \quad dom(E,x{:}T) = dom(E) \cup \{x\} \quad dom(\varnothing) = \varnothing$$

$$\begin{aligned}
clauses(\varnothing) &= \varnothing \\
clauses(E,x{:}T) &= clauses(E) && \text{(if } T \neq \mathsf{Ok}(S)\text{)} \\
clauses(E,C) &= clauses(E) \cup \{C\} \\
clauses(E,x{:}\mathsf{Ok}(S)) &= clauses(E) \cup S \\
env(P \mid Q)^{\widetilde{x},\widetilde{y}} &= env(P)^{\widetilde{x}}, env(Q)^{\widetilde{y}} && \text{(where } \{\widetilde{x},\widetilde{y}\} \cap fn(P \mid Q) = \varnothing\text{)} \\
env(\mathsf{new}\ x{:}T;P)^{x,\widetilde{x}} &= x{:}T, env(P)^{\widetilde{x}} && \text{(where } \{\widetilde{x}\} \cap fn(P) = \varnothing\text{)} \\
env(C)^{\varnothing} &= C \\
env(P)^{\varnothing} &= \varnothing && \text{(otherwise)}
\end{aligned}$$

Convention: $env(P) \triangleq env(P)^{\widetilde{x}}$ for some distinct $\widetilde{x}$ such that $env(P)^{\widetilde{x}}$ is defined.

We assume a standard notion $E \vdash \diamond$ of well-formedness for environments, defined formally in Appendix A.

3.2 Typing Rules

For each message constructor there are two typing rules, one to give it an informative type, and one to give it type Un. Rules of the second kind are useful to show that any opponent process can be typed.

Rule (Msg Encrypt) shows that an encryption under a trusted key does not need to be trusted, in the sense that it can be sent to an opponent. Rules (Msg Proc) and (Msg Proc Un) invoke the typing relation for processes in an environment that assumes respectively type T or type Un for the process parameter x. Rule (Msg Ok) is typical of this typed approach to verification: in order for an ok token to convey the effects S, it must be the case that the clauses contained in the environment (which include the policy and all the facts consequently accumulated by Ok types) entail each of the clauses in S.

Rules for Messages: $E \vdash M : T$

(Msg x)
$$\frac{E \vdash \diamond \quad x \in dom(E)}{E \vdash x : E(x)}$$

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

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

(Msg Pair)
$$\frac{E \vdash M : T \quad E \vdash N : U\{M/x\}}{E \vdash (M,N) : (x{:}T, U)}$$

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

(Msg Ok Un)
$$\frac{E \vdash \diamond}{E \vdash \mathsf{ok} : \mathsf{Un}}$$

(Msg Proc)
$$\frac{E, x{:}T \vdash P}{E \vdash (x{:}T)P : \mathsf{Pr}(T)}$$

(Msg Proc Un)
$$\frac{E, x{:}\mathsf{Un} \vdash P}{E \vdash (x{:}\mathsf{Un})P : \mathsf{Un}}$$

(Msg Ok)
$$\frac{E, S \vdash \diamond \quad clauses(E) \models C \quad \forall C \in S}{E \vdash \mathsf{ok} : \mathsf{Ok}(S)}$$

Rule (Proc Res) requires to type P in an environment with the additional binding $x{:}T$. Correspondingly, the reduction rule (Red Res) assumes the binding in the run-time environment of its premise. Rule (Proc Par) collects the effects of process Q to typecheck P, and vice versa. Similarly, the premise of (Red Par) assumes $env(Q)$ in the environment of its premise. Rule (Proc Expect) requires an expected clause to be entailed by the environment, much in the same way as (Msg Ok). Rule (Proc Typecase) is somewhat subtle. It corresponds to an Un rule if we pick U and T to be Un. Moreover, the type U is not related *a priori* to the type T. In typical examples, the rule allows us to check a message M received at type Un and bind a variable y of some more useful type T to this message if the check succeeds. The remaining rules come in pairs, with one rule that assumes informative types and one that assumes Un types. Most of them are straightforward. For example, (Proc Output) says that a message of type T can be sent on a channel of type $\mathsf{Ch}(T)$, and (Proc Decrypt) says that the variable y that represents the payload of a ciphertext of type Un decrypted with a key of type $\mathsf{Key}(T)$ can be assumed to have type T in the continuation process. The rules for split and match are in Appendix A.

Rules for Processes: $E \vdash P$

(Proc Nil)
$$\frac{E \vdash \diamond}{E \vdash \mathbf{0}}$$

(Proc Res)
$$\frac{E, x{:}T \vdash P \quad T \text{ generative}}{E \vdash \mathsf{new}\ x{:}T; P}$$

(Proc Fact)
$$\frac{E, C \vdash \diamond}{E \vdash C}$$

(Proc Expect)
$$\frac{E, C \vdash \diamond \quad clauses(E) \models C}{E \vdash \mathsf{expect}\ C}$$

(Proc Par)
$$\frac{E, env(Q) \vdash P \qquad E, env(P) \vdash Q \qquad fn(P \mid Q) \subseteq dom(E)}{E \vdash P \mid Q}$$

(Proc Typecase)
$$\frac{E \vdash M : U \qquad E, x : T \vdash P}{E \vdash \mathsf{typecase}\ M\ \mathsf{of}\ x{:}T; P}$$

(Proc Spawn)
$$\frac{E \vdash M : \mathsf{Pr}(T) \quad E \vdash N : T}{E \vdash \mathsf{spawn}\ M\ \mathsf{with}\ N}$$

(Proc Spawn Un)
$$\frac{E \vdash M : \mathsf{Un} \quad E \vdash N : \mathsf{Un}}{E \vdash \mathsf{spawn}\ M\ \mathsf{with}\ N}$$

(Proc Input)
$$\frac{E \vdash M : \mathsf{Ch}(T) \qquad E, x{:}T \vdash P}{E \vdash [!]\mathsf{in}\ M(x{:}T); P}$$

(Proc Input Un)
$$\frac{E \vdash M : \mathsf{Un} \qquad E, x{:}\mathsf{Un} \vdash P}{E \vdash [!]\mathsf{in}\ M(x{:}\mathsf{Un}); P}$$

(Proc Output)
$$\frac{E \vdash M : \mathsf{Ch}(T) \quad E \vdash N : T}{E \vdash \mathsf{out}\ M(N)}$$

(Proc Output Un)
$$\frac{E \vdash M : \mathsf{Un} \quad E \vdash N : \mathsf{Un}}{E \vdash \mathsf{out}\ M(N)}$$

(Proc Decrypt)
$$\frac{E \vdash M : \mathsf{Un} \quad E \vdash N : \mathsf{Key}(T) \quad E, y{:}T \vdash P}{E \vdash \mathsf{decrypt}\ M\ \mathsf{as}\ \{y{:}T\}N; P}$$

(Proc Decrypt Un)
$$\frac{E \vdash M : \mathsf{Un} \quad E \vdash N : \mathsf{Un} \quad E, y{:}\mathsf{Un} \vdash P}{E \vdash \mathsf{decrypt}\ M\ \mathsf{as}\ \{y{:}\mathsf{Un}\}N; P}$$

Notation: brackets denote optional constructs.

As a simple example, we can show that for $E = Bar{:}-Foo, b{:}\mathsf{Ch}(\mathsf{Ok}(Bar))$, the typing judgment $E \vdash Foo \mid \mathsf{out}\ b(\mathsf{ok})$ is valid. The judgment follows by an instance of (Proc

Par), from $E \vdash Foo$ and $E, Foo \vdash \mathsf{out}\ b(\mathsf{ok})$. The latter in turn follows by (Proc Output) and (Msg Ok), where the second rule uses the logical inference $clauses(E, Foo) \models Bar$. Section 4 includes a longer, detailed example of how the interplay between static and dynamic typechecking makes this type system expressive.

3.3 Results

We obtain a type preservation result and a safety theorem that guarantees that typability implies safety.

Lemma 1 (Type Preservation). *If $E \vdash P$ and $P \rightarrow_E^{*\equiv} P'$ then $E \vdash P'$.*

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

The safety theorem makes explicit the connection between the environment used for typing (existentially quantified in related work), and the run-time environment.

In order to show that our notion of opponent is not restrictive in a typed setting, we prove that any opponent can be typed in an environment that does not make trust assumptions. Finally, we prove that if a process P is safe for a security policy S and an untrusted environment, then it is robustly safe.

Lemma 2 (Opponent Typability). *For opponent O, $\widetilde{x}{:}\widetilde{\mathsf{Un}} \vdash O$, where $fn(O) \subseteq \{\widetilde{x}\}$.*

Theorem 2 (Robust Safety). *If $\widetilde{x}{:}\widetilde{\mathsf{Un}}, S \vdash P$ then P is robustly safe for $\widetilde{x}{:}\widetilde{\mathsf{Un}}, S$.*

For example, let us consider process $Q = \mathsf{out}\ b(a, \mathsf{ok}) \mid \mathsf{in}\ b(x,y); \mathsf{expect}\ A(x)$. It is easy to see that given $E = a{:}\mathsf{Un}, b{:}\mathsf{Ch}(x{:}\mathsf{Un}, A(x)), A(a)$ we have $E \vdash Q$, so Q is safe for E. On the other hand it is not possible to derive $a{:}\mathsf{Un}, b{:}\mathsf{Un}, A(a) \vdash Q$, so we cannot prove robust safety (which does not hold).

3.4 Dynamic Verification

We define a derived construct and typing rule to verify that a piece of code M, when passed a parameter N of type T enforces property S. The idea is to typecheck dynamically M, against the parameter type T and an implicit parameter c that is a channel used to return the result of verification, namely an ok token carrying the effects S. The continuation process P will execute only if verification succeeds, that is M sends an ok on channel c.

