Affine Refinement Types for Secure Distributed Programming

MICHELE BUGLIESI and STEFANO CALZAVARA, Università Ca' Foscari Venezia
FABIENNE EIGNER and MATTEO MAFFEI, CISPA, Saarland University

Recent research has shown that it is possible to leverage general-purpose theorem proving techniques to develop powerful type systems for the verification of a wide range of security properties on application code. Although successful in many respects, these type systems fall short of capturing resource-conscious properties that are crucial in large classes of modern distributed applications. In this paper, we propose the first type system that statically enforces the safety of cryptographic protocol implementations with respect to authorization policies expressed in affine logic. Our type system draws on a novel notion of "exponential serialization" of affine formulas, a general technique to protect affine formulas from the effect of duplication. This technique allows to formulate an expressive logical encoding of the authentication mechanisms underpinning distributed resource-aware authorization policies. We discuss the effectiveness of our approach on two case studies: the EPMO e-commerce protocol and the Kerberos authentication protocol. We finally devise a sound and complete type-checking algorithm, which is the key to achieving an efficient implementation of our analysis technique.

1. INTRODUCTION

Verifying the security of modern distributed applications is an important and complex challenge, which has attracted the interest of a growing research community audience over the last decade. Recent research has shown that it is possible to leverage general-purpose theorem proving techniques to develop powerful type systems for the verification of a wide range of security properties on application code, thus narrowing the gap between the formal model designed for the analysis and the actual implementation of the protocols [Bengtson et al. 2011; Backes et al. 2011; Swamy et al. 2011]. The integration between type systems and theorem proving is achieved by resorting to a form of dependent types, known as *refinement* types. A refinement type $\{x : T \mid F(x)\}$ qualifies the structural information of the type T with a property specified by the logical formula F: a value M of this type is a value of type T such that $F(M)$ holds true.

Authorization systems based on refinement types use the refinement formulas to express (and gain static control of) the credentials associated with the data and the

cryptographic keys involved in the authorization checks. Clearly, the expressiveness of the resulting analysis hinges on the choice of the underlying logic, and indeed several logics have been proposed for the specification and verification of security properties [Chapin et al. 2008]. A number of proposals have thus set logic *parametricity* as a design goal, to gain modularity and scalability of the resulting systems. Though logic parametricity is in principle a sound and wise design choice, current attempts in this direction draw primarily (if not exclusively) on classical (or intuitionistic) logical frameworks. That, in turn, is a choice that makes the resulting systems largely ineffective on large classes of resource-aware authorization policies, such as those based on consumable credentials, or predicating over access counts and/or usage bounds.

The natural choice for expressing and reasoning about such classes of policies are instead *substructural* logics, such as linear and affine logic [Girard 1995; Troelstra 1992]. On the other hand, integrating substructural logics with existing refinement type systems for distributed authorization is challenging, as one must build safeguards against the ability of an attacker to duplicate the data exchanged over the network, and correspondingly duplicate the associated credentials, thus undermining their bounded nature [Bugliesi et al. 2011].

Contributions. In this paper, we present an *affine refinement type system* for RCF [Bengtson et al. 2011], a concurrent λ-calculus which can be directly mapped to a large subset of a real functional programming language like F#. The type system guarantees that well-typed programs comply with any given authorization policy expressed in affine logic, even in the presence of an active opponent.

This type system draws on the novel concept of *exponential serialization*, a general technique to protect affine formulas from the effect of duplication. This technique makes it possible to factor the authorization-relevant invariants of the analysis out of the type system, and to characterize them directly as proof obligations for the underlying affine logical system. This leads to a rather general and modular design of our proposal, and sheds new light on the logical foundations of standard cryptographic patterns underpinning distributed authorization frameworks. Furthermore, the concept of serialization enhances the expressiveness of the type system, capturing programming patterns out of the scope of many substructural type systems.

The clean separation between typing and logical entailment has the additional advantage of enabling the formulation of an algorithmic version of our system, in which the non-deterministic proof search distinctive of substructural type systems can be dispensed with. Intuitively, we can shift all the burden related to substructural resource management into a single proof obligation to be discharged to an external theorem prover. This proof obligation can be efficiently generated from a program in a syntax-directed way: this is the key to achieve a practical implementation of our framework.

We show the effectiveness of our approach on two case studies, namely the *EPMO* e-commerce protocol [Guttman et al. 2004] and the *Kerberos* authentication protocol [Steiner et al. 1988]. For both case studies we discuss the advantages in expressiveness enabled by the adoption of an underlying substructural logic.

Structure of the paper. Section 2 overviews the challenges and the most important aspects of our theory on a simple example. Section 3 reviews intuitionistic affine logic. Section 4 presents the meta-theory of exponential serialization. Section 5 reviews RCF and defines our notion of safety. Section 6 outlines the type system. Section 7 discusses encodings of network communication and our treatment of formal cryptography. Sections 8-9 present the case studies. Section 10 discusses the algorithmic formulation of our type system. Section 11 overviews the related work. Section 12 concludes.

The proofs of the main theorems are provided in the appendixes: Appendix A establishes the soundness of exponential serialization; Appendix B details a soundness

proof for our type system; Appendix C provides proofs of the soundness and completeness of the algorithmic type system. The full proofs of all lemmas and other auxiliary results are contained in the electronic appendix for this article, which can be accessed in the ACM Digital Library.

Unpublished content. The present work extends and revises a conference paper published at POST 2013 [Bugliesi et al. 2013], which received the EATCS award for the best theory paper at ETAPS. In this extended version we present full details of the formalization, including a complete presentation of the type system, its algorithmic variant, and complete soundness proofs for our main results: all this material was not published before, due to space constraints. Moreover, the Kerberos case study in Section 9 is new and required us to define an encoding of "self-dependent key types" in our type system, which we believe to be of independent interest.

2. OVERVIEW OF THE FRAMEWORK

Our protocol specification language is an affine variant of RCF, a concurrent λ-calculus with message passing and refinement types originally introduced in [Bengtson et al. 2011]. We anticipate that RCF is very expressive and can be mapped to a large subset of F#. For better readability, in the examples we use F#-like syntax with polymorphic types: our theoretical framework lacks full-fledged polymorphism, but that can be recovered by duplicating definitions at multiple monomorphic types when needed.

2.1. Protocol verification with (affine) refinement types

Verifying distributed authorization protocols with refinement types presupposes that protocols be annotated with security *assumptions* and *assertions*. The former are formulas that are assumed to hold at a given point in time, and they are employed to specify authorization policies and to encode the credentials available to request authorization. In contrast, assertions act as guards defining the properties to be entailed by the assumptions and the underlying policy, to grant authorization [Fournet et al. 2005; 2007; Bengtson et al. 2011].

An example will help in making the discussion concrete. We introduce a system to place and ship orders in a distributed online service governed by a simple authorization policy, establishing that an order can be cleared for shipping to a user only if that user has indeed placed the order. For example, we could start by assuming the authorization policy encoded by the first-order formula: $\mathcal{P} \triangleq \forall x, y.(\mathsf{Order}(x,y) \Rightarrow \mathsf{Ship}(x,y))$. The security-annotated code corresponding to the online service scenario is given below:

let *place_order* = **fun** *ch id item skey* $\rightarrow$
  assume Order*(id,item);*
  let *pkt* = *sign skey (id,item)* **in** *send ch pkt*

let *ship_order* = **fun** *ch vkey* $\rightarrow$
  let *pkt* = *recv ch* **in**
    let *(xc, xit)* = *verify vkey pkt* **in**
      assert Ship*(xc,xit)*

The assumption Order*(id,item)* makes the required credential available to the *place_order* function, enabling the subsequent code to sign a request with the key *skey* and send it off over channel *ch*. Upon receiving the message, *ship_order* verifies the signature using the verification key *vkey*, retrieves the two components *xc* and *xit* of the request and asserts the formula Ship*(xc,xit)*.

A client and a server will execute the two functions, communicating on a shared channel *ch* and using a pair of corresponding signing and verification keys, as shown below (the server runs *ship_order* recursively to serve multiple requests):

let *prot_spec ch* =
  assume $\mathcal{P}$;
  let *sk* = *mksigkey ()* **in**
    let *vk* = *mkverkey sk* **in**
      let *client* = *(place_order ch "alice" "book" sk)* **in**
        let rec *server* = *(ship_order ch vk)* ⇌ *(server ch vk)* **in**
          client ⇌ *server*

The protocol specification given above may be proved *robustly* safe by existing refinement type systems: this ensures that the conjunction of all the assertions which will become active at runtime (i.e., $\mathsf{Ship}(\textit{"alice"},\textit{"book"})$), is entailed by the active assumptions (i.e., $\mathcal{P}$, $\mathsf{Order}(\textit{"alice"},\textit{"book"})$), despite the best efforts of an arbitrary opponent. Unfortunately, a closer look reveals that the authorization policy $\mathcal{P}$ is too weak to enforce desirable *resource-aware* access constraints: for instance, in our example the online service is presumably interested in ensuring that each user's order can be cleared and shipped only *once*, but in first-order logic we can prove:

$$\forall x, y.(\mathsf{Order}(x, y) \Rightarrow \mathsf{Ship}(x, y)), \mathsf{Order}(id, item) \vdash \mathsf{Ship}(id, item) \wedge \mathsf{Ship}(id, item),$$

i.e., a single payment by the user can lead to the same order being shipped twice, without violating the previous authorization policy and (robust) safety.

Remarkably, the desired resource-aware authorization policy can be naturally encoded in affine logic by assuming the formula: $\mathcal{P}_{okay} \triangleq\ !\forall x, y.(\mathsf{Order}(x, y) \multimap \mathsf{Ship}(x, y))$, where the bang modality (!) allows using the authorization policy arbitrarily many times in a proof, while the multiplicative implication ($\multimap$) ensures that formulas of the form $\mathsf{Order}(id, item)$ are *consumed* when proving $\mathsf{Ship}(id, item)$. Verifying the desired *injective* correspondence between placed and shipped orders amounts then just to reinterpreting the standard notion of (robust) safety by taking into account the *multiplicative* conjunction ($\otimes$) of the top-level assertions rather than the standard conjunction of first-order logic: roughly, this ensures that the (multi-)set of assumptions can be partitioned in different (multi-)sets, each proving one specific assertion, hence the same assumption is never used in the proof of two different assertions.