Derived Syntax and Derived Typing Rule for Verification: $\mathsf{verify}\ M\langle\widetilde{N{:}T}\rangle{:}S;P$

$$\mathsf{verify}\ M\langle[\widetilde{z{:}\mathsf{Un}}, N{:}T]\rangle{:}S;P \triangleq \mathsf{new}\ c{:}\mathsf{Ch}(\mathsf{Ok}(S)); \big(\mathsf{typecase}\ M\ \mathsf{of}\ y{:}\mathsf{Pr}([\widetilde{z{:}\mathsf{Un}}, T,]\mathsf{Ch}(\mathsf{Ok}(S)))); \mathsf{spawn}\ y\ \mathsf{with}\ ([\widetilde{z}, N,]c) \mid \mathsf{in}\ c(x{:}\mathsf{Ok}(S)); P\big)$$
(where $\{c,y,x\} \cap fn(P, M, [N,]S) = \varnothing$, and $\{\widetilde{z}\} \subseteq fn(S)$)

(Proc Verify)
$$\frac{E \vdash M : U \quad [E \vdash \widetilde{N : T}] \quad E, S \vdash P}{E \vdash \mathsf{verify}\ M\langle[\widetilde{N{:}T}]\rangle{:}S;P}$$

With this derived typing rule, it is easy to see that if $E \vdash \mathsf{verify}\ M\langle\widetilde{N{:}T}\rangle{:}S;P$ then $E \vdash \mathsf{verify}\ M\langle\widetilde{N{:}T}\rangle{:}S;(\mathsf{expect}\ S \mid P)$. Hence, it is not necessary to annotate code with an expectation after a verification step.

One may wonder whether it is prudent to run the code of an untrusted verifier that is guaranteed to enforce a certain policy. Although additional precautions may be appropriate,

this guarantee is substantial. By lexical scoping, the code of the verifier cannot contain capabilities that are not already known by its generator; other capabilities can only be passed explicitly as parameters. Moreover, the verifier must be well-typed in the run-time typing environment, which can be restricted conveniently to further limit potential side effects. On the other hand, this guarantee does not cover other kinds of attacks (such as information leaks or denial-of-service attacks), which may be addressed independently.

4 Examples: a Conference Program Committee

As a benchmark for the effectiveness of CCA, we revisit the conference program committee example of [12]. We first review the idealized electronic conference system, then present examples that illustrate the benefits of CCA.

4.1 Review: an Electronic Conference Reviewing System

There are three kinds of principals: the program committee chair (pc-chair), identified with the server, the program committee members (pc-members), and potential reviewers. The last two are clients of the server. We model only the portion of the conference reviewing system for delegating and filing reviews. The authorization policy S, from the subjective viewpoint of the pc-chair, is:

$$
\begin{aligned}
S = \; & \mathit{Review}(U,ID,R) :- \mathit{Reviewer}(U,ID),\mathit{Opinion}(U,ID,R) \\
& \mathit{Review}(U,ID,R) :- \mathit{PCMember}(U),\mathit{Opinion}(U,ID,R) \\
& \mathit{Reviewer}(V,ID) :- \mathit{Reviewer}(U,ID),\mathit{Delegate}(U,V,ID) \\
& \mathit{Delegate}(U,W,ID) :- \mathit{Delegate}(U,V,ID),\mathit{Delegate}(V,W,ID) \\
& \mathit{Delegate}(U,U,ID) :- \mathit{Opinion}(U,ID,R)
\end{aligned}
$$

The predicate $\mathit{Opinion}(u,id,r)$ states that principal u holds opinion r on paper id, and is under the control of u itself (that is, the code identified with u can freely assert that predicate). The predicate $\mathit{Delegate}(u,v,id)$ states that principal u delegates its capability to review paper id to principal v, and is also under the control of u. All the other predicates are controlled by the pc-chair, and should be asserted only within server code.

Cryptographic keys can be associated with each of these predicates to convey authorization facts through untrusted messages. Thus, the pc-chair may appoint *alice* as a pc-member by sending her a token $\{alice\}kp$ encrypted under a key that carries the effect $\mathit{PCMember}(alice)$, and similarly for the other predicates. We define the type of the keys that correspond to each effect, and the type of a channel that implements a database where the pc-chair stores the keys of all potential users:

$$
\begin{aligned}
KA &= \mathsf{Key}(u{:}\mathsf{Un},id{:}\mathsf{Un},\mathsf{Ok}(\mathit{Reviewer}(u,id))) \\
KP &= \mathsf{Key}(u{:}\mathsf{Un},\mathsf{Ok}(\mathit{PCMember}(u))) \\
KD &= \mathsf{Key}(z{:}\mathsf{Un},id{:}\mathsf{Un},\mathsf{Ok}(\mathit{Delegate}(v,z,id))) \\
KR &= \mathsf{Key}(id{:}\mathsf{Un},r{:}\mathsf{Un},\mathsf{Ok}(\mathit{Opinion}(v,id,r))) \\
T &= \mathsf{Ch}(v{:}\mathsf{Un},(KD,KR))
\end{aligned}
$$

Keys of type KA or KP are used by the pc-chair only, to assign a paper to a reviewer or to appoint a pc-member respectively. Keys of type KD or KR (parametric in v) can be used by principal v to convey either an opinion or a delegation effect. Type T is the type of a channel used to retrieve the keys of each registered user. Note that it is a dependent type that binds the free parameter v of types KD and KR.

4.2 Off-line Delegation

Our first example presents a system that lets reviewers appoint sub-reviewers without involving the pc-chair in the process. A typical solution that does not use CCA is to have a reviewer present to the server a request that contains her opinion, together with some evidence that represents a chain of delegation. The server then runs an algorithm to traverse the chain and check corresponding permissions, and grants access if the evidence is satisfactory. This solution commits the server to a specific verification algorithm (or a fixed number thereof). Using CCA instead, the server code can be simpler and parametric. For example, the server is defined by the same code whether or not the delegation chain is ordered, has limited length, or delegation is permitted at all. Along with each request to file a review, the server receives the code of a verifier and some evidence. It verifies that the code enforces the desired authorization policy, and grants access without further checks.

The relevant portion of the server code is:

```
Server(pwdb:T,ka:KA,kp:KP) =
S | !in filereview(v,id,r,p,e);
    verify p⟨(v,r,e,(pwdb,ka,kp)):(v:Un,r:Un,Un,(T,KA,KP))⟩:Review(v,id,r); [...]
```

It contains the assertion of policy S, and a process always ready to accept messages on the public channel *filereview*. Parameters v, id, and r are interpreted as a request from principal v to file review r on paper id. Parameter p is the code of a verifier that must be run to grant authorization (i.e., prove $Review(v,id,r)$) on data including the evidence received as the last parameter e, and local credentials provided by the server. The parameters passed by the server to the verifier p are the name v of the principal issuing the request, the report r, the evidence e, and a triple $(pwdb, ka, kp)$. Channel $pwdb$ can be used to retrieve user credentials. Keys ka and kp are the secret keys used by the pc-chair to appoint reviewers and pc-members. If verification succeeds, authorization is granted, and r is a valid review for id.

A delegate v receives from a reviewer a request to review paper id, with additional parameters p (the verifier code to be passed on to the server), and dc (the evidence that represents a chain of delegation). The delegate may appoint another sub-reviewer, adding a delegation step to the chain $(v, \{u, id, ok\}kdv, dc)$, or file a review, adding evidence of its opinion to the top of the chain:

```
Delegate(v:Un,krv:KR,kdv:KD) =
!in reviewrequest(=v,id,p,dc);
(in accept(r); Opinion(v,id,r) | out filereview(v,id,r,p,({id,r,ok}krv,dc)) |
(in delegate(u); Delegate(v,u,id) | out reviewrequest(u,id,p,(v,{u,id,ok}kdv,dc)))
```

The pc-member can embed its logical effects directly in the verification code. For that reason, it transmits as evidence ok tokens with empty logical effects. The verifier *fver*, used to file a review ignores the principal name and the evidence, states that v holds opinion r on id, parses the server credentials to get the key to appoint pc-members, proves that v is a pc-member, by decrypting the appointment token (passed by the server earlier on), and finally signals success.

The (commented) verifier code *dver* involves a loop to gather and verify all the elements of the delegation chain:

```
PCMember(v:Un,pctoken:Un) =
!in paperassign(=v,id,idtoken);
(in review(r); out filereview(v,id,r,fver,ok) |
```

```
(in delegate(u); out reviewrequest(u,id,dver,ok))

fver = (_,_,keys,return) (Opinion(v,id,r) | tuple keys as (_,_,kp);
       decrypt pctoken as {=v,_}kp; out return(ok))

dver = (z,r,evid,keys,return) (Delegate(v,u,id) | // implicit effect
       tuple evid as (op,dc); // parse the evidence
       tuple keys as (pwdb,ka,_); // parse server keys
       in pwdb(=z,kdz,krz); // retrieve the credentials for z
       decrypt op as {=id,=r,_}krz; // check that z has opinion r on id
       new link:Ch(u:Un,dc:Un,Ok(Delegate(u,z,id)));
       out link(z,dc,ok) | !in link(w,dc,_); // start the delegation chain loop
         ( tuple dc as (t,del,dc); in pwdb(=t,kdt,_);
           decrypt del as {=w,=id,_}kdt; out link(t,dc,ok)) |
         // loop repetition: check delegation step
         ( tuple dc as =ok; decrypt idtoken as {=v,=id,_}ka; out return(ok))
         // end of the loop: check reviewer token
```

This code, and a few additional code fragments not shown here, can be assembled into a program that represents the entire conference reviewing system. This program typechecks in an environment of the form $\widetilde{x}{:}\widetilde{\mathsf{Un}}$ (according to the rules of Section 3). Therefore, Theorem 2 applies, and guarantees robust safety. In this particular case, this theorem implies that expectations in the server code, such as *Review*(*v*,*id*,*r*), are always satisfied at run-time when they occur, even in an untrusted environment.

4.3 Server-Side Proxy

Our second example illustrates the use of verifiers as server-side proxies installed by clients. It illustrates the flexibility of using typecase and spawn independently from the derived verify construct.

We modify our previous example so that the pc-member sends the delegation verifier *dver* directly to the server, which can use it to authorize requests from delegated reviewers. We show the code for dealing with delegated reviews, which is the most interesting. The server registers proxies for each pc-member, and accepts requests on each proxy. A message on the public channel *newproxy* causes the server to typecheck the code *dver* and install it as a handler and verifier for requests coming from reviewers delegated by pc-member *u*:

```
Server(pwdb:T,ka:KA,kp:KP) =
S | new protectedfilereview:V;
      (!in newproxy(dver); typecase dver is y:Pr(U);
            spawn y with ((pwdb,ka,kp),protectedfilereview)
      |!in protectedfilereview(v,id,r,_); expect Review(v,id,r); [...])
U = ((T,KA,KP),V)
V = Ch(v:Un,id:Un,r:Un,Ok(Review(v,id,r)))
```

Once appointed, a pc-member installs its delegation proxy on the server. The proxy receives requests from delegates on a dedicated channel and authorizes them. Upon delegation, the pc-member needs to send to the delegate a request that contains the name of the dedicated channel and evidence of delegation. The evidence consists of a delegation chain that contains a delegation step $\{u,id,\mathsf{ok}\}kdv$ (the name of the delegate and the paper id encrypted under the delegation key of the pc-member, and an ok token) and the list terminator (another ok token):

```
PCMember(v:Un,pctoken:Un) =
!in paperassign(=v,id,idtoken);
new filesubreview:Un;
    out newproxy(dver) |
    (in delegate(u); out reviewrequest(u,id,filesubreview,({u,id,ok}kdv,ok)))
```

The verifier *dver* now installs a process ready to listen to delegate requests on channel *filesubreview*, and then verifies requests similarly to the code shown above for off-line delegation. The main differences are that, in this case, the result returned by the verification process needs to contain the parameters v, id, r of the effect *Review*(*v*,*id*,*r*) to be enforced, and the code does not contain the implicit delegation effect *Delegate*(*v*,*u*,*id*):

```
dver = (keys,return) !filereview(z,id,r,evid);
       tuple evid as (op,dc); // parse the evidence
       tuple keys as (pwdb,ka,_); // parse server keys
       in pwdb(=z,kdz,krz); // retrieve the credentials for z
       decrypt op as {=id,=r,_}krz; // check that z has opinion r on id
       new link:Ch(u:Un,dc:Un,Ok(Delegate(u,z,id)));
       out link(z,dc,ok) | !in link(w,dc,_); // start the delegation chain loop
         ( tuple dc as (t,del,dc); in pwdb(=t,kdt,_);
           decrypt del as {=w,=id,_}kdt; out link(t,dc,ok)) |
         // loop repetition: check delegation step
         ( tuple dc as =ok; decrypt idtoken as {=v,=id,_}ka; out return(z,id,r,ok))
         // end of the loop: check reviewer token
```

The code for the delegate is little changed. It files reviews on the dedicated channels, or delegates further:

```
Delegate(v:Un,krv:KR,kdv:KD) =
!in reviewrequest(=v,id,filereview,dc);
(in accept(r); Opinion(v,id,r) | out filereview(v,id,r,({id,r,ok}krv,dc)) |
(in delegate(u); Delegate(v,u,id) | out reviewrequest(u,id,filereview,(v,{u,id,ok}kdv,dc)))
```

4.4 Best-Effort Evidence

Our third example presents a system that supports the possibility for reviewers to appoint sub-reviewers, without needing immediate access to their delegation credentials. In a completely static type system, a typical delegation protocol such as the one presented in the previous section needs to record in a delegation chain the causal relation between delegation steps. Hence, a reviewer that momentarily does not have access to its delegation key cannot appoint a sub-reviewer.

We present a protocol that is well-typed, hence guarantees that, each time authorization to file a review is granted, the requesting principal is provably a reviewer. Yet, the protocol is "best-effort", in that authorization can be denied at run-time if the server has not yet received all the delegation messages necessary to reconstruct a valid delegation chain.

To simplify the presentation, and to illustrate another advantage of CCA, we present code that does not use cryptography. Suppose that the machine of the reviewer is down, so she picks up the phone and asks a sub-reviewer to review a paper and to send his opinion (in the form of a simple verifier) to the server, trusting that the review will be accepted. The sub-reviewer can do so, or delegate further by issuing another informal request and by separately contacting the server to communicate his delegation decision:

```
Delegate(v:Un) =
!in phonereviewrequest(=v,id);
  (in accept(r); out filereview(v,id,r,fver))
|(in delegate(u); out phonereviewrequest(u,id) | out latedelegation(v,u,id,dver))
fver = (return)(Opinion(v,id,r)|out return(ok))
dver = (return)(Delegate(v,u,id)|out return(ok))
```

The server independently accepts requests for filing reviews and messages that state delegation decisions. In the first case, the server simply verifies that the review can be filed; in the second case it verifies that it is safe to assert a delegation step. At run-time the server authorizes the request to file a review from a delegate only if it has already verified enough delegation evidence to form a chain that originates from an appointed reviewer:

```
Server() =
S | PCMember(alice) | Reviewer(bob,42)
  | (!in filedreview(v,id,r,fver); verify fver⟨⟩:Review(v,id,r); [...])
  | (!in latedelegation(v,u,id,dver); verify dver⟨⟩:Delegate(v,u,id);Delegate(v,u,id))
```

In previous static systems, this sort of best-effort code was not possible. The code had to be written so that the expectation *Review*(*v*,*id*,*r*) could occur only after code that would check the necessary delegation facts.

5 From Theorem Proving to Proof Checking

We have shown how to pass and dynamically check the code of a verifier process. The dynamic check may involve invoking a theorem prover, potentially a costly operation. On the other hand, passing proofs only requires the receiving side to have a proof checker, reducing both the trusted computing base and the performance cost of verification. For this reason, we extend our framework with the capability to pass also *hints*, that can help the receiver of a reference monitor with the logical proofs involved in dynamic typechecking. Hints could be proofs, in the formal sense of the word, or any other kind of information which may (or may not) be helpful. In particular, hints could be incomplete proofs, that simplify rather than eliminate theorem proving.

5.1 From ok s to Hints

The ok token can already be interpreted as an empty hint, that leaves to the typechecker the burden of finding a proof. We parametrize ok tokens by a generic language of (possibly empty) proof hints H. Hints may contain variables, so that they can be combined at run-time to form larger hints. Expectations now mention a term that can be used as a hint to prove C.

Syntax for Hints

$M,N ::= \mathsf{ok}\,H \mid \ldots$	proof hint H replaces ok
$P,Q,R ::= \mathsf{expect}\ C\ \mathsf{by}\ M \mid \ldots$	expectation that clause C is derivable by M replaces expect C

The notion of type-safety does not change (just replace expect C by expect C by M), since the final result that we desire is still that any expectation is justified by logical entailment. It is the verification process that can be made simpler by adopting a verification relation, which naturally should imply entailment.

Verification Relation: $\mathscr{V}(M,C,S)$

Given an authorization logic $(\mathscr{C}, fn, \models)$, we assume an abstract verification predicate $\mathscr{V}$ that holds only if a message M is a proof of clause C starting from policy S, and such that $\mathscr{V}(M,C,S) \Rightarrow S \models C$.

We use hints and the verification relation in the typing rules that involve logical effects. In particular, we only need to replace (Msg Ok), (Msg Ok Un), and (Proc Expect) by the corresponding typing rules given below.

Typing Rules for Hints

(Msg Hint)
$$\frac{E,S \vdash \diamond \quad fn(H) \subseteq dom(E) \quad \mathscr{V}(H,C,clauses(E)) \quad \forall C \in S}{E \vdash \mathsf{ok}\, H : \mathsf{Ok}(S)}$$

(Msg Hint Un)
$$\frac{E \vdash \diamond \quad fn(H) \subseteq dom(E)}{E \vdash \mathsf{ok}\, H : \mathsf{Un}}$$

(Proc Expect Hint)
$$\frac{E,C \vdash \diamond \quad E \vdash M : \mathsf{Ok}(S) \quad C \in S}{E \vdash \mathsf{expect}\ C\ \mathsf{by}\ M}$$

The rules for hints are the obvious adaptations of the corresponding rules for ok. Note that verification can assume as lemmas the effects of hints that are just variables, because they are included by $clauses(E)$ in the premise of (Msg Hint). Rule (Proc Expect Hint) no longer involves verification directly. It is the premise needed to give M the $\mathsf{Ok}(S)$ type that may involve proof-checking.

This type system conservatively extends the one without hints. In fact, the type system presented in Section 3 correspond exactly to the instance of the current type system where H is empty, each expectation is of the form $\mathsf{expect}\ C\ \mathsf{by}\ \mathsf{ok}$, and $\mathscr{V}(M,C,S)$ is defined as $S \models C$.

Theorem 3 (Safety with Hints). *(i) If $E \vdash P$ then P is safe for E. (ii) If $\widetilde{x{:}\mathsf{Un}}, S \vdash P$ then P is robustly safe for $\widetilde{x{:}\mathsf{Un}}, S$.*

The theorem holds for any choice of the verification function $\mathscr{V}$, as long as it implies logical entailment $\models$. In particular, if $\mathscr{V}$ is not monotonic or not closed under substitutions, the type system is safe but may not enjoy type preservation.