Extending refinement type systems to show compliance with respect to affine logic policies like $\mathcal{P}_{okay}$ is challenging. Technically, these type systems support a form of compositional reasoning enabled by the structure of the cryptographic key types, and the typing discipline enforced on them. Briefly, cryptographic key types are associated with refinement types of the form $\mathsf{Key}(\{x : T \mid F\})$, enforcing the following invariants: (i) to package a value $M : T$ with a key of this type, one must be able to prove $F(M)$ and consequently, (ii) upon extracting a value $w : T$ packaged under a key of this type, one may in turn assume the formula $F(w)$ to hold. These two invariants are enough to derive static proofs of robust safety in traditional refinement type systems drawing on classical and intuitionistic logics, but they fall short of providing the necessary guarantees in resource-conscious settings such as the one we consider here.

2.2. Exponential serialization for protecting affine formulas

Given the nature of affine formulas as consumable resources, an affine refinement type system must additionally provide protection against an unconstrained assumption of the refinement formulas conveyed by the key types [Bugliesi et al. 2011]. For instance, when receiving a packet signed with a key of type $\mathsf{SigKey}(x : T, \{x : U \mid \mathsf{Order}(x, y)\})$, we must ensure that *each time* we verify the signature (and assume $\mathsf{Order}(x, y)$) at the

receiver side, a corresponding assumption has indeed been introduced at the sender side.

Ensuring this kind of injective correspondence in distributed settings is known to require some protective measures, as an adversary may easily break it by mounting a replay attack and fool a receiver into deriving multiple assertions corresponding to one single assumption. We can see that in our running example: given the protocol specification defined above, assume we let it run over an untrusted network by passing the function *prot_spec* as a parameter to the function *adversary* defined below, which intercepts the message by the client and sends it twice to the server:

let *adversary prot* =
  let *ch = mkchan ()* **in**
    prot ch;
    let *m = recv ch* **in** *(send ch m)* ⚡ *(send ch m)*

The replay attack mounted by the adversary breaks the desired injective correspondence between assumptions and assertions, since the system admits a run in which the adversary intercepts the message exchanged on *ch* and duplicates it, leading to two assertions $\mathsf{Ship}(\textit{"alice"},\textit{"book"})$ being made against just one assumption $\mathsf{Order}(\textit{"alice"},\textit{"book"})$. More technically, in affine logic we have:

$$!\forall x,y.(\mathsf{Order}(x,y) \multimap \mathsf{Ship}(x,y)), \mathsf{Order}(id,item) \nvdash \mathsf{Ship}(id,item) \otimes \mathsf{Ship}(id,item),$$

hence the protocol above is *not* robustly safe in our affine setting.

The problem we just outlined is, in fact, rather general and may be stated as follows: data exchanged over the network is inherently exposed to replays, hence their credentials, occurring as refinements of cryptographic key types, must be protected so that replicating the data does not duplicate the credentials. In the type system, this may be achieved by guarding the refinements of the key types with *control* formulas, which are guaranteed to be assumed in at most one point of the protocol code.

The resulting typing discipline leverages the underlying computational measures to counter replay attacks. Though the details vary for the different computational mechanisms, the intuition applies uniformly. The types of cryptographic keys are built around *guarded* refinements of the form:

$$\{\tilde{w} : \tilde{T}, \tilde{x} : \tilde{U} \mid !(C(\tilde{w}) \multimap F(\tilde{x}))\},$$

protecting the credential $F(\tilde{x})$ with the control formula $C(\tilde{w})$. In a nonce-handshake protocol, for instance, $\tilde{w}$ may represent a challenger-generated nonce, call it n, and $C(n)$ may be the corresponding guard assumed by the challenger, modeling that the nonce has been freshly generated. Upon receiving the nonce, a responder willing to transmit M will package the pair (n, M) under a key with the above type as a payload: intuitively, the receiver can then open the cryptographic packet to assume the implication above and derive the desired formula $F(M)$ by *consuming* the formula $C(n)$, which was never sent on the network and remained thus under the control of the challenger.

Notice that guarded refinement types as the one above contain an *exponential* formula prefixed by the bang modality, hence opening messages packaged under a key with this type more than once does not really provide additional information to the receiver and is perfectly safe. We call this packaging technique *exponential serialization*, as it provides us with a safe way to transmit payload with affine refinement types over an untrusted network, using an encoding based on exponential formulas.

2.3. Serializers for security type-checking

There is one problem left with the intuition above. A responder possessing the credential $F(M)$ and willing to prove it to the challenger will not be able to do so, as

in affine logic an assumption $F(M)$ does not entail the guarded exponential formula $!(C(n) \multimap F(M))$, which the responder would need to prove to type-check the response. To close this gap, each affine assumption in the code must be associated with a corresponding *serializer*, to enable its use in the guarded refinements of the key types. Serializers have the general form:

$$!\forall \tilde{x}.\tilde{w}(F(\tilde{x}) \multimap !(C(\tilde{w}) \multimap F(\tilde{x})))),$$

and explicitly enable the transformation of the credential $F(\tilde{M})$ into its serialized form $!(C(\tilde{n} \multimap F(\tilde{M})))$, for appropriate terms $\tilde{n}$ and $\tilde{M}$.

Back to our example, assume we extend the protocol to include the nonce-handshake mentioned above:

let *place_order'* = **fun** *ch1 ch2 id item skey* $\rightarrow$
  let *nonce = recv ch1* **in**
    assume $\mathsf{Order}(id, item)$;
    let *pkt = sign skey (id,item,nonce)* **in** *send ch2 pkt*

let *ship_order'* = **fun** *ch1 ch2 vkey* $\rightarrow$
  let *mknonce = (***fun** *()* $\rightarrow$ **let** *x = mkfresh ()* **in assume** $\mathsf{N}(x)$; *x)* **in**
    let *nonce = mknonce ()* **in**
      send ch1 nonce;
      let *pkt = recv ch2* **in**
        let *(xc, xit, xn) = verify vkey pkt* **in**
          if *(xn = nonce)* **then**
            assert $\mathsf{Ship}(xc, xit)$
          else
            failwith *"unauthorized"*

We assume to be given access to a library function *mkfresh* : unit $\rightarrow$ bytes, which generates fresh bit-strings. The function *mknonce*: unit $\rightarrow$ $\{x : \text{bytes} \mid \mathsf{N}(x)\}$ is a wrapper around *mkfresh*, which additionally assumes the control formula $\mathsf{N}(x)$ over the returned value x. The new assumption is reflected by the *refined* return type of *mknonce*. Then, the typing of the signing and verification keys may be structured as follows:

$$skey \; : \; \mathsf{SigKey}(\{x : \text{string}, y : \text{string}, z : \text{bytes} \mid !(\mathsf{N}(z) \multimap \mathsf{Order}(x,y))\})$$
$$vkey \; : \; \mathsf{VerKey}(\{x : \text{string}, y : \text{string}, z : \text{bytes} \mid !(\mathsf{N}(z) \multimap \mathsf{Order}(x,y))\})$$

conveying the affine formula $\mathsf{Order}(xc, xit)$ conditionally to the *guard* $\mathsf{N}(nonce)$ assumed by *ship_order'*. If the guard can be proved only once, $\mathsf{Order}(xc, xit)$ can also be retrieved only once, irrespectively of the number of signature verifications performed. To type-check the protocol, we need to assume the expected serializer:

$$\mathcal{S} \triangleq !\forall x, y.z.(\mathsf{Order}(x,y) \multimap !(\mathsf{N}(z) \multimap \mathsf{Order}(x,y))).$$

Overall, we get the following revised protocol:

let *prot_spec' ch1 ch2* =
  assume $\mathcal{P}_{okay}$;
  assume $\mathcal{S}$;
  let *sk = mksigkey ()* **in**
    let *vk = mkverkey sk* **in**
      let *client'* = *(place_order' ch1 ch2 "alice" "book" sk)* **in**
        let rec *server'* = *(ship_order' ch1 ch2 vk)* ⇂ *(server' ch vk)* **in**
          client' ⇂ *server'*

We briefly discuss how the two protocol components type-check. We start from the server. Upon creating *nonce*, *server'* assumes the control formula $\mathsf{N}(nonce)$ based on the return type of the function *mknonce* by calling *ship_order'*. Upon verifying the received signature, it extracts the refinement $!(\mathsf{N}(xn) \multimap \mathsf{Order}(xc, xit))$ based on the type of the verification key. Then, from the assumption $\mathsf{N}(nonce)$ and the nonce-checking test $xn = nonce$ that protects the assertion, it derives $\mathsf{N}(xn)$. Now, with two $\multimap$ elimination steps, using the refinement above and the policy $\mathcal{P}_{okay}$, it derives the asserted formula $\mathsf{Ship}(xc, xit)$. As to the client, upon receiving the challenge, by calling the function *place_order'*, *client'* assumes the formula $\mathsf{Order}("alice", "book")$ and then signs the triple *(cid,item,nonce)* with *vk*. Typing the signature requires the serializer, which provides a direct way to prove the desired formula.

We notice here that serializers may be generated automatically for any given affine formula, and we prove that introducing them as additional assumptions is sound, in that it does not affect the set of entailed assertions, under the sufficient conditions discussed in Section 4. Furthermore, serializers capture a rather general class of mechanisms for ensuring timely communications, like session keys or timestamps, which are all based on the consumption of an affine resource to assess the freshness of an exchange. We discuss these patterns in our case studies in Sections 8-9.

3. REVIEW: AFFINE LOGIC

In our framework we focus on a simple, yet expressive, fragment of intuitionistic affine logic [Troelstra 1992]. We presuppose an underlying signature Σ of predicate symbols, ranged over by p, and function symbols, ranged over by f. The syntax of terms t and formulas F is defined by the following productions:

$$
\begin{array}{llll}
t & ::= & x \mid f(t_1, \ldots, t_n) & \text{terms } (f \text{ of arity } n \text{ in } \Sigma) \\
A & ::= & p(t_1, \ldots, t_n) \mid t = t' & \text{atoms } (p \text{ of arity } n \text{ in } \Sigma) \\
F & ::= & A \mid F \otimes F \mid F \multimap F \mid \forall x.F \mid !F \mid \mathbf{0} & \text{formulas}
\end{array}
$$

This is the multiplicative fragment of affine logic with conjunction ($\otimes$) and implication ($\multimap$), the universal quantifier ($\forall$), the exponential modality (!) to express persistent truths, logical falsity ($\mathbf{0}$) to express negation, and syntactic equality. The logical truth is written $\mathbf{1}$ and encoded as $() = ()$, where $()$ is the nullary function symbol encoding the RCF "unit" value[1]. The negation of F, written $F^{\perp}$, is encoded as $F \multimap \mathbf{0}$, while inequality, written $t \neq t'$, is encoded as $(t = t')^{\perp}$. For simplicity, we do not consider disjunction and existential quantification: the logic considered here suffices for our purposes and we leave further extensions as future work.