The syntactic sugar from Section 4 can be adapted easily to hints by making explicit the variable x that is bound to the hint that results from the verification process, so that it can be used in subsequent expectations, or to build more complex hints.

Derived Typing Rule for Verification with Hints: $\mathsf{verify}\ M\langle N{:}T\rangle{:}x{:}S; P$

(Proc Verify Hint)
$$\frac{E \vdash M : U \quad \widetilde{[E \vdash N : T]} \quad E, x{:}S \vdash P}{E \vdash \mathsf{verify}\ M\langle \widetilde{[N{:}T]}\rangle{:}x{:}S; P}$$

5.2 Verification in Datalog

For the examples, we use the simple hint language and logical verification relation for Datalog defined below, where $S \models_1 C$ is the single-step entailment relation.

For example, considering $S = D{:-}C, C{:-}B, B{:-}A, A$ and $S_1 = D{:-}C, C{:-}B, B$ and $S_2 = C, D{:-}C$, we have that $\mathscr{V}(\mathsf{ok}\,(S_1,S_2), D, S)$ follows by an instance of (Verify Pair) with premises $\mathscr{V}(\mathsf{ok}\, S_1, C, S)$, $\mathscr{V}(\mathsf{ok}\, S_1, D{:-}C, S)$, and $\mathscr{V}(\mathsf{ok}\, S_2, D, S_1)$.

Hints and Verification

$H ::= S \mid M$ proof hint: clauses S or message M

(Verify S)
$$\frac{S \models_1 C' \quad \forall C' \in S' \qquad S' \models_1 C}{\mathscr{V}(\mathsf{ok}\,S', C, S)}$$

(Verify Pair)
$$\frac{\mathscr{V}(\mathsf{ok}\,M_1, C', S) \quad \forall C' \in \overline{M_2} \qquad \mathscr{V}(\mathsf{ok}\,M_2, C, \overline{M_1})}{\mathscr{V}(\mathsf{ok}\,(M_1, M_2), C, S)}$$

$\overline{\mathsf{ok}\,S} = S \qquad \overline{(M_1, M_2)} = \overline{M_1} \cup \overline{M_2} \qquad \overline{M} = \varnothing$ otherwise

5.3 Example: Best-Effort Evidence Revisited

We revisit the example of Section 4. In the system without automatic theorem prover, it is not enough to perform the operational checks that grant authorization. It is also necessary to provide the logical engine with hints on how to derive the right authorization facts.

For example, a reviewer v for paper id that decides to appoint a sub-reviewer u, needs to tell the server how to derive from the policy the authorization fact $Reviewer(u,id)$, based on the facts that may be available by the time the request is submitted. In particular, the hint H in the verifier code $dver$ contains the facts $Delegate(v,u,id)$, stated by v itself, $Reviewer(v,id)$ which v cannot state, but that it can assume to be asserted by the time the delegation request is filed, and the rule needed to conclude $Reviewer(u,id)$. The (simpler) case for filing reviews is given in Appendix B.

$H = Reviewer(U,ID) :- Reviewer(V,ID), Delegate(V,U,ID); Reviewer(v,id); Delegate(v,u,id)$
$dver = (return)(Delegate(v,u,id) \mid \mathsf{out}\ return(\mathsf{ok}(H))$

The server code needs to change the effects obtained by verifying a delegation request, essentially stating a lemma useful to prove further authorization.

$S \mid PCMember(alice) \mid Reviewer(bob,id) \mid ...$
$\quad \mid (!\mathsf{in}\ latedelegation(v,u,id,dver);\ \mathsf{verify}\ dver\langle\rangle{:}Reviewer(u,id); Reviewer(u,id))$

6 Conclusions

In this paper, we introduce "Code-Carrying Authorization" as a discipline for passing fragments of a reference monitor rather than proofs in order to perform run-time authorization. These fragments are themselves checked dynamically, since in general they are not trusted. We present a typing discipline that statically enforces safety with respect to authorization logics, and explore the notion of passing (proof) hints as a way to alleviate the dynamic verification process. The recent literature contains other type systems for authorization policies. While we base our work on that of Fournet et al. [12], because of its simplicity, the ideas that we explore should carry over to more elaborate languages. In particular, these variants would address the problem of partial trust [13]. They may also enable us to instantiate CCA in a general-purpose programming language such as F# [6] (a dialect of ML). Going beyond the present exploration (in which we emphasize concepts and theory over practice), such extensions are important for the further study of CCA and its applications.

Acknowledgments. We thank Gordon Plotkin for useful comments and suggestions. Sergio Maffeis is supported by EPSRC grant EP/E044956/1. This work was done while Maffeis was visiting Microsoft Research, Silicon Valley, whose hospitality is gratefully acknowledged.

References

1. M. Abadi. Access control in a core calculus of dependency. In *Computation, Meaning, and Logic: Articles dedicated to Gordon Plotkin*, pages 5–31. Elsevier, 2007. Volume 172 of ENTCS.
2. M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically-typed language. In *POPL'89: Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages*, pages 213–227. ACM, 1989.
3. M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Inf. and Comp.*, 148:1–70, 1999.
4. A. W. Appel and E. W. Felten. Proof-carrying authentication. In *CCS'99: Proceedings of the 6th ACM Conference on Computer and Communications Security*, pages 52–62, 1999.
5. L. Bauer, M. A. Schneider, and E. W. Felten. A general and flexible access-control system for the Web. In *Proceedings of the 11th USENIX Security Symposium*, pages 93–108, 2002.
6. J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffeis. Refinement types for secure implementations. In *21st IEEE Computer Security Foundations Symposium (CSF'08)*, pages 17–32. IEEE, June 2008.
7. S. Ceri, G. Gottlob, and L. Tanca. What you always wanted to know about Datalog (and never dared to ask). *IEEE Trans. Knowl. Data Eng.*, 1(1):146–166, 1989.
8. B.-Y. E. Chang, A. J. Chlipala, G. C. Necula, and R. R. Schneck. The Open Verifier framework for foundational verifiers. In *ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI 2005)*, pages 1–12. ACM, 2005.
9. A. Cirillo, R. Jagadeesan, C. Pitcher, and J. Riely. Do As I SaY! Programmatic access control with explicit identities. In *CSF'07: 20th IEEE Computer Security Foundation Symposium*, pages 16–30. IEEE, 2007.
10. A. Cirillo and J. Riely. Access control based on code identity for open distributed systems. In *TGC'07: Trustworthy Global Computing*, volume 4912 of *Lecture Notes in Computer Science*, pages 169–185. Springer, 2007.
11. J. DeTreville. Binder, a logic-based security language. In *IEEE Computer Society Symposium on Research in Security and Privacy*, pages 105–113. IEEE, 2002.
12. C. Fournet, A. D. Gordon, and S. Maffeis. A type discipline for authorization policies. *ACM Trans. Program. Lang. Syst.*, 29(5):25, 2007.
13. C. Fournet, A. D. Gordon, and S. Maffeis. A type discipline for authorization policies in distributed systems. In *CSF'07: 20th IEEE Computer Security Foundation Symposium*, pages 31–45. IEEE, 2007.
14. M. Hennessy, J. Rathke, and N. Yoshida. safeDpi: a language for controlling mobile code. *Acta Inf.*, 42(4-5):227–290, 2005.
15. C. Lesniewski-Laas, B. Ford, J. Strauss, R. Morris, and M. F. Kaashoek. Alpaca: extensible authorization for distributed services. In *CCS '07: Proceedings of the 14th ACM conference on Computer and communications security*, pages 432–444. ACM, 2007.
16. G. C. Necula. Proof-carrying code. In *POPL '97: Proceedings of the 24th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 106–119. ACM, 1997.
17. J. Riely and M. Hennessy. Trust and partial typing in open systems of mobile agents. *J. Autom. Reas.*, 31(3-4):335–370, 2003.
18. D. Sangiorgi. From pi-calculus to higher-order pi-calculus - and back. In *TAPSOFT'93: Theory and Practice of Software Development*, pages 151–166, 1993.
19. J. A. Vaughan, L. Jia, K. Mazurak, , and S. Zdancewic. Evidence-based audit. In *21st IEEE Computer Security Foundations Symposium (CSF'08)*, pages 163–176. IEEE, June 2008.

A Tables

Authorization Logic: $(\mathscr{C}, fn, \models)$

An *authorization logic* $(\mathscr{C}, fn, \models)$ is a set of clauses $C \in \mathscr{C}$ closed by substitutions σ of messages for names, with finite sets of *free names* $fn(C)$ such that $C\sigma = C$ if $dom(\sigma) \cap fn(C) = \varnothing$ and $fn(C\sigma) \subseteq (fn(C) \setminus dom(\sigma)) \cup fn(\sigma)$; and with an *entailment relation* $S \models C$, between sets of clauses $S \subseteq \mathscr{C}$ and clauses $C, C' \in \mathscr{C}$, such that *(Mon)* $S \models C \Rightarrow S \cup \{C'\} \models C$ and *(Subst)* $S \models C \Rightarrow S\sigma \models C\sigma$.