The entailment relation $\Delta \vdash F$ from multiset of formulas to formulas is given in Table I. Observe that, in affine logic, rule (WEAK) can be liberally applied to disregard formulas along a proof derivation, while rule (CONTR) is restricted to exponential formulas, allowing for their unbounded duplication. Intuitively, the combination of the two rules enforces the following usage policy for formulas: "every formula must be used *at most* once in a proof, with the exception of exponential formulas, which can be used arbitrarily many times". This is in contrast with linear logic, where each formula must be used *exactly* once [Girard 1995].

As informally discussed before, affine logic provides multiplicative counterparts of standard logical connectives: for instance, rule ($\otimes$-RIGHT) states that to prove the multiplicative conjunction $F_1 \otimes F_2$ from the hypotheses $\Delta = \Delta_1, \Delta_2$, we have to prove F_1 from Δ_1 and F_2 from Δ_2, thus each affine hypothesis in Δ is used either to prove F_1 or to prove F_2. Analogously, rule ($\multimap$-LEFT) formalizes the intuition that the multi-

---

[1]We mention here that RCF terms can be encoded into the logic using the locally nameless representation of syntax with binders [de Bruijn 1972], as shown in [Bengtson et al. 2011].

Table I The entailment relation $\Delta \vdash F$

(IDENT) $F \vdash F$	(WEAK) $\dfrac{\Delta \vdash F'}{\Delta, F \vdash F'}$	(CONTR) $\dfrac{\Delta, !F, !F \vdash F'}{\Delta, !F \vdash F'}$	($\otimes$-LEFT) $\dfrac{\Delta, F_1, F_2 \vdash F'}{\Delta, F_1 \otimes F_2 \vdash F'}$	($\otimes$-RIGHT) $\dfrac{\Delta_1 \vdash F_1 \quad \Delta_2 \vdash F_2}{\Delta_1, \Delta_2 \vdash F_1 \otimes F_2}$

($\multimap$-LEFT) $\dfrac{\Delta_1 \vdash F_1 \quad \Delta_2, F_2 \vdash F'}{\Delta_1, F_1 \multimap F_2, \Delta_2 \vdash F'}$	($\multimap$-RIGHT) $\dfrac{\Delta, F_1 \vdash F_2}{\Delta \vdash F_1 \multimap F_2}$	($\forall$-LEFT) $\dfrac{\Delta, F\{t/x\} \vdash F'}{\Delta, \forall x.F \vdash F'}$	($\forall$-RIGHT) $\dfrac{\Delta \vdash F \quad x \notin fv(\Delta)}{\Delta \vdash \forall x.F}$

(!-LEFT) $\dfrac{\Delta, F \vdash F'}{\Delta, !F \vdash F'}$	(!-RIGHT) $\dfrac{!\Delta \vdash F}{!\Delta \vdash !F}$	(FALSE) $\mathbf{0} \vdash F$	(=-SUBST) $\dfrac{\exists \sigma = mgu(t, t') \Rightarrow \Delta\sigma \vdash F\sigma}{\Delta, t = t' \vdash F}$	(=-REFL) $\Delta \vdash t = t$

plicative implication $F_1 \multimap F_2$ acts as a sort of reaction, which consumes the resources needed to prove the premise F_1 to produce the conclusion F_2.

Rule (!-LEFT) is often called the *dereliction* rule and allows exponential assumptions to be degraded to affine assumptions, which can be used at most once. Rule (!-RIGHT), instead, is typically referred to as the *promotion* rule, which allows one to prove exponential formulas starting from the proof of an affine formula: the notation $!\Delta$ means that every formula in Δ must be of the form $!F$. The two rules for equality (=-SUBST) and (=-REFL) are borrowed from [Tiu and Momigliano 2012]; in rule (=-SUBST), if the terms t and t' are not unifiable, then we consider the premise as trivially fulfilled.

4. METATHEORY OF EXPONENTIAL SERIALIZATION

Recall from Section 2.3 that we had to explicitly assume a serializer $\mathcal{S}$ to make our example protocol type-check. In principle, the introduction of this serializer among the assumed hypotheses could alter the intended semantics of the authorization policy $\mathcal{P}_{okay}$, due to the subtle interplay of formulas through the entailment relation defined in Table I. Here, we isolate sufficient conditions under which exponential serialization leads to a sound protection mechanism for affine formulas.

We presuppose that the signature Σ of predicate symbols is partitioned in two sets Σ_A and Σ_C. Atomic formulas A have the form $p(t_1, \ldots, t_n)$ for some $p \in \Sigma_A$; control formulas C have the same form, though with $p \in \Sigma_C$. We identify various categories of formulas defined by the following productions:

$$\begin{array}{llll} B & ::= & A \mid B \otimes B \mid B \multimap B \mid \forall x.B \mid !B & \text{base formulas} \\ P & ::= & B \mid C \mid P \otimes P & \text{payload formulas} \\ G & ::= & C \multimap P \mid !G & \text{guarded formulas} \end{array}$$

Base formulas B are formulas of an authorization policy, built from atomic formulas using logical connectives. We use base formulas as security annotations in the application code. For simplicity, we dispense in this section with equalities and $\mathbf{0}$ to ensure logical consistency: these elements are used in our typed analysis, but we stipulate that they are never directly assumed in the protocol code (and thus never serialized).

Payload formulas P are formulas which we want to serialize for communication over the untrusted network. Importantly, payload formulas comprise both base formulas and control formulas, which allows, e.g., for the transmission of fresh nonces to remote verifiers: this pattern is present in several authentication protocols [Gordon and Jeffrey 2003]. Finally, guarded formulas G are used to model the serialized version of payload formulas, suitable for transmission. Notice also that serializers are not generated by any of the previous productions, so we let S stand for any serializer of the form $!\forall \tilde{x}.(P \multimap !(C \multimap P))$. We write $\Delta \vdash F^n$ for $\Delta \vdash F \otimes \ldots \otimes F$ (n times), with the proviso that $\Delta \vdash F^0$ stands for $\Delta \nvdash F$.

Intuitively, given a multiset of assumptions Δ, the extension of Δ with the serializers $S_1, \ldots, S_n$ is sound if Δ and its extension derive the same payload formulas. As it turns out, this is only true when Δ satisfies additional conditions, which we formalize next.

Definition 4.1 (*Rank*). Let $rk : \Sigma_C \to \mathbb{N}$ be a total, injective function. Given a formula F, we define the *rank* of F with respect to rk, denoted by $rk(F)$, as follows:

$$
\begin{aligned}
rk(p(t_1, \ldots, t_n)) &= rk(p) & &\text{if } p \in \Sigma_C \\
rk(F_1 \otimes F_2) &= min\ \{rk(F_1), rk(F_2)\} & & \\
rk(F) &= +\infty & &\text{otherwise}
\end{aligned}
$$

Definition 4.2 (*Stratification*). A formula F is *stratified* with respect to a rank function rk if and only if: (i) $F = C \multimap P$ implies $rk(C) < rk(P)$; (ii) $F = P \multimap G$ implies that G is stratified; (iii) $F = \forall x.F'$ implies that F' is stratified; (iv) $F = \,!F'$ implies that F' is stratified. We assume F to be stratified in all the other cases. We say that a multiset of formulas Δ is stratified if and only if there exists a rank function rk such that each formula in Δ is stratified with respect to rk.

For instance, the multiset $C_1 \multimap C_2, C_2 \multimap C_3$, where C_1, C_2, C_3 are built over distinct predicate symbols, is stratified, given an appropriate choice of a rank function, while the multiset $C_1 \multimap C_2, C_2 \multimap C_1$ is not stratified. Stratification is required precisely to disallow these circular dependencies among control formulas and simplify the proof of our soundness result, Theorem 4.4 below. To prove that result, we need a further definition:

Definition 4.3 (*Controlled Multiset*). Let $\Delta = P_1, \ldots, P_m, S_1, \ldots, S_n$ be a stratified multiset of formulas. We say that Δ is *controlled* if and only if $\Delta \vdash C^k$ implies $k \leq 1$ for any control formula C.

The intuition underlying the definition may be explained as follows. Consider a multiset Δ, a payload formula P such that $\Delta \vdash P$ and let $S = \,!\forall \tilde{x}.(P \multimap \,!(C \multimap P))$ be a serializer for P. Now, the only way that S may affect derivability is by allowing for the duplication of the payload formula P via the exponential implication $!(C \multimap P)$, since the latter can be used arbitrarily often in a proof derivation. However, this effect is prevented if we are guaranteed that the control formula C guarding P is derived at most once in Δ: that is precisely what the condition above ensures.

THEOREM 4.4 (SOUNDNESS OF SERIALIZATION). *Let* $\Delta = P_1, \ldots, P_m$. *If* $\Delta' = \Delta, S_1, \ldots, S_n$ *is controlled and* $\Delta' \vdash P$, *then* $\Delta \vdash P$ *for all payload formulas* P.

PROOF. See Appendix A. □

Notice that checking if a multiset of formulas is controlled may be difficult, since this depends on logical entailment, hence it may be not obvious when the theorem above can be applied. Fortunately, however, we can isolate a sufficient criterion to decide whether a multiset of formulas is controlled, based on a simple syntactic check.

PROPOSITION 4.5 (CHECKING CONTROL). *If* $\Delta = P_1, \ldots, P_m, S_1, \ldots, S_n$ *is stratified and the control formulas occurring in* $P_1, \ldots, P_m$ *are pairwise distinct, then* Δ *is controlled.*

PROOF. See Appendix A. □

5. RCF AND SAFETY

We now review RCF [Bengtson et al. 2011], a concurrent λ-calculus with message passing primitives, which provides the core language around which our theory is developed.

Table II Syntax of RCF expressions

$M, N ::=$		*values*
	x	variable
	$()$	unit
	(M, N)	pair
	$\lambda x.\,E$	function
	$h\,M$	construction ($h \in \{\mathsf{inl}, \mathsf{inr}, \mathsf{fold}\}$)
$D, E ::=$		*expressions*
	M	value
	$M\,N$	application
	$M = N$	syntactic equality
	let $x = E$ in E'	let (scope of x is E')
	let $(x, y) = M$ in E	pair split (scope of x, y is E)
	match M with $h\,x$ then E else E'	match (scope of x is E)
	$(\nu a)E$	restriction (scope of a is E)
	$E \upharpoonright E'$	fork
	$a!M$	message send
	$a?$	message receive
	assume F	assumption
	assert F	assertion

We also formally introduce the resource-aware variant of the standard notion of safety for RCF, which we have been mentioning.