Free Names: $fn(M), fn(T)$

$fn(x) = \{x\} \quad fn(\{M\}N) = fn((M,N)) = fn(M) \cup fn(N) \quad fn(\mathsf{ok}\,H) = fn(H)$
$fn((x{:}T)P) = fn(T) \cup (fn(P) \setminus \{x\})$

$fn(\mathsf{Un}) = \varnothing \quad fn(\mathsf{Key}(T)) = fn(\mathsf{Ch}(T)) = fn(T) \quad fn(\mathsf{Pr}(T)) = fn(T)$

$fn((x{:}T,U) = fn(T) \cup (fn(U) \setminus \{x\}) \quad fn(\mathsf{Ok}(S)) = fn(S)$

Rules for Structural Equivalence: $P \equiv Q$

$P \equiv P$	(Struct Refl)
$Q \equiv P \Rightarrow P \equiv Q$	(Struct Symm)
$P \equiv Q, Q \equiv R \Rightarrow P \equiv R$	(Struct Trans)
$P \equiv P' \Rightarrow \mathsf{new}\ x{:}T; P \equiv \mathsf{new}\ x{:}T; P'$	(Struct Res)
$P \equiv P' \Rightarrow P \mid R \equiv P' \mid R$	(Struct Par)
$P \mid \mathbf{0} \equiv P$	(Struct Par Zero)
$P \mid Q \equiv Q \mid P$	(Struct Par Comm)
$(P \mid Q) \mid R \equiv P \mid (Q \mid R)$	(Struct Par Assoc)
$\mathsf{new}\ x{:}T; (P \mid Q) \equiv P \mid \mathsf{new}\ x{:}T; Q$	(Struct Res Par) (for $x \notin fn(P)$)
$\mathsf{new}\ x_1{:}T_1; \mathsf{new}\ x_2{:}T_2; P \equiv$ $\mathsf{new}\ x_2{:}T_2; \mathsf{new}\ x_1{:}T_1; P$	(Struct Res Res) (for $x_1 \neq x_2, x_1 \notin fn(T_2), x_2 \notin fn(T_1)$)

Syntactic Sugar: Input, Decryption and Pattern-Matching

$\mathsf{in}\ M(\widetilde{M}); P = \mathsf{in}\ M(y{:}Ty_C(M)); \mathsf{tuple}\ y\ \mathsf{as}\ (\widetilde{M}); P$ (S Input)
(where $y \notin fn(\widetilde{M}) \cup fn(P)$)

$\mathsf{decrypt}\ M\ \mathsf{as}\ \{\widetilde{N}\}N; P = \mathsf{decrypt}\ M\ \mathsf{as}\ \{y{:}Ty_K(N)\}N; \mathsf{tuple}\ y\ \mathsf{as}\ (\widetilde{N}); P$ (S Decrypt)
(where $y \notin fn(\widetilde{M}) \cup fn(P)$)

$\mathsf{tuple}\ M\ \mathsf{as}\ (z, \widetilde{M}); P = \mathsf{split}\ M\ \mathsf{as}\ (z{:}Ty_L(M), y{:}Ty_R(M)); \mathsf{tuple}\ y\ \mathsf{as}\ (\widetilde{M}); P$ (S Split)
(where $y \notin fn(\widetilde{M}) \cup fn(P) \cup \{z\}$)

$\mathsf{tuple}\ M\ \mathsf{as}\ (z); P = \mathsf{split}\ (M,M)\ \mathsf{as}\ (z{:}Ty(M), y{:}Ty(M)); P$ (S Split 0)
(where $y \notin fn(P) \cup \{z\}$)

$\mathsf{tuple}\ M\ \mathsf{as}\ (=N, \widetilde{N}); P = \mathsf{match}\ M\ \mathsf{as}\ (N, y{:}Ty_R(M)); \mathsf{tuple}\ y\ \mathsf{as}\ (\widetilde{N}); P$ (S Match)
(where $y \notin fn(\widetilde{M}) \cup fn(P)$)

$\mathsf{tuple}\ M\ \mathsf{as}\ (=N); P = \mathsf{match}\ (M,M)\ \mathsf{as}\ (N, y{:}Ty(M)); P$ (S Match 0)
(where $y \notin fn(P)$)

When an environment E is fixed, the macro $Ty_{[C/K/L/R]}(M)$ can be translated to T
if $E \vdash M : T'$ where T' is respectively $T, \mathsf{Ch}(T), \mathsf{Key}(T), (x : T, U)$ or $(x : U, T)$.

Substitution for Types:

$\mathsf{Un}\sigma = \mathsf{Un} \quad \mathsf{Key}(T)\sigma = \mathsf{Key}(T\sigma) \quad (x{:}T,U)\{M/x\} = (x{:}T\{M/x\}, U)$
$(x{:}T,U)\{M/y\} = (x{:}T\{M/y\}, U\{M/y\}) \quad (\text{if } x \neq y)$
$\mathsf{Ch}(T)\sigma = \mathsf{Ch}(T\sigma) \quad \mathsf{Ok}(S)\sigma = \mathsf{Ok}(S\sigma) \quad \mathsf{Pr}(T)\sigma = \mathsf{Pr}(T\sigma)$

Rules for Environments: $E \vdash \diamond$

(Env $\varnothing$) $\dfrac{}{\varnothing \vdash \diamond}$

(Env x) $\dfrac{E \vdash \diamond \quad fn(T) \subseteq dom(E) \quad x \notin dom(E)}{E, x{:}T \vdash \diamond}$

(Env C) $\dfrac{E \vdash \diamond \quad fn(C) \subseteq dom(E)}{E, C \vdash \diamond}$

Additional Typing Rules for Processes: $E \vdash P$

(Proc Match) $\dfrac{E \vdash M : (x{:}T, U) \quad E \vdash N : T \quad E, y{:}U\{N/x\} \vdash P}{E \vdash \mathsf{match}\ M\ \mathsf{as}\ (N, y{:}U\{N/x\}); P}$

(Proc Match Un) $\dfrac{E \vdash M : \mathsf{Un} \quad E \vdash N : \mathsf{Un} \quad E, y{:}\mathsf{Un} \vdash P}{E \vdash \mathsf{match}\ M\ \mathsf{as}\ (N, y{:}\mathsf{Un}); P}$

(Proc Split) $\dfrac{E \vdash M : (x{:}T, U) \quad E, x{:}T, y{:}U \vdash P}{E \vdash \mathsf{split}\ M\ \mathsf{as}\ (x{:}T, y{:}U); P}$

(Proc Split Un) $\dfrac{E \vdash M : \mathsf{Un} \quad E, x{:}\mathsf{Un}, y{:}\mathsf{Un} \vdash P}{E \vdash \mathsf{split}\ M\ \mathsf{as}\ (x{:}\mathsf{Un}, y{:}\mathsf{Un}); P}$

Derived Syntax and Derived Typing Rule for Verification: $\mathsf{verify}\ M\langle \widetilde{N{:}T} \rangle{:}x{:}S; P$

$$\mathsf{verify}\ M\langle [\widetilde{z{:}\mathsf{Un}}, N{:}T] \rangle{:}x{:}S; P \triangleq \mathsf{new}\ c{:}\mathsf{Ch}(\mathsf{Ok}(S)); \big(\mathsf{typecase}\ M\ \mathsf{of}\ y{:}\mathsf{Pr}([\widetilde{z{:}\mathsf{Un}}, T,]\mathsf{Ch}(\mathsf{Ok}(S))); \mathsf{spawn}\ y\ \mathsf{with}\ ([\widetilde{z}, N,]c) \mid \mathsf{in}\ c(x{:}\mathsf{Ok}(S)); P\big)$$
$$(\{c, y, x\} \cap fn(P, M, [N,]S) = \varnothing,\ \{\widetilde{z}\} \subseteq fn(S),\ x \notin fn(M[, N]))$$

(Proc Verify) $\dfrac{E \vdash M : U \quad [E \vdash \widetilde{N : T}] \quad E, x{:}\mathsf{Ok}(S) \vdash P}{E \vdash \mathsf{verify}\ M\langle [\widetilde{N{:}T}] \rangle{:}x{:}S; P}$

B Code from Examples

The complete code for the best-effort-evidence example, with hints, is reported below.

```
Delegate(v:Un) =
!in phonereviewrequest(=v,id);
  (in accept(r); out filereview(v,id,r,fver))
|(in delegate(u); out phonereviewrequest(u,id)
                  | out latedelegation(v,u,id,dver))

fver = (return)(Opinion(v,id,r)|out return(ok(HF)))
dver = (return)(Delegate(v,u,id)|out return(ok(HD))

HD = (Reviewer(U,ID) :- Reviewer(V,ID),Delegate(V,U,ID)),
     Reviewer(v,id),
     Delegate(v,u,id)
HF = (Review(U,ID,R) :- Reviewer(U,ID),Opinion(U,ID,R)),
     Reviewer(v,id),
     Opinion(v,id,r)

S | PCMember(alice) | Reviewer(bob,id)
  | (!in filedreview(v,id,r,fver); verify fver⟨⟩:Review(v,id,r); [...])
  | (!in latedelegation(v,u,id,dver); verify dver⟨⟩:Reviewer(u,id);Reviewer(u,id))
```

C Proofs

In order to prove the main theorems in Section 3 and Section 5, we first prove subject reduction, safety and robust safety for a reference type system obtained by replacing rule (Msg Hint) with the rule below.

Typing Rule for the Reference Type System

(Msg Hint Gen)
$$\frac{E,S \vdash \diamond \quad fn(H) \subseteq dom(E) \quad clauses(E) \models C \quad \forall C \in S}{E \vdash \mathsf{ok}\, H : \mathsf{Ok}(S)}$$

We then show that the type system of Section 3 coincides with the reference type system for the sub-language where each hint is the empty set of clauses, and each expectation has the token ok as a parameter. Finally, we show that for each sound verification predicate, rule (Proc Expect Hint) is derivable in the reference type system, hence safety and robust safety (but not necessarily subject reduction) follow for the type system of Section 5.

C.1 Results for the Reference Typing System

We proceed to show the main properties of the type system, in particular subject congruence and subject reduction, which together give type preservation (Lemma 1).