5.1. Review of RCF

We assume collections of names (a, b, c, m, n) and variables (x, y, z). The syntax of values and expressions of RCF is introduced in Table II. The notions of free names and free variables arise as expected, according to the scope defined in the table.

Values include variables, unit, pairs, functions and constructions; constructors account for the creation of standard tagged unions and iso-recursive types. We also encode the boolean values true $\triangleq$ inl() and false $\triangleq$ inr(). Expressions of RCF include standard λ-calculus constructs like values, applications, equality checks, lets, pair splits, and pattern matching, as well as primitives for concurrent, message-passing computations in the style of process algebras.

The semantics is mostly standard. The function application $(\lambda x.\,E)\,N$ evaluates to $E\{N/x\}$; the syntactic equality check $M = N$ evaluates to true when M is equal to N and to false otherwise; the let expression let $x = E$ in E' first evaluates E to a value N and then behaves as $E'\{N/x\}$; the pair splitting let $(x, y) = (M, N)$ in E evaluates to $E\{M/x\}\{N/y\}$; and the pattern matching match M with $h\,x$ then E else E' evaluates to $E\{N/x\}$ when M is equal to $h\,N$ for some N, while it evaluates to E' otherwise. We then have some constructs reminiscent of process algebras: expression $(\nu a)E$ generates a globally fresh channel name a and then behaves as E. Expression $E \upharpoonright E'$ evaluates E and E' in parallel, and returns the result of E'. Expression $a!M$ asynchronously outputs M on channel a and returns $()$. Expression $a?$ waits until a term N is available on channel a and returns N. These message-passing expressions can be used to model the sending and receiving functions "*send*" and "*recv*" that are used in the code of our examples and that we further explain in Section 7.1. Assumptions and assertions are stuck expressions, which are just needed to state our safety notion (see below). The formal semantics of RCF expressions is defined by the reduction rules in Table III.

The reduction semantics depends upon the heating relation $E \Rightarrow E'$, an asymmetric version of the standard structural congruence, to perform some syntactic rearrangements of expressions and allow reductions. We write $E \equiv E'$ to denote that both $E \Rightarrow E'$ and $E' \Rightarrow E$. The definition of the heating relation is presented in Table IV, the only

Table III Reduction semantics for RCF

$(\lambda x.\, E)\, N \to E\{N/x\}$	(RED FUN)
let $(x,y) = (M,N)$ in $E \to E\{M/x\}\{N/y\}$	(RED SPLIT)
match M with $h\, x$ then E else $E' \to$	(RED MATCH)
$\begin{cases} E\{N/x\} & \text{if } M = h\, N \text{ for some } N \\ E' & \text{otherwise} \end{cases}$	
$M = N \to \begin{cases} \text{true} & \text{if } M = N \\ \text{false} & \text{otherwise} \end{cases}$	(RED EQ)
$a!M \restriction a? \to M$	(RED COMM)
let $x = M$ in $E \to E\{M/x\}$	(RED LET VAL)
let $x = E$ in $E'' \to$ let $x = E'$ in E'' if $E \to E'$	(RED LET)
$(\nu a)E \to (\nu a)E'$ if $E \to E'$	(RED RES)
$E \restriction E'' \to E' \restriction E''$ if $E \to E'$	(RED FORK 1)
$E'' \restriction E \to E'' \restriction E'$ if $E \to E'$	(RED FORK 2)
$E \to E'$ if $E \Rightarrow D, D \to D', D' \Rightarrow E'$	(RED HEAT)

difference with respect to the original RCF presentation is the introduction of the rule (HEAT ASSERT ()), which simplifies our definition of safety.

Table IV Heating relation for RCF

$E \Rightarrow E$	(HEAT REFL)
$E \Rightarrow E''$ if $E \Rightarrow E'$ and $E' \Rightarrow E''$	(HEAT TRANS)
let $x = E$ in $E'' \Rightarrow$ let $x = E'$ in E'' if $E \Rightarrow E'$	(HEAT LET)
$(\nu a)E \Rightarrow (\nu a)E'$ if $E \Rightarrow E'$	(HEAT RES)
$E \restriction E'' \Rightarrow E' \restriction E''$ if $E \Rightarrow E'$	(HEAT FORK 1)
$E'' \restriction E \Rightarrow E'' \restriction E'$ if $E \Rightarrow E'$	(HEAT FORK 2)
$() \restriction E \equiv E$	(HEAT FORK ())
$a!M \Rightarrow a!M \restriction ()$	(HEAT MSG ())
assume $F \Rightarrow$ assume $F \restriction ()$	(HEAT ASSUME ())
assert $F \Rightarrow$ assert $F \restriction ()$	(HEAT ASSERT ())
$E' \restriction (\nu a)E \Rightarrow (\nu a)(E' \restriction E)$ if $a \notin \mathit{fn}(E')$	(HEAT RES FORK 1)
$(\nu a)E \restriction E' \Rightarrow (\nu a)(E \restriction E')$ if $a \notin \mathit{fn}(E')$	(HEAT RES FORK 2)
let $x = (\nu a)E$ in $E' \Rightarrow (\nu a)(\text{let } x = E \text{ in } E')$ if $a \notin \mathit{fn}(E')$	(HEAT RES LET)
$(E \restriction E') \restriction E'' \equiv E \restriction (E' \restriction E'')$	(HEAT FORK ASSOC)
$(E \restriction E') \restriction E'' \Rightarrow (E' \restriction E) \restriction E''$	(HEAT FORK COMM)
let $x = (E \restriction E')$ in $E'' \equiv E \restriction (\text{let } x = E' \text{ in } E'')$	(HEAT FORK LET)

5.2. Resource-aware safety

We are now ready to adapt the formal notion of safety defined for RCF expressions to our resource-aware setting. Intuitively, an expression E is safe if, for all runs, the *multiplicative* conjunction of the top-level assertions is entailed by the top-level assumptions. Giving a precise definition, however, is somewhat tricky and it is convenient to introduce the notion of *structure* for this purpose.

Let e denote an *elementary* expression, i.e., any expression that is not an assumption, assertion, restriction, let, fork, or send. Structures formalize the idea that a computation state has four components: (1) a multiset of assumed formulas F_i; (2) a multiset of asserted formulas F_j'; (3) a series of messages M_k sent on channels but not yet received; and (4) a series of elementary expressions e_ℓ being evaluated in parallel contexts. The definition of a structure **S** is given in Table V. Structures are convenient, since their syntactic form already exhibits all the necessary ingredients to state a simple notion of *static safety*, the basic building block for safety.

We can prove that every expression E can be transformed into a structure by heating, hence we can define a suitable notion of safety for any expression.

Table V Structures and static safety

$$\Pi_{i \in [1,n]} E_i \triangleq () \upharpoonright E_1 \upharpoonright \cdots \upharpoonright E_n$$
$$\mathcal{L}[e] ::= e \mid \mathsf{let}\ x = \mathcal{L}[e]\ \mathsf{in}\ E$$
$$\mathbf{S} ::= (\nu \widetilde{a})((\Pi_{i \in [1,m]} \mathsf{assume}\ F_i) \upharpoonright (\Pi_{j \in [1,n]} \mathsf{assert}\ F'_j) \upharpoonright (\Pi_{k \in [1,o]} c_k! M_k) \upharpoonright (\Pi_{\ell \in [1,p]} \mathcal{L}_\ell[e_\ell]))$$

The structure $\mathbf{S}$ above is *statically safe* if and only if $F_1, \ldots, F_m \vdash F'_1 \otimes \ldots \otimes F'_n$.

LEMMA 5.1 (STRUCTURE). *For every expression E, there exists a structure $\mathbf{S}$ such that $E \Rightarrow \mathbf{S}$.*

PROOF. By induction on the structure of E. □

Definition 5.2 (*Safety*). A closed expression E is *safe* if and only if, for all E' and $\mathbf{S}$, if $E \to^* E'$ and $E' \Rightarrow \mathbf{S}$, then $\mathbf{S}$ is statically safe.

The real property of interest, however, is stronger than the previous one: we desire protection despite the best efforts of an active opponent. We let an *opponent* be any closed expression of RCF which does not contain any assumption or assertion. The latter is a standard restriction, since opponents containing arbitrary assertions could vacuously falsify the property we target; this does not involve any loss of generality in practice, since we want to verify application code with respect to the security annotations placed therein. We note that security annotations are simply considered a tool for verification but that they hold no semantic meaning and are thus not necessary for the opponent code.

Definition 5.3 (*Robust Safety*). A closed expression E is *robustly safe* if and only if, for any opponent O, the application $O\ E$ is safe[2].

6. THE TYPE SYSTEM

Our refinement type system builds on previous work by Bengtson et al. [Bengtson et al. 2011], extending it to guarantee the correct usage of affine formulas and to enforce our revised notion of (robust) safety.

6.1. Types, typing environments, and base judgements

The syntax of types is defined in Table VI. Again the notions of free names and free variables arise as expected, according to the scope defined in the table.

Table VI Syntax of types

T, U, V ::=		*types*
	unit	unit type
	$x : T \to U$	dependent function type (scope of x is U)
	$x : T * U$	dependent pair type (scope of x is U)
	$T + U$	sum type
	$\mu\alpha.\,T$	iso-recursive type (scope of α is T)
	α	type variable
	$\{x : T \mid F\}$	refinement type (scope of x is F)

The unit value $()$ is given type unit. Sum types have the form $T + U$, iso-recursive types are denoted by $\mu\alpha.\,T$, and type variables are denoted by α. There exist various forms of dependent types: a function of type $x : T \to U$ takes as an input a value M of type T and returns a value of type $U\{M/x\}$; a pair (M, N) has type $x : T * U$ if M has type T and N has type $U\{M/x\}$; a value M has a refinement type $\{x : T \mid F\}$ if M

---

[2]Here, we use the standard syntactic sugar $O\ E$ for the expression let $x = O$ in let $y = E$ in $x\ y$.

has type T and the formula $F\{M/x\}$ holds true. We use type $\mathsf{Un} \triangleq \mathsf{unit}$ to model data that may come from, or be sent to the opponent, as it is customary for security type systems.[3] Type $\mathsf{bool} \triangleq \mathsf{unit} + \mathsf{unit}$ is inhabited by $\mathsf{true} \triangleq \mathsf{inl}()$ and $\mathsf{false} \triangleq \mathsf{inr}()$.