Before proving subject congruence and subject reduction in detail, we state without proof a series of fairly standard technical properties of the type system. In the rest of this section, let $\mathscr{J} \in \{\diamond, M : T, P\}$.

Lemma 3 (Well-formedness). *If $E \vdash \mathscr{J}$ then $E \vdash \diamond$ and $fv(\mathscr{J}) \subseteq dom(E)$.*

Proof By induction on the derivation of $E \vdash \mathscr{J}$. □

Lemma 4 (Exchange). *If $E_1, E_2, E_3, E_4 \vdash \mathscr{J}$ and $dom(E_2) \cap fv(E_3) = \varnothing$ then $E_1, E_3, E_2, E_4 \vdash \mathscr{J}$.*

Proof By induction on the derivation of $E_1, E_2, E_3, E_4 \vdash \mathscr{J}$. □

Lemma 5 (Weakening). *If $E_1, E_2 \vdash \mathscr{J}$, $E \vdash \diamond$ and $fv(E) \subseteq dom(E_1)$ and $dom(E) \cap dom(E_1, E_2) = \varnothing$, then $E_1, E, E_2 \vdash \mathscr{J}$.*

Proof By induction on the derivation of $E_1, E_2 \vdash \mathscr{J}$, using property (Mon) of the authorization logic.

□

Lemma 6 (Strengthening). *(i) If $E, x{:}T, E' \vdash \mathscr{J}$ and T is generative and $x \notin fn(\mathscr{J}) \cup fn(E')$ then $E, E' \vdash \mathscr{J}$. (ii) If $E, C, E' \vdash \diamond$ then $E, E' \vdash \diamond$.*

Proof By induction on the derivations of $E, x{:}T, E' \vdash \mathscr{J}$ and $E, C, E' \vdash \diamond$. □

Lemma 7 (Substitution). *If $E_1, x{:}T, E_2 \vdash \mathscr{J}$ and $E_1 \vdash M : T$ then $E_1, E_2\{M/x\} \vdash \mathscr{J}\{M/x\}$.*

Proof By induction on the derivation of $E_1, x{:}T, E_2 \vdash \mathscr{J}$, using property (Subst) of the authorization logic. □

Lemma 8 (Environment Change). *If $E_1, E, E_2 \vdash \mathscr{J}$ and $fv(\mathscr{J}) \cap dom(E) = \varnothing$ and $E_1, E', E_2 \vdash \diamond$ and $dom(E) \subseteq dom(E')$ and $clauses(E) \subseteq clauses(E')$ then $E_1, E', E_2 \vdash \mathscr{J}$.*

Proof By induction on the derivation of $E_1, E, E_2 \vdash \mathscr{J}$. □

Lemma 9 (Normal Form). *If $E \vdash P$ and $clauses(env(P)^{\widetilde{x}}) = S$ then $E, env(P)^{\widetilde{x}} \vdash \diamond$ and there exists $\widetilde{T}, P'$ such that $P \equiv \mathsf{new}\ \widetilde{x}{:}\widetilde{T}; (S \mid P')$.*

Proof By induction on the derivation of $E \vdash P$. □

Lemma 10 (Opponent Typability). *Let $\mathscr{J} \in \{M{:}\mathsf{Un}, P\}$. If $fn(\mathscr{J}) = \{\widetilde{x}\}$ and $\mathscr{J}$ does not contain expectations, then $\widetilde{x}{:}\mathsf{Un} \vdash \mathscr{J}$.*

Proof By induction on the derivation of $\widetilde{x}{:}\mathsf{Un} \vdash \mathscr{J}$. □

Proof of Lemma 2. *For any opponent O, $\widetilde{x}{:}\widetilde{\mathsf{Un}} \vdash O$, where $fn(O) \subseteq \{\widetilde{x}\}$.*

Proof Follows directly from Lemma 10. □

Lemma 11 (Subject Congruence). *If $E \vdash P$ and $P \equiv P'$ then $E \vdash P'$ and $clauses(env(P)) = clauses(env(P'))$.*

Proof Let $\rho(\widetilde{a})$ be a permutation of $\widetilde{a}$.
By induction on the derivation of $P \equiv P'$ we show:

(1) if $E \vdash P$ then $E \vdash P'$ and $clauses(env(P)^{\widetilde{a}}) = clauses(env(P')^{\rho(\widetilde{a})})$;
(2) if $E \vdash P'$ then $E \vdash P$ and $clauses(env(P')^{\rho(\widetilde{a})}) = clauses(env(P)^{\widetilde{a}})$.

(Struct Refl) Suppose $P \equiv P$.
Both (1) and (2) are immediate.
(Struct Symm) Suppose $P \equiv Q$.
By hypothesis, $Q \equiv P$.
Both (1) and (2) follow immediately applying the inductive hypotheses (2) and (1).
(Struct Trans) Suppose $P \equiv R$.
By hypothesis, $P \equiv Q, Q \equiv R$.
Both cases follow easily from transitivity of implication and the inductive hypotheses.
(Struct Res) Suppose $\mathsf{new}\ a{:}T; P \equiv \mathsf{new}\ a{:}T; P'$.
By hypothesis, $P \equiv P'$.
By hypothesis of (1), $E \vdash \mathsf{new}\ a{:}T; P$ and
$clauses(env(\mathsf{new}\ a{:}T; P)^{a,\widetilde{a}}) = clauses(env(P)^{\widetilde{a}})$.
By (Proc Res), $E, a{:}T \vdash P$.
By inductive hypothesis, $E, a{:}T \vdash P'$ and $clauses(env(P)^{\widetilde{a}}) = clauses(env(P')^{\rho(\widetilde{a})})$.
By (Proc Res), $E \vdash \mathsf{new}\ a{:}T; P'$.
By definition, $clauses(env(\mathsf{new}\ a{:}T; P')^{a,\rho(\widetilde{a})}) = clauses(env(P')^{\rho(\widetilde{a})})$ and
$clauses(env(P)^{\widetilde{a}}) = clauses(env(\mathsf{new}\ a{:}T; P)^{a,\widetilde{a}})$.
The proof for (2) is symmetric.

(Struct Par) Suppose $P \mid Q \equiv P' \mid Q$.
By hypothesis, $P \equiv P'$.
By hypothesis of (1), $E \vdash P \mid Q$.
By (Proc Par), $E, env(Q)^{\widetilde{c}} \vdash P$ and $E, env(P)^{\widetilde{a}} \vdash Q$ and $fv(P \mid Q) \subseteq dom(E)$.
By inductive hypothesis, $E, env(Q)^{\widetilde{c}} \vdash P'$ and
$clauses(env(P)^{\widetilde{a}}) = clauses(env(P')^{\rho(\widetilde{a})})$.
By Lemma 9, $E, env(P')^{\rho(\widetilde{a})} \vdash \diamond$.
By definition of *env*, $dom(env(P)^{\widetilde{a}}) = dom(env(P')^{\rho(\widetilde{a})}) = \{\widetilde{a}\}$.
By Lemma 8, $E, env(P')^{\rho(\widetilde{a})} \vdash Q$.
By (Proc Par), $E \vdash P' \mid Q$.
By definition of *env*, $clauses(env(P \mid Q)^{\widetilde{a},\widetilde{c}}) = clauses(env(P' \mid Q)^{\rho(\widetilde{a}),\widetilde{c}})$.
The proof for (2) is symmetric.

(Struct Par Zero) Suppose $P \mid \mathbf{0} \equiv P$.
By hypothesis of (1), $E \vdash P \mid \mathbf{0}$.
By (Proc Par), $E, env(\mathbf{0}) \vdash P$ and $E, env(P) \vdash \mathbf{0}$ and $fv(P \mid \mathbf{0}) \subseteq dom(E)$.
By definition of *env*, $env(\mathbf{0}) = \varnothing$, hence $E \vdash P$.
By definition of *clauses*, $clauses(env(\mathbf{0})) = \varnothing$, hence
$clauses(env(P \mid \mathbf{0})) = clauses(env(P))$.
The proof for (2) is similar.

(Struct Par Comm) Suppose $P \mid Q \equiv Q \mid P$.
By hypothesis of (1), $E \vdash P \mid Q$.
By (Proc Par), $E, env(Q)^{\widetilde{c}} \vdash P$ and $E, env(P)^{\widetilde{a}} \vdash Q$ and $fv(P \mid Q) \subseteq dom(E)$.
By (Proc Par), $E \vdash Q \mid P$.
By definition of *clauses* and *env*, $clauses(env(P \mid Q)^{\widetilde{a},\widetilde{c}}) = clauses(env(Q \mid P)^{\widetilde{c},\widetilde{a}})$.
The proof for (2) is symmetric.

(Struct Par Assoc) Suppose $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$.
By hypothesis of (1), $E \vdash (P \mid Q) \mid R$.
By (Proc Par), $E, env(R)^{\widetilde{c}} \vdash P \mid Q, E, env(P \mid Q)^{\widetilde{a},\widetilde{b}} \vdash R$ and $fv((P \mid Q) \mid R) \subseteq dom(E)$.
By (Proc Par), $E, env(R)^{\widetilde{c}}, env(Q)^{\widetilde{b}} \vdash P$, $E, env(R)^{\widetilde{c}}, env(P)^{\widetilde{a}} \vdash Q$ and
$fv(P \mid Q) \subseteq dom(E, env(R))$.
By Lemma 4, $E, env(Q)^{\widetilde{b}}, env(R)^{\widetilde{c}} \vdash P$ and $E, env(P)^{\widetilde{a}}, env(R)^{\widetilde{c}} \vdash Q$.
By (Proc Par), $E, env(P)^{\widetilde{a}} \vdash Q \mid R$.
By (Proc Par), $E \vdash P \mid (Q \mid R)$.
By definition of *clauses* and *env*, $clauses(env((P \mid Q) \mid R)^{\widetilde{a},\widetilde{b},\widetilde{c}}) = clauses(env(P \mid (Q \mid R))^{\widetilde{a},\widetilde{b},\widetilde{c}})$.
The proof for (2) is similar.