The type system comprises several typing judgements of the form $\Gamma; \Delta \vdash \mathcal{J}$, where $\Gamma; \Delta$ is a typing environment collecting all the information which can be used to derive $\mathcal{J}$. In particular, Γ contains the type bindings, while Δ comprises logical formulas that are supposed to hold at run-time. Formally, we let Γ be an ordered list of entries $\mu_1, \ldots, \mu_n$ and Δ be a multiset of affine logic formulas. Each entry μ_i in Γ denotes either a type variable (α), a kinding annotation ($\alpha :: k$), or a type binding for channels ($a \updownarrow T$) or variables ($x : T$). We let ε denote the empty list and $\emptyset$ the empty multiset. The *domain* of Γ, written $dom(\Gamma)$, is defined as follows: $dom(\alpha) = \{\alpha\}$; $dom(\alpha :: k) = \{\alpha\}$; $dom(a \updownarrow T) = \{a\}$; $dom(x : T) = \{x\}$; and $dom(\mu_1, \ldots, \mu_n) = dom(\mu_1) \cup \ldots \cup dom(\mu_n)$. The set of *free* variables and free names is denoted by *fnfv*. The definition is standard.

We first discuss the base judgements of the type system. We use the judgement $\Gamma; \Delta \vdash \diamond$ to denote that the typing environment $\Gamma; \Delta$ is well-formed, i.e., it satisfies some standard syntactic conditions (for instance, it does not contain duplicate type bindings for the same variable). The only remarkable point in the definition of $\Gamma; \Delta \vdash \diamond$ is that we forbid variables in Γ to be mapped to a refinement type: indeed, when extending a typing environment with a new type binding $x : T$, we will use the function ψ to place the structural type information in Γ and the function *forms* to place the associated refinements in Δ. We also write $\Gamma; \Delta \vdash T$ to denote that type T is well-formed in $\Gamma; \Delta$ and $\Gamma; \Delta \vdash F$ when the formulas in Δ entail the formula F. We often abuse notation and write $\Gamma; \Delta \vdash F_1, \ldots, F_n$ to stand for $\Gamma; \Delta \vdash F_1 \otimes \ldots \otimes F_n$, with the proviso that $\Gamma; \Delta \vdash \emptyset$ is equivalent to $\Gamma; \Delta \vdash \mathbf{1}$. A complete formal definition of the described elements is given in Table VII below.

Table VII Auxiliary functions and base judgements

$$\psi(U) = \begin{cases} \psi(T) & \text{if } U = \{x : T \mid F\} \\ U & \text{otherwise} \end{cases} \qquad forms(y : U) = \begin{cases} F\{y/x\}, forms(y : T) & \text{if } U = \{x : T \mid F\} \\ \emptyset & \text{otherwise} \end{cases}$$

(ENV EMPTY)
$\varepsilon; \emptyset \vdash \diamond$

$$\frac{\text{(TYPE ENV ENTRY)}}{\Gamma; \Delta \vdash \diamond \qquad dom(\mu) \cap dom(\Gamma) = \emptyset \qquad \mu = x : T \Rightarrow T = \psi(T) \wedge fnfv(T) \subseteq dom(\Gamma)}{\Gamma, \mu; \Delta \vdash \diamond}$$

$$\frac{\text{(FORM ENV ENTRY)}}{\Gamma; \Delta \vdash \diamond \qquad fnfv(F) \subseteq dom(\Gamma)}{\Gamma; \Delta, F \vdash \diamond} \qquad \frac{\text{(TYPE)}}{\Gamma; \Delta \vdash \diamond \qquad fnfv(T) \subseteq dom(\Gamma)}{\Gamma; \Delta \vdash T}$$

$$\frac{\text{(DERIVE)}}{\Gamma; \Delta \vdash \diamond \qquad fnfv(F) \subseteq dom(\Gamma) \qquad \Delta \vdash F}{\Gamma; \Delta \vdash F}$$

6.2. Environment rewriting

We stipulate that all the type information stored in Γ can be used arbitrarily often in the derivation of any judgement of our type system, hence we dispense with affine types[4]. The treatment of the formulas in Δ is subtler, since affine resources must be

---

[3]Note that other types built over Un are available to the opponent through subtyping.

[4]In Section 6.8 we thoroughly discuss why this does not involve any loss in expressiveness, by showing an encoding of affine types through exponential serialization.

used at most once during type-checking: in particular, we need to split the environment Δ among subderivations to avoid the unbounded duplication of the formulas therein. However, a simple splitting of the formulas in Δ would lead to a very restrictive type system. To illustrate, let $\Delta \triangleq A, A \multimap !B$: if we just distributed the formulas A and $A \multimap !B$ between two distinct subderivations, then the formula $!B$ would be available only in (at most) one subderivation, despite it being an exponential formula, which we may want to use arbitrarily often during type-checking.

The general structure of the rules of our system then looks as follows:

$$\frac{\Gamma;\Delta_1 \vdash \mathcal{J}_1 \qquad \ldots \qquad \Gamma;\Delta_n \vdash \mathcal{J}_n \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\ldots,\Delta_n}{\Gamma;\Delta \vdash \mathcal{J}}$$

where $\Gamma;\Delta \hookrightarrow \Gamma;\Delta'$ denotes the *environment rewriting* of $\Gamma;\Delta$ to $\Gamma;\Delta'$. This relation is defined by rule (REWRITE) below:

$$\begin{array}{c} (\text{REWRITE}) \\ \dfrac{\Delta \vdash \Delta' \qquad \Gamma;\Delta \vdash \diamond \qquad \Gamma;\Delta' \vdash \diamond}{\Gamma;\Delta \hookrightarrow \Gamma;\Delta'} \end{array}$$

where we write $\Delta \vdash F_1,\ldots,F_n$ to denote that $\Delta \vdash F_1 \otimes \ldots \otimes F_n$, again with the proviso that $\Delta \vdash \emptyset$ stands for $\Delta \vdash \mathbf{1}$. Coming back to our previous example, notice that we have $A, A \multimap !B \vdash !B \otimes !B$ in affine logic, hence we can obtain two copies of $!B$ upon rewriting and distribute them between two distinct subderivations upon type-checking. As we will explain in Section 6.5, for soundness reasons we will often rely on rewriting of the form $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta'$, where $!\Delta'$ is a so-called *exponential* environment, i.e., an environment of the form $!F_1,\ldots,!F_n$.

The adoption of the environment rewriting relation as an house-keeping device for the formulas in Δ greatly improves the expressiveness of the type system in a very natural way. This idea of extending to the typing environment a number of context manipulation rules from the underlying substructural logic was first proposed by Mandelbaum et al. [Mandelbaum et al. 2003], even though their solution is technically different from ours. Namely, the authors of [Mandelbaum et al. 2003] allow for applications of arbitrary left rules from the logic inside the typing environment, while our proposal is reminiscent of the (CUT) rule typical of sequent calculi. We find this solution simpler to present and more convenient to prove sound.

Interestingly, all the non-determinism introduced by the application of the rewriting rules and the splitting of the logical formulas among the premises of the type rules can be effectively tamed by the algorithmic type system discussed in Section 10.

6.3. Kinding

Security type systems often rely on a kinding relation to discriminate whether or not messages of a specific type may be sent to the attacker or generated by it. The kinding judgement $\Gamma;\Delta \vdash T :: k$ denotes that type T is of kind k. We distinguish between two kinds: kind $k = \mathsf{pub}$ denotes that the inhabitants of a given type are public and may be sent to the attacker, while kind $k = \mathsf{tnt}$ denotes that the inhabitants of a given type are tainted and may come from the attacker. We let $\overline{\mathsf{pub}} \triangleq \mathsf{tnt}$ and $\overline{\mathsf{tnt}} \triangleq \mathsf{pub}$.

The complete kinding relation is given in Table VIII. Most of the rules resemble those presented in other security type systems [Bengtson et al. 2011; Backes et al. 2011] and only differ in the treatment of affine formulas, which is similar to the one we employ for typing values and expressions. We postpone the discussion on this point until the next section, where it will be easier to provide an intuitive understanding. Here, we just point out some simple observations, which should hopefully guide the reader in understanding a few important aspects.

Table VIII Kinding rules

$$
\frac{\Gamma;\Delta \vdash \diamond \quad (\alpha :: k) \in \Gamma}{\Gamma;\Delta \vdash \alpha :: k} \text{ (KIND VAR)}
\qquad
\frac{\Gamma;\Delta \vdash \diamond}{\Gamma;\Delta \vdash \mathsf{unit} :: k} \text{ (KIND UNIT)}
$$

$$
\text{(KIND FUN)}
\qquad
\frac{\Gamma;!\Delta_1 \vdash T :: \overline{k} \qquad \Gamma, x : \psi(T);!\Delta_2 \vdash U :: k \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2}{\Gamma;\Delta \vdash x : T \rightarrow U :: k}
$$

$$
\text{(KIND PAIR)}
\qquad
\frac{\Gamma;!\Delta_1 \vdash T :: k \qquad \Gamma, x : \psi(T);!\Delta_2 \vdash U :: k \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2}{\Gamma;\Delta \vdash x : T * U :: k}
$$

$$
\text{(KIND SUM)}
\qquad
\frac{\Gamma;!\Delta_1 \vdash T :: k \qquad \Gamma;!\Delta_2 \vdash U :: k \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2}{\Gamma;\Delta \vdash T + U :: k}
$$

$$
\text{(KIND REC)}
\qquad
\frac{\Gamma, \alpha :: k;!\Delta' \vdash T :: k \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta'}{\Gamma;\Delta \vdash \mu\alpha.\,T :: k}
$$

$$
\text{(KIND REFINE PUBLIC)}
\qquad
\frac{\Gamma;\Delta \vdash \{x : T \mid F\} \qquad \Gamma;\Delta \vdash T :: \mathsf{pub}}{\Gamma;\Delta \vdash \{x : T \mid F\} :: \mathsf{pub}}
$$

$$
\text{(KIND REFINE TAINTED)}
\qquad
\frac{\Gamma;\Delta_1 \vdash \psi(T) :: \mathsf{tnt} \qquad \Gamma, y : \psi(T);\Delta_2 \vdash \mathit{forms}(y : T) \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2 \qquad T \text{ refined}}{\Gamma;\Delta \vdash T :: \mathsf{tnt}}
$$

The type unit is assumed to be both public and tainted by (KIND UNIT). According to (KIND PAIR), a pair type is public if both its components are public and can be disclosed to the opponent. Conversely, by the same rule, a pair type is tainted if both its components are tainted, since, if even a single component of the pair is untainted, then the pair cannot come from the opponent. The kinding of sum types (KIND SUM) behaves analogously. By rule (KIND FUN) a function type is public (thus available to the attacker) only if its return type is public (otherwise $\lambda x.\, M_{\text{secret}}$ could be public and leak a secret to the attacker) and its argument type is tainted such that it can be called by the attacker. The treatment of tainted function types is dual. To give kind k to an iso-recursive type with a bound variable α, the rule (KIND REC) proceeds recursively and extends the typing environment in the premise with the kinding annotation $\alpha :: k$. These kinding annotations are used when kinding a type variable (KIND VAR). By (KIND REFINE PUBLIC) a refinement type is public if the structural type it refines is public, while by (KIND REFINE TAINTED) it is tainted if its structural information is tainted and its refinements are entailed by the typing environment.

6.4. Subtyping

The subtyping judgment $\Gamma;\Delta \vdash T <: U$ expresses the fact that T is a subtype of U and, thus, values of type T can be safely used in place of values of type U. The complete presentation of the subtyping relation can be found in Table IX.

We first note that subtyping is reflexive by (SUB REFL). Furthermore, the subtyping judgment makes public types subtype of tainted types through rule (SUB PUB TNT), and further describes standard subtyping relations for types sharing the same structure: for instance, pair and sum types are covariant (cf. (SUB PAIR) and (SUB SUM)), while function types are contravariant in their arguments and covariant in their return types (cf. (SUB FUN)). Intuitively, this means that a function can safely replace another function if it is "more liberal" in the types it accepts and "more conservative" in the types it returns.

The rule for iso-recursive types (SUB POS REC) is borrowed from [Backes et al. 2011] and it differs from the standard Amber rule proposed in the original presentation of RCF: the rule we consider here is easier to prove sound and the loss of expressiveness is very mild. We refer the interested reader to [Backes et al. 2011] for further discussion on this technical point.

The most interesting subtyping rule in Table IX is (SUB REFINE), which subsumes the rules (SUB REFINE LEFT) and (SUB REFINE RIGHT) from the original presentation

Table IX Subtyping rules

$$(\textsc{Sub Refl}) \qquad \frac{\Gamma;\Delta \vdash T}{\Gamma;\Delta \vdash T <: T}$$

$$(\textsc{Sub Pub Tnt}) \qquad \frac{\Gamma;\Delta_1 \vdash T :: \mathsf{pub} \qquad \Gamma;\Delta_2 \vdash U :: \mathsf{tnt} \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash T <: U}$$

$$(\textsc{Sub Fun}) \qquad \frac{\Gamma;!\Delta_1 \vdash T' <: T \qquad \Gamma, x:\psi(T');!\Delta_2 \vdash U <: U' \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2}{\Gamma;\Delta \vdash x:T \to U <: x:T' \to U'}$$

$$(\textsc{Sub Pair}) \qquad \frac{\Gamma;!\Delta_1 \vdash T <: T' \qquad \Gamma, x:\psi(T);!\Delta_2 \vdash U <: U' \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2}{\Gamma;\Delta \vdash x:T * U <: x:T' * U'}$$

$$(\textsc{Sub Sum}) \qquad \frac{\Gamma;!\Delta_1 \vdash T <: T' \qquad \Gamma;!\Delta_2 \vdash U <: U' \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2}{\Gamma;\Delta \vdash T + U <: T' + U'}$$

$$(\textsc{Sub Pos Rec}) \qquad \frac{\Gamma,\alpha;!\Delta' \vdash T <: T' \qquad \alpha \text{ occurs only positively in } T \text{ and } T' \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta'}{\Gamma;\Delta \vdash \mu\alpha.\,T <: \mu\alpha.\,T'}$$

$$(\textsc{Sub Refine}) \qquad \frac{\Gamma;\Delta_1 \vdash \psi(T) <: \psi(U) \qquad \Gamma, y:\psi(T);\Delta_2, forms(y:T) \vdash forms(y:U) \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2 \qquad T \text{ and/or } U \text{ refined}}{\Gamma;\Delta \vdash T <: U}$$

of RCF, which are shown below:

$$(\textsc{Sub Refine Left}) \qquad \frac{\Gamma \vdash \{x:T \mid F\} \qquad \Gamma \vdash T <: U}{\Gamma \vdash \{x:T \mid F\} <: U}$$

$$(\textsc{Sub Refine Right}) \qquad \frac{\Gamma \vdash T <: U \qquad \Gamma, x:T \vdash F}{\Gamma \vdash T <: \{x:U \mid F\}}$$

The first rule allows discarding unneeded logical formulas and conforms to the core idea of "refinement" typing: values of type $\{x:T \mid F\}$ can be safely replaced for values of type T, since they are just values of type T further qualified by the information encoded by the formula F. The second rule, instead, generalizes the substitution principle underlying subtyping to the refinement formulas: for instance, we have $\emptyset;\varepsilon \vdash \{x:\mathsf{Un} \mid x = 5\} <: \{x:\mathsf{Un} \mid x > 0\}$, since the logical condition $x = 5$ is stronger than the condition $x > 0$.

A natural adaptation of (Sub Refine Right) to our affine setting would be:

$$(\textsc{Sub Refine Wrong}) \qquad \frac{\Gamma;\Delta_1 \vdash T <: U \qquad \Gamma, x:\psi(T);\Delta_2, forms(x:T) \vdash F \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash T <: \{x:U \mid F\}}$$

Unfortunately, this rule is unsound, since the affine formulas of T could actually be duplicated and we could prove, for instance: $\emptyset;\varepsilon \vdash \{x:\mathsf{Un} \mid F\} <: \{z:\{x:\mathsf{Un} \mid F\} \mid F\}$ by using (Sub Refl) in the left premise of (Sub Refine Wrong). This cannot happen with our new rule, since $F \nvdash F \otimes F$ in affine logic.

While it is in principle possible to find out other sound counterparts of (Sub Refine Right) in an affine setting, previous work [Bugliesi et al. 2011] highlighted that the technical treatment of these rules is rather complicated, and we find rule (Sub Refine) more convenient for proofs. The previous discussion should have also provided an intuition on the reasons behind a slightly more restrictive treatment for subtyping pairs and functions with respect to the original RCF paper, i.e., we must take care in applying the refinement stripping function ψ before extending the typing environment in the second premise of the corresponding rules.

6.5. Typing values

The typing judgement $\Gamma; \Delta \vdash M : T$ denotes that value M is given type T under environment $\Gamma; \Delta$. The typing rules for values are given in Table X.

Table X Typing rules for values

$$
\text{(Val Var)} \quad \frac{\Gamma; \Delta \vdash \diamond \quad (x : T) \in \Gamma}{\Gamma; \Delta \vdash x : T}
$$

$$
\text{(Val Unit)} \quad \frac{\Gamma; \Delta \vdash \diamond}{\Gamma; \Delta \vdash () : \mathsf{unit}}
$$

$$
\text{(Val Fun)} \quad \frac{\Gamma, x : \psi(T); !\Delta', forms(x : T) \vdash E : U \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta'}{\Gamma; \Delta \vdash \lambda x.\, E : x : T \to U}
$$

$$
\text{(Val Pair)} \quad \frac{\Gamma; !\Delta_1 \vdash M : T \qquad \Gamma; !\Delta_2 \vdash N : U\{M/x\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, !\Delta_2}{\Gamma; \Delta \vdash (M, N) : x : T * U}
$$

$$
\text{(Val Refine)} \quad \frac{\Gamma; \Delta_1 \vdash M : T \qquad \Gamma; \Delta_2 \vdash F\{M/x\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; \Delta_1, \Delta_2}{\Gamma; \Delta \vdash M : \{x : T \mid F\}}
$$

$$
\text{(Val Inl)} \quad \frac{\Gamma; !\Delta' \vdash M : T \qquad \Gamma; !\Delta' \vdash U \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta'}{\Gamma; \Delta \vdash \mathsf{inl}\ M : T + U}
$$

$$
\text{(Val Inr)} \quad \frac{\Gamma; !\Delta' \vdash M : U \qquad \Gamma; !\Delta' \vdash T \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta'}{\Gamma; \Delta \vdash \mathsf{inr}\ M : T + U}
$$

$$
\text{(Val Fold)} \quad \frac{\Gamma; !\Delta' \vdash M : T\{\mu a.\, T/\alpha\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta'}{\Gamma; \Delta \vdash \mathsf{fold}\ M : \mu \alpha.\, T}
$$

The rules for variable and unit typing are standard: variables are typed by looking up their type binding in the typing environment Γ using (Val Var); the unit value can be given type unit under any well-formed environment using (Val Unit). Rule (Val Refine) is a natural adaptation to an affine setting of the standard rule for refinement types: a value M has type $\{x : T \mid F\}$ if M has type T and the formula $F\{M/x\}$ holds true. Rules (Val Fun) and (Val Pair) are more interesting: recall, in fact, that our type system does not include affine types, since the type information in Γ is propagated to all the premises of a typing rule. It is then crucial for soundness that both pairs and functions are type-checked in an exponential environment, i.e., an environment of the form $!F_1, \ldots, !F_n$. Indeed, using an affine formula F from the typing environment to give a pair (M, N) type $x : T * \{y : U \mid F\}$ would lead to an unbounded duplication of F upon repeated pair splitting operations on (M, N). Similar restrictions apply also to sum types (cf. (Val Inl) and (Val Inr)) and iso-recursive types (cf. (Val Fold)).

Notice that allowing for affine refinements, but forbidding affine types, confines the problem of resource management to the formula environment Δ, thus simplifying the technical development of the type system, as well as its algorithmic variant. In Section 6.8 we explain how our exponential serialization technique can be leveraged to encode affine types in our framework, hence our choice does not lead to any loss of expressiveness.

6.6. Typing expressions

The typing judgement $\Gamma; \Delta \vdash E : T$ denotes that expression E is given type T under environment $\Gamma; \Delta$. The typing rules for expressions are given in Table XI.

Several typing rules make use of the *extraction* relation $E \rightsquigarrow [\Delta \mid D]$ that destructively collects all the assumed formulas Δ from the expression E and returns the expression D obtained by purging E of its assumptions. The relation is defined in Table XII and will be explained further in the context of rule Exp Fork.