(Struct Res Par) Suppose $\mathsf{new}\ a{:}T; (P \mid Q) \equiv P \mid \mathsf{new}\ a{:}T; Q$.
By hypothesis, $a \notin fn(P)$.
By hypothesis of (1), $E \vdash \mathsf{new}\ a{:}T; (P \mid Q)$ and
$env(\mathsf{new}\ a{:}T; (P \mid Q)) = a{:}T, env(P)^{\widetilde{a}}, env(Q)^{\widetilde{c}}$.
By (Proc Res), $E, x{:}T \vdash P \mid Q$.
By (Proc Par), $E, a{:}T, env(Q)^{\widetilde{c}} \vdash P$ and $E, a{:}T, env(P)^{\widetilde{a}} \vdash Q$.
By (Proc Res), $E, env(P)^{\widetilde{a}} \vdash \mathsf{new}\ a{:}T; Q$.
Since $a \notin fn(P)$, by Lemma 6, $E, env(Q)^{\widetilde{c}} \vdash P$.
By (Proc Par), $E \vdash P \mid \mathsf{new}\ a{:}T; Q$.
By definition of *env*, $env(P \mid \mathsf{new}\ a{:}T; Q)^{\widetilde{a},a,\widetilde{c}} = env(P)^{\widetilde{a}}, a{:}T, env(Q)^{\widetilde{c}}$.
By definition of *clauses*,
$clauses(env(\mathsf{new}\ a{:}T; (P \mid Q))^{a,\widetilde{a},\widetilde{c}}) = clauses(env(P \mid \mathsf{new}\ a{:}T; Q)^{\widetilde{a},a,\widetilde{c}})$.
The proof for (2) is similar, using Lemma 5 instead of Lemma 6.

(Struct Res Res) Suppose $\mathsf{new}\ a_1{:}T_1;\mathsf{new}\ a_2{:}T_2;P \equiv \mathsf{new}\ a_2{:}T_2;\mathsf{new}\ a_1{:}T_1;P$.
By hypothesis, $a_1 \neq a_2, x_1 \notin fn(T_2), a_2 \notin fn(T_1)$.
By (Proc Res), $E, a_1{:}T_1 \vdash \mathsf{new}\ a_2{:}T_2;P$.
By (Proc Res), $E, a_1{:}T_1, a_2{:}T_2 \vdash P$.
Since $a_1 \neq a_2, a_1 \notin fn(T_2), a_2 \notin fn(T_1)$, by Lemma 4, $E, a_2{:}T_2, a_1{:}T_1 \vdash P$.
By two applications of (Proc Res), $E \vdash \mathsf{new}\ a_2{:}T_2;, \mathsf{new}\ a_1{:}T_1;P$.
By definition of *clauses*, $clauses(env(\mathsf{new}\ a_1{:}T_1;\mathsf{new}\ a_2{:}T_2;P)^{a_1,a_2,\widetilde{a}}) = clauses(env(\mathsf{new}\ a_2{:}T_2;\mathsf{new}\ a_1{:}T_1;P)^{a_2,a_1,\widetilde{a}}) = clauses(env(P)^{\widetilde{a}})$.
The proof for (2) is symmetric. □

Lemma 12 (Subject Reduction). *If $E \vdash P$ and $P \rightarrow_E P'$ then $E \vdash P'$ and $clauses(env(P)^{\widetilde{a}}) \subseteq clauses(env(P')^{\widetilde{c}})$ and $\{\widetilde{a}\} \subseteq \{\widetilde{c}\}$.*

Proof The proof is by induction on the derivation of $P \rightarrow_E P'$.

(Red Par) Suppose $P \mid Q \rightarrow_E P' \mid Q$.
By hypothesis, $P \rightarrow_{E,env(Q)} P'$.
By hypothesis of the lemma, $E \vdash P \mid Q$.
This must follow from applying (Proc Par), with the premises $E, env(Q)^{\widetilde{c}} \vdash P$, $E, env(P)^{\widetilde{a}} \vdash Q$ and $fn(P \mid Q) \subseteq dom(E)$.
By inductive hypothesis, $E, env(Q)^{\widetilde{c}} \vdash P'$ and $clauses(env(P)^{\widetilde{a}}) \subseteq clauses(env(P')^{\widetilde{b}})$ and $\{\widetilde{a}\} \subseteq \{\widetilde{b}\}$.
By Lemma 9, $E, env(P')^{\widetilde{b}} \vdash \diamond$.
By definition of *env*, $dom(env(P)^{\widetilde{a}}) = \{\widetilde{a}\}$ and $dom(env(P')^{\widetilde{b}}) = \{\widetilde{b}\}$.
By Lemma 8, $E, env(P')^{\widetilde{b}} \vdash Q$.
By (Proc Par), $E \vdash P' \mid Q$.
By definition, $clauses(env(P \mid Q)^{\widetilde{a},\widetilde{c}}) \subseteq clauses(env(P' \mid Q)^{\widetilde{b},\widetilde{c}})$ and $\{\widetilde{a},\widetilde{c}\} \subseteq \{\widetilde{b},\widetilde{c}\}$.

(Red Res) Suppose $\mathsf{new}\ a{:}T;P \rightarrow_E \mathsf{new}\ a{:}T;P'$.
By hypothesis, $P \rightarrow_{a:T} P'$.
By hypothesis of the lemma, $E \vdash \mathsf{new}\ a{:}T;P$.
This must follow from applying (Proc Res), with the premise $E, a{:}T \vdash P$.
By inductive hypothesis, $E, a{:}T \vdash P'$ and $clauses(env(P)^{\widetilde{a}}) \subseteq clauses(env(P')^{\widetilde{c}})$ and $\{\widetilde{a}\} \subseteq \{\widetilde{c}\}$.
By (Proc Res), $E \vdash \mathsf{new}\ a{:}T;P'$.
By definition, $clauses(env(\mathsf{new}\ a{:}T;P)^{a,\widetilde{a}}) \subseteq clauses(env(\mathsf{new}\ a{:}T;P')^{a,\widetilde{c}})$ and $\{a,\widetilde{a}\} \subseteq \{a,\widetilde{c}\}$.

(Red Struct) Suppose $P \rightarrow_E P'$.
By hypothesis, $P \equiv Q, Q \rightarrow_E Q', Q' \equiv P'$.
By Lemma 11 on $E \vdash P$, $E \vdash Q$ and $clauses(env(P)^{\widetilde{a}}) = clauses(env(Q)^{\rho_1(\widetilde{a})})$.
By inductive hypothesis on $E \vdash Q$, $E \vdash Q'$ and $clauses(env(Q)^{\rho_1(\widetilde{a})}) \subseteq clauses(env(Q')^{\widetilde{b}})$ and $\{\rho(\widetilde{a})\} \subseteq \{\widetilde{b}\}$, hence $\{\widetilde{a}\} \subseteq \{\widetilde{b}\}$.
By Lemma 11 on $E \vdash Q'$, $E \vdash P'$ and $clauses(env(Q')^{\widetilde{b}}) = clauses(env(P')^{\rho_2(\widetilde{b})})$.
By transitivity, $clauses(env(P)^{\widetilde{a}}) \subseteq clauses(env(P')^{\rho_2(\widetilde{b})})$ and $\{\widetilde{a}\} \subseteq \{\rho_2(\widetilde{b})\}$.

(Red Comm) Suppose $\mathsf{out}\ a(M) \mid \mathsf{in}\ a(x{:}T);P \rightarrow_E P\{M/x\}$.
By hypothesis of the lemma, $E \vdash \mathsf{out}\ a(M) \mid \mathsf{in}\ a(x{:}T);P$.
This must follow from applying (Proc Par), with the premises $E \vdash \mathsf{out}\ a(M)$, $E \vdash \mathsf{in}\ a(x{:}T);P$ and $fn(\mathsf{out}\ a(M) \mid \mathsf{in}\ a(x{:}T);P) \subseteq dom(E)$, where we used the fact that $env(\mathsf{out}\ a(M)) = env(\mathsf{in}\ a(x{:}T);P) = \varnothing$).
We split the proof in two cases, depending on the hypothesis used to derive $E \vdash \mathsf{in}\ a(x{:}T);P$.

– Suppose $E \vdash \mathsf{in}\ a(x{:}T);P$ follows from applying (Proc Input) with the premises $E \vdash a : \mathsf{Ch}(T)$ and $E, x{:}T \vdash P$.
It must be the cases that $E \vdash \mathsf{out}\ a(M)$ follows from applying (Proc Output) with the premises $E \vdash a : \mathsf{Ch}(T)$ and $E \vdash M : T$.
By Lemma 7, $E \vdash P\{M/x\}$.
By definition, $\mathit{clauses}(\mathit{env}(\mathsf{out}\ a(M);Q \mid \mathsf{in}\ a(x);P)) = \varnothing$, and $\varnothing \subseteq \mathit{clauses}(\mathit{env}(Q \mid P\{M/x\}))$.
– Suppose $E \vdash \mathsf{in}\ a(x{:}T);P$ follows from applying (Proc Input Un) with the premises $E \vdash a : \mathsf{Un}$ and $E, x{:}T \vdash P$, where $T = \mathsf{Un}$.
It must be the cases that $E \vdash \mathsf{out}\ a(M)$ follows from applying (Proc Output Un) with the premises $E \vdash a : \mathsf{Un}$ and $E \vdash M : \mathsf{Un}$.
By Lemma 7, $E \vdash P\{M/x\}$.
By definition, $\mathit{clauses}(\mathit{env}(\mathsf{out}\ a(M);Q \mid \mathsf{in}\ a(x);P)) = \varnothing$, and $\varnothing \subseteq \mathit{clauses}(\mathit{env}(Q \mid P\{M/x\}))$.