Rule (Exp Subsum) is a standard subsumption rule for expressions: if E can be given type T, then it can be conservatively given any supertype of T. The rule for typing function applications (Exp Appl) divides the formula environment Δ among its premises and checks that the type of the argument corresponds to the expected

Table XI Typing rules for expressions

(EXP SUBSUM)
$$\frac{\Gamma;\Delta_1 \vdash E : T \qquad \Gamma;\Delta_2 \vdash T <: T' \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash E : T'}$$

(EXP APPL)
$$\frac{\Gamma;\Delta_1 \vdash M : x : T \to U \qquad \Gamma;\Delta_2 \vdash N : T \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash M\ N : U\{N/x\}}$$

(EXP LET)
$$\frac{E \leadsto^{\emptyset} [\Delta' \mid D] \quad \Gamma;\Delta_1 \vdash D : T \quad \Gamma, x : \psi(T);\Delta_2, forms(x:T) \vdash E' : U \quad x \notin fv(U) \qquad \Gamma;\Delta,\Delta' \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash \mathsf{let}\ x = E\ \mathsf{in}\ E' : U}$$

(EXP SPLIT)
$$\frac{\Gamma;\Delta_1 \vdash M : x : T * U \qquad \Gamma, x : \psi(T), y : \psi(U);\Delta_2, forms(x:T), forms(y:U), !((x,y) = M) \vdash E : V \quad \{x,y\} \cap fv(V) = \emptyset \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash \mathsf{let}\ (x,y) = M\ \mathsf{in}\ E : V}$$

(EXP MATCH)
$$\frac{\Gamma;\Delta_1 \vdash M : T \quad \Gamma, x : \psi(H);\Delta_2, forms(x:H), !(h\ x = M) \vdash E : U \quad \Gamma;\Delta_2 \vdash E' : U \quad (h,H,T) \in \{(\mathsf{inl}, T_1, T_1 + T_2), (\mathsf{inr}, T_2, T_1 + T_2), (\mathsf{fold}, T'\{\mu\alpha.T'/\alpha\}, \mu\alpha.T')\} \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash \mathsf{match}\ M\ \mathsf{with}\ h\ x\ \mathsf{then}\ E\ \mathsf{else}\ E' : U}$$

(EXP EQ)
$$\frac{\Gamma;\Delta_1 \vdash M : T \quad \Gamma;\Delta_2 \vdash N : U \quad x \notin fv(M) \cup fv(N) \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2}{\Gamma;\Delta \vdash M = N : \{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = N)\}}$$

(EXP ASSUME)
$$\frac{\Gamma;\Delta, F \vdash \mathsf{assume}\ \mathbf{1} : T \qquad F \neq \mathbf{1}}{\Gamma;\Delta \vdash \mathsf{assume}\ F : T}$$

(EXP TRUE)
$$\frac{\Gamma;\Delta \vdash \diamond}{\Gamma;\Delta \vdash \mathsf{assume}\ \mathbf{1} : \mathsf{unit}}$$

(EXP ASSERT)
$$\frac{\Gamma;\Delta \vdash F}{\Gamma;\Delta \vdash \mathsf{assert}\ F : \mathsf{unit}}$$

(EXP RES)
$$\frac{E \leadsto^a [\Delta' \mid D] \qquad \Gamma, a \updownarrow T;\Delta,\Delta' \vdash D : U \qquad a \notin fn(U)}{\Gamma;\Delta \vdash (\nu a)E : U}$$

(EXP SEND)
$$\frac{\Gamma;\Delta \vdash M : T \qquad (a \updownarrow T) \in \Gamma}{\Gamma;\Delta \vdash a!M : \mathsf{unit}}$$

(EXP RECV)
$$\frac{\Gamma;\Delta \vdash \diamond \qquad (a \updownarrow T) \in \Gamma}{\Gamma;\Delta \vdash a? : T}$$

(EXP FORK)
$$\frac{E_1 \leadsto^{\emptyset} [\Delta_1 \mid D_1] \quad E_2 \leadsto^{\emptyset} [\Delta_2 \mid D_2] \quad \Gamma;\Delta_1' \vdash D_1 : T_1 \quad \Gamma;\Delta_2' \vdash D_2 : T_2 \qquad \Gamma;\Delta,\Delta_1,\Delta_2 \hookrightarrow \Gamma;\Delta_1',\Delta_2'}{\Gamma;\Delta \vdash E_1 \upharpoonright E_2 : T_2}$$

---

function argument type; in the return type we substitute the argument to the variable bound in the function type, thus implementing a form of value dependent typing. In rule (EXP SPLIT) we exploit the logic to keep track of the performed pair splitting operation and make type-checking more precise; a similar technique is used also in (EXP MATCH) and (EXP EQ). The treatment of channels is mostly standard: For each new channel a, a message type is determined ($a \updownarrow T$) and added to the typing environment Γ (cf. EXP RES) that is used to type-check the remaining expression. The rules for sending (EXP SEND) and receiving (EXP RECV) messages on such channel assure that the sent/received messages have the correct type. Rule (EXP ASSERT) is standard and requires an asserted formula F to be derivable from the formulas collected by the typing environment: in fact, these formulas under-approximate the formulas which will be assumed at runtime. As we will see in the explanation of the rule (EXP FORK) below, due to the affine nature of the logic, the treatment of assumptions is a delicate

task. Assumptions can be typed using either rule (EXP TRUE) or (EXP ASSUME). The former describes the trivial case of a truth assumption $\mathbf{1}$ that is always given type unit, the latter is used for more complex formulas F, which are added to the formula environment Δ. Intuitively, the intended usage of these rules to type-check an assumption assume F with type T is as follows: (1) Prove $\Gamma; \Delta, F \vdash$ assume $\mathbf{1}$: unit by rule (EXP TRUE); (2) Refine the type unit into T by subtyping; (3) Use (EXP ASSUME) to conclude $\Gamma; \Delta \vdash$ assume $F : T$.

The most complex rule is (EXP FORK): intuitively, when type-checking the parallel expressions $E_1 \upharpoonright E_2$, assumptions in E_1 can be safely used to type-check assertions in E_2 and vice-versa. On the other hand, we need to prevent an affine assumption in E_1 from being used twice to justify assertions in both E_2 *and* E_1. This is achieved by the extraction relation, i.e., through the premises of the form $E_i \rightsquigarrow [\Delta_i \mid D_i]$: the extraction operation destructively collects all the assumptions from the expression E_i and returns the expression D_i obtained by purging E_i of its assumptions. The typing environment is then extended with the collected assumptions and partitioned to type-check the purged expressions D_1 and D_2. For instance, we can show that the expression assume $F \upharpoonright$ assert F is well-typed, while the expression (assume $F \upharpoonright$ assert F) $\upharpoonright$ assert F is not: indeed, notice that the latter is not safe according to Definition 5.2.

The extraction relation $E \rightsquigarrow^{\widetilde{a}} [\Delta \mid D]$ is formally defined in Table XII. Note that we annotate the arrow with a list of names $\widetilde{a}$ to prevent formulas containing free names from being extracted outside the scope of the respective binders. For instance, in the expression $((\nu a)\text{assume } F(a)) \upharpoonright$ assert $F(a)$ we do not want to use the assumption to type-check the parallel assertion, since the scope of the name a is limited to the assumption itself. The extraction relation is used to type-check any expression possibly containing "active" assumptions, i.e., lets (cf. (EXP LET)), restrictions (cf. (EXP RES)), and assumptions themselves (cf. (EXP ASSUME), which hardcodes the extraction).

Table XII The extraction relation

(EXTR FORK)	(EXTR LET)
$\dfrac{E_1 \rightsquigarrow^{\widetilde{a}} [\Delta_1 \mid D_1] \qquad E_2 \rightsquigarrow^{\widetilde{a}} [\Delta_2 \mid D_2]}{E_1 \upharpoonright E_2 \rightsquigarrow^{\widetilde{a}} [\Delta_1, \Delta_2 \mid D_1 \upharpoonright D_2]}$	$\dfrac{E_1 \rightsquigarrow^{\widetilde{a}} [\Delta \mid D_1]}{\text{let } x = E_1 \text{ in } E_2 \rightsquigarrow^{\widetilde{a}} [\Delta \mid \text{let } x = D_1 \text{ in } E_2]}$

(EXTR RES)	(EXTR ASSUME)	(EXTR EXP)
$\dfrac{E \rightsquigarrow^{a,\widetilde{b}} [\Delta \mid D]}{(\nu a)E \rightsquigarrow^{\widetilde{b}} [\Delta \mid (\nu a)D]}$	$\dfrac{F \neq \mathbf{1} \qquad fn(F) \cap \{\widetilde{a}\} = \emptyset}{\text{assume } F \rightsquigarrow^{\widetilde{a}} [F \mid \text{assume } \mathbf{1}]}$	$\dfrac{\text{no other rule applies}}{E \rightsquigarrow^{\widetilde{a}} [\emptyset \mid E]}$

6.7. Formal results

The main soundness results for our type system are given below.

THEOREM 6.1 (SAFETY). *If $\varepsilon; \emptyset \vdash E : T$, then E is safe.*

PROOF. See Appendix B. □

THEOREM 6.2 (ROBUST SAFETY). *If $\varepsilon; \emptyset \vdash E : $ Un, then E is robustly safe.*

PROOF. See Appendix B. □

Theorem 6.2 above and Theorem 4.4 (Soundness of Serialization) constitute the two building blocks of our static verification technique, which we may finally summarize as follows. Given any expression E, we identify the payload formulas assumed in E, and construct their serializers $S_1, \ldots, S_n$. Let then $E^\star = $ assume $S_1 \otimes \cdots \otimes S_n \upharpoonright E$ be the original expression extended with the serializers. By Theorem 6.2, if $\varepsilon; \emptyset \vdash E^\star : $ Un,

then $E^\star$ is robustly safe. By Theorem 4.4, so is the original expression E, provided that a further invariant holds for $E^\star$, namely that all multisets of formulas assumed during the evaluation of $E^\star$ are controlled.

While this latter invariant is not enforced by our type system, the desired guarantees may be achieved by requiring that the assumption of control formulas be confined within system code packaged into library functions, providing certified access and management of the capabilities associated with those formulas. The certification of the system code provided by the library function, in turn, may be achieved with limited effort, based on the sufficient condition provided by Proposition 4.5. Actually, we observe that the syntactic criterion proposed by the proposition becomes a *semantic* property of the program to type-check, since programs contain variables to be replaced at runtime: we will discuss for our examples how we verify that the typing environment satisfies the conditions required for robust safety.

6.8. Discussion: encoding affine types

We now discuss how we can take advantage of exponential serialization to encode affine types in our type system. For the sake of simplicity, we focus on the encoding of affine pairs, but the same ideas applies uniformly to other data types (i.e., tagged unions and iso-recursive types).