(Red !Comm) Similar to the previous case.

(Red Decrypt) Suppose $\mathsf{decrypt}\ \{M\}k\ \mathsf{as}\ \{y{:}T\}k;P \to P\{M/y\}$.
If $E \vdash \mathsf{decrypt}\ \{M\}k\ \mathsf{as}\ \{y{:}T\}k;P$ is derived by (Proc Decrypt) then $E \vdash M : T$, $E \vdash k : \mathsf{Key}(T)$, and $E, y{:}T \vdash P$.
By Lemma 7, $E \vdash P\{M/y\}$.
Note that $\mathit{env}(\mathsf{decrypt}\ \{M\}k\ \mathsf{as}\ \{y{:}T\}k;P) = \varnothing$.
The case for rule (Proc Decrypt Un) is similar.

(Red Split) Suppose $\mathsf{split}\ (M,N)\ \mathsf{as}\ (x{:}T, y{:}U);P \to P\{M/x\}\{N/y\}$.
If $E \vdash \mathsf{split}\ (M,N)\ \mathsf{as}\ (x{:}T, y{:}U);P$ is derived by (Proc Split) then $E \vdash (M,N) : (x{:}T, U)$ and $E, x{:}T, y{:}U \vdash P$.
By (Msg Pair), $E \vdash M : T$ and $E \vdash N : U\{M/x\}$.
By Lemma 7, $E, y{:}U\{M/x\} \vdash P\{M/x\}$.
By Lemma 7, $E \vdash P\{M/x\}\{N/y\}$.
Note that $\mathit{env}(\mathsf{split}\ (M,N)\ \mathsf{as}\ (x{:}T, y{:}U);P) = \varnothing$.
The case for rule (Proc Split Un) is similar.

(Red Match) Suppose $\mathsf{match}\ (M,N)\ \mathsf{as}\ (M, y{:}U);P \to P\{N/y\}$.
If $E \vdash \mathsf{match}\ (M,N)\ \mathsf{as}\ (M, y{:}U);P$ is derived by (Proc Match) then $E \vdash (M,N) : (x{:}T, U)$, $E \vdash M : T$ and $E, y{:}U\{M/x\} \vdash P$.
By (Msg Pair), $E \vdash N : U\{M/x\}$.
By Lemma 7, $E \vdash P\{N/y\}$.
Note that $\mathit{env}(\mathsf{match}\ (M,N)\ \mathsf{as}\ (M, y{:}U);P) = \varnothing$.
The case for rule (Proc Match Un) is similar.

(Red Spawn) Suppose $\mathsf{spawn}\ (x{:}T)P\ \mathsf{with}\ M \to_E P\{M/x\}$.
We split the proof in two cases, depending on the hypothesis used to derive $E \vdash \mathsf{spawn}\ (x{:}T)P\ \mathsf{with}\ M$.

(Proc Spawn) The judgment must follow from the premises $E \vdash (x{:}T)P : \mathsf{Pr}(T)$ and $E \vdash M : T$.
It must be the case that $E \vdash (x{:}T)P : \mathsf{Pr}(T)$ follows by (Msg Proc), from the premise $E, x{:}T \vdash P$.

(Proc Spawn Un) The judgment must follow from the premises $E \vdash (x{:}\mathsf{Un})P : \mathsf{Un}$ and $E \vdash M : \mathsf{Un}$.
It must be the case that $E \vdash (x{:}\mathsf{Un})P : \mathsf{Un}$ follows by (Msg Proc Un), from the premise $E, x{:}\mathsf{Un} \vdash P$.

By Lemma 7, $E \vdash P\{M/x\}$.
By definition, spawn $(x{:}T)P$ with $M = \varnothing$.

(Red Typecase) Suppose typecase M of $y{:}T;P \rightarrow_E P\{M/y\}$ because $E \vdash M : T$.
It must be the case that $E \vdash$ typecase M of $y{:}T;P$ is derived by (Proc Typecase) from the premises $E \vdash M : U$ and $E, y{:}T \vdash P$.
By Lemma 7, $E \vdash P\{M/y\}$.
Note that $env(\text{typecase } M \text{ of } y{:}T;P) = \varnothing$.

□

Lemma 13 (Safety). *If $E \vdash P$ and then P is safe for E.*

Proof We need to show that whenever $P \rightarrow_E^{*\equiv}$ new $\tilde{x}{:}\tilde{T};(\text{expect } C \text{ by } M \mid P')$, we can refactor P' so that $P' \equiv$ new $\tilde{y}{:}\tilde{U};(S \mid P'')$, and $S \cup clauses(E) \models C$, with $\{\tilde{y}\} \cap fn(C) = \varnothing$.
By hypothesis, $E \vdash P$.
By Lemma 12 and Lemma 11, if $P \rightarrow_E^{*\equiv}$ new $\tilde{x}{:}\tilde{T};(\text{expect } C \text{ by } M \mid P')$ then
$E \vdash$ new $\tilde{x}{:}\tilde{T};(\text{expect } C \text{ by } M \mid P')$.
This must follow from repeatedly applying (Proc Res) from the premise
$E, \tilde{x}{:}\tilde{T} \vdash$ expect C by $M \mid P'$.
This must follow from (Proc Par), from the premises $E, \tilde{x}{:}\tilde{T} \vdash P'$ and $E, \tilde{x}{:}\tilde{T}, env(P')^{\tilde{y}} \vdash$ expect C by M, where $fn(\text{expect } C \text{ by } M) \subseteq (dom(E) \cup \{\tilde{x}\})$.
By well-formedness of the environment, $\{\tilde{y}\} \cap fn(C) = \varnothing$. This must follow by (Proc Expect Hint), from the premise $E, \tilde{x}{:}\tilde{T}, env(P')^{\tilde{y}} \vdash M : \mathsf{Ok}(S)$, where $C \in S$.
This must follow by (Msg Hint Gen) from the premise $clauses(E, \tilde{x}{:}\tilde{T}, env(P')^{\tilde{y}}) \models C'$, for all $C' \in S$, hence $clauses(E, \tilde{x}{:}\tilde{T}, env(P')^{\tilde{y}}) \models C$.
Remember that since all the $\tilde{x}$ come from nested restrictions, all the $\tilde{T}$ are generative, hence $clauses(\tilde{x}{:}\tilde{T}) = \varnothing$.
Assume, without loss of generality, that $clauses(env(P')^{\tilde{y}}) = S$.
By definition, $S \cup clauses(E) \models C$.
By Lemma 9 on (ii), $P' \equiv$ new $\tilde{y}{:}\tilde{U};(S \mid P'')$. □

Lemma 14 (Robust Safety). *If $\tilde{x}{:}\widetilde{\mathsf{Un}}, S \vdash P$ then P is robustly safe for $\tilde{x}{:}\widetilde{\mathsf{Un}}, S$.*

Proof Consider an arbitrary opponent O, and let $\{\tilde{z}\} = fn(O) \cup \{\tilde{x}\}$.
Let $E = \tilde{z}{:}\widetilde{\mathsf{Un}}, S$.
By Lemma 10, and Lemma 5, $E, env(P) \vdash O$.
By hypothesis $\tilde{x}{:}\widetilde{\mathsf{Un}}, S \vdash P$.
By Lemma 5 and Lemma 4, $E, env(O) \vdash P$.
By (Proc Par), $E \vdash P \mid O$.
By Theorem 1, $P \mid O$ is safe for E. □

C.2 Main Results

In this Section, we let $E \vdash P$ range over typing judgments holding for the rules (and corresponding processes) of the reference type system, or of the type systems of Section 3 and Section 5. The menaing will be clear form the context.

Lemma 15 (Derived Typing Rules).

(1) *If H is defined as the empty set of clauses, and for each* expect C by M *we have* $M =$ ok$\varnothing$*, then (Msg Hint Gen), (Msg Hint Un) and (Proc Expect Hint) are equivalent to (Msg Ok), (Msg Ok Un) and (Proc Expect).*

(2) *If $\mathscr{V}$ is a sound verification predicate, (Msg Hint) is a sound rule in the reference type system.*

Proof

(1) By inspection of the rules.
(2) By (Msg Hint Gen) and by soundness of the verification predicate with respect to logical entailment, which follows by definition of $\mathscr{V}$. □

Proof of Lemma 1 (Type Preservation). *If $E \vdash P$ and $P \rightarrow_E^{*\equiv} P'$ then $E \vdash P'$.*

Proof By Lemmas 11 and 12 noting that the reduction rules trivially preserve the conditions on hints and expectations as of point (1) of Lemma 15.

□

Proof of Theorem 1. *If $E \vdash P$ and then P is safe for E.*

Proof By point (1) of Lemma 15 and Lemma 13. □

Proof of Theorem 2. *If $\widetilde{x}{:}\widetilde{\mathsf{Un}}, S \vdash P$ then P is robustly safe for $\widetilde{x}{:}\widetilde{\mathsf{Un}}, S$.*

Proof By point (1) of Lemma 15 and Lemma 14. □

Proof of Theorem 3. *(i) If $E \vdash P$ and then P is safe for E. (ii) If $\widetilde{x}{:}\widetilde{\mathsf{Un}}, S \vdash P$ then P is robustly safe for $\widetilde{x}{:}\widetilde{\mathsf{Un}}, S$.*

Proof By point (2) of Lemma 15, Lemma 13 and Lemma 14. □