Consider the typing environment $\Gamma; \Delta \triangleq x : \mathsf{Un}, y : \mathsf{Un}; A(x), B(y)$. Standard refinement type systems [Bengtson et al. 2011] allow for the following type judgement:

$$\Gamma; \Delta \vdash (x, y) : \{x : \mathsf{Un} \mid A(x)\} * \{y : \mathsf{Un} \mid B(y)\}$$

If the formulas $A(x)$ and $B(y)$ are interpreted as affine resources, however, the previous type assignment is sound only as long as the pair (x, y) can be split only once, since every application of rule (EXP SPLIT) for pair destruction introduces the formulas $A(x), B(y)$ into the typing environment of the continuation. Since our type system does not feature affine types and has no way to enforce a single deconstruction of a pair, it conservatively forbids the previous type judgement, in that the premises of rule (VAL PAIR) require an exponential typing environment.

Nevertheless, the following type judgement is allowed by our type system:

$$x : \mathsf{Un}, y : \mathsf{Un}; A(x), B(y), S_1, S_2 \vdash (x, y) : \{x : \mathsf{Un} \mid A'(x)\} * \{y : \mathsf{Un} \mid B'(y)\}$$

where $A'(x) \triangleq\ !(C_1(x) \multimap A(x))$ and $B'(y) \triangleq\ !(C_2(y) \multimap B(y))$ are the serialized variants of $A(x)$ and $B(y)$ respectively, while $S_1 \triangleq\ !\forall x.(A(x) \multimap A'(x))$ and $S_2 \triangleq\ !\forall y.(B(y) \multimap B'(y))$ are the corresponding serializers. Here, the main idea for type-checking is to appeal to environment rewriting to consume the affine formulas $A(x)$ and $B(y)$, and introduce their exponential counterparts $A'(x)$ and $B'(y)$ into the typing environment before assigning a type to the components of the pair. In fact, notice that we have:

$$x : \mathsf{Un}, y : \mathsf{Un}; A(x), B(y), S_1, S_2 \hookrightarrow x : \mathsf{Un}, y : \mathsf{Un}; A'(x), B'(y),$$

hence we can prove the following type judgement:

$$\frac{x : \mathsf{Un}, y : \mathsf{Un}; A'(x) \vdash x : \{x : \mathsf{Un} \mid A'(x)\} \qquad x : \mathsf{Un}, y : \mathsf{Un}; B'(y) \vdash y : \{y : \mathsf{Un} \mid B'(y)\}}{x : \mathsf{Un}, y : \mathsf{Un}; A(x), B(y), S_1, S_2 \vdash (x, y) : \{x : \mathsf{Un} \mid A'(x)\} * \{y : \mathsf{Un} \mid B'(y)\}}$$

The interesting point now is that the pair (x, y) can be split arbitrarily often, but the affine formulas $A(x)$ and $B(y)$ can be retrieved at most once, as long as the control formulas $C_1(x)$ and $C_2(y)$ are assumed at most once in the application code. In this way, we recover the expressiveness provided by affine types. We actually even go beyond that, allowing for a liberal usage of the value itself, as opposed to enforcing the affine usage of any data structure which contains an affine component, as dictated by many

earlier substructural frameworks (see [Fähndrich and DeLine 2002] for a thorough discussion on this point).

7. A LIBRARY FOR COMMUNICATION AND CRYPTOGRAPHY

In this section we describe the primitives for communication that we use throughout the examples in this work and discuss how we encode cryptography using sealing. We note that our encoding of both communication and cryptography benefits from the notion of exponential serialization: we will never use channels, references, or cryptographic operations directly for messages with affine refinements, but we instead rely on exponentially serialized versions of such refinements. Formally, our libraries build on so-called *exponential types* that do not carry an affine refinement and are defined in Table XIII. Since these types do not need to be protected from replication we can immediately leverage existing non-affine libraries [Bengtson et al. 2011].

For the sake of simplicity, the definitions of the necessary functions and types are parametric in a type variable γ used to denote exponential types. We recall, however, that our system does not support full polymorphism, but we can recover its effects by replicating library code to specialize it to the different types we need. Most of the content of this section is taken from [Bengtson et al. 2011] and included for the reader's convenience to make the paper self-contained.

Table XIII Exponential types

T exponential if $\begin{cases} T \in \{\text{unit}, \alpha\} \\ U \text{ exponential} \\ T_1 \text{ exponential and } T_2 \text{ exponential} \\ T_1 \text{ exponential and } T_2 \text{ exponential} \\ T_1 \text{ exponential and } T_2 \text{ exponential} \\ U \text{ exponential} \end{cases}$	$\begin{aligned} & \\ & \text{for } T = \{x : U \mid !F\} \\ & \text{for } T = x : T_1 \to T_2 \\ & \text{for } T = x : T_1 * T_2 \\ & \text{for } T = T_1 + T_2 \\ & \text{for } T = \mu\alpha. U \end{aligned}$

7.1. An encoding of channels and messaging

In RCF channels are not values, hence they cannot be shared dynamically among principals. That same effect may however be recovered with the following encoding of channels for messages of exponential type γ (and the associated primitives for message passing).

We report both the communication interface and its implementation below.

type $\text{Ch}(\gamma) = (\gamma \to \text{unit}) * (\text{unit} \to \gamma)$
val *mkchan :* $\text{unit} \to \text{Ch}(\gamma)$
val *send :* $\text{Ch}(\gamma) \to \gamma \to \text{unit}$
val *recv :* $\text{Ch}(\gamma) \to \gamma$

let *mkchan* = **fun** _ $\to$ **(new** *a***)(fun** $x \to a!x$, **fun** _ $\to a?$**)**
let *send* = **fun** $c\ x \to$ **let** *(s, r) = c* **in** $s\ x$
let *recv* = **fun** $c \to$ **let** *(s, r) = c* **in** $r\ ()$

We note that references can be encoded analogously.

type $\text{Ref}(\gamma) = \text{Ch}(\gamma)$
val *mkref :* $\gamma \to \text{Ref}(\gamma)$
val *setref :* $\text{Ref}(\gamma) \to \gamma \to \text{unit}$
val *deref :* $\text{Ref}(\gamma) \to \gamma$

let *mkref* = **fun** $x \rightarrow$ **let** r = *mkchan ()* **in** *send r x; r*
let *setref* = **fun** $r\ x$ $\rightarrow$ **let** _ = *recv r* **in** *send r x*
let *deref* = **fun** r $\rightarrow$ **let** x = *recv r* **in** *send r x; x*

In the following we typically write "*r := v*" for "*setref r v*" and "*!r*" for "*deref r*".

Note that the code for dereferencing a reference will not type-check for types that are not exponential, since the value retrieved from the reference is used twice: it is stored back into the reference and returned. Without the serialization approach, we would thus need to change the implementation, for instance, by using destructive references that erase their content after a read.

7.2. A sealing-based encoding of cryptography

Formal cryptography can be encoded inside RCF in terms of *sealing* [Morris 1973; Sumii and Pierce 2007]. A *seal* for a type T is a pair of functions: a sealing function $T \rightarrow \mathsf{Un}$ and an unsealing function $\mathsf{Un} \rightarrow T$. Intuitively, for symmetric cryptography, these functions model encryption and decryption operations, respectively. A payload of type T can be sealed to type Un and sent over the untrusted network; conversely, a message retrieved from the network with type Un can be unsealed to its correct type T. This mechanism is implemented in terms of a list of pairs, which is stored in a global reference that can only be accessed using the sealing and unsealing functions. Upon sealing, the payload p is paired with a fresh, public value h (the *handle*) representing its sealed version, and the pair (p, h) is stored in the list; conversely, the unsealing function looks for the handle h in the list and returns the associated payload p.

Since for symmetric cryptography the possession of the key allows to perform both encryption and decryption operations, for such cryptographic schemes we identify the key with the seal, i.e., we give access to both the sealing and the unsealing functions to any owner of the key and we let $\mathsf{SymKey}(T) \triangleq (T \rightarrow \mathsf{Un}) * (\mathsf{Un} \rightarrow T)$. Different cryptographic primitives, like public key encryptions and signature schemes, can be encoded following the same recipe: for instance, since the owner of a signing key is typically able to verify her own signature, the sealing-based abstraction of a signing key may consist of both the sealing and the unsealing functions, and be given type $\mathsf{SigKey}(T) \triangleq (T \rightarrow \mathsf{Un}) * (\mathsf{Un} \rightarrow T)$. The corresponding verification key, instead, should comprise only the unsealing function and be given type $\mathsf{VerKey}(T) \triangleq \mathsf{Un} \rightarrow T$. The functions "*sign*" and "*verify*" introduced in Section 2 can then be straightforwardly implemented: *sign M N* just extracts the first component of M and calls it with parameter N, while *verify M N* simply invokes M with parameter N.

As stated above, another important benefit of exponential serialization is that we can immediately leverage the sealing-based cryptographic library proposed by Bengtson et al. [Bengtson et al. 2011], since we will define cryptographic operations to be performed only on messages of exponential type. Without the serialization approach, we would need to define a different implementation of the sealing/unsealing functions: namely, we would have to enforce that an affine payload is never extracted more than once from the list stored in the global reference, hence the dereferencing/unsealing function would have to remove the payload from the secret list. This would complicate the sealing-based abstraction of cryptography and require additional reasoning to justify its soundness [Backes et al. 2010]. Instead, with our approach, the unsealing function does not need to be changed: we can invoke it an arbitrary number of times to retrieve the payload, but the associated refinements will be retrieved at most once through exponential serialization.

We give full details of the cryptographic API used throughout this paper (just the types, not the code) below.

type $\mathsf{Seal}(\gamma) = (\gamma \to \mathsf{Un}) * (\mathsf{Un} \to \gamma)$
type $\mathsf{SealRef}(\gamma) = \mathsf{Ref}(\mathsf{List}(\gamma * \mathsf{Un}))$

val *mkseal* $:$ string $\to \mathsf{Seal}(\gamma)$
val *seal* $: \mathsf{SealRef}(\gamma) \to \gamma \to \mathsf{Un}$
val *unseal* $: \mathsf{SealRef}(\gamma) \to \mathsf{Un} \to \gamma$

type $\mathsf{SymKey}(\gamma) = \mathit{Sym}$ **of** $\mathsf{Seal}(\gamma)$
val *mksymkey* $:$ unit $\to \mathsf{SymKey}(\gamma)$
val *sencrypt* $: \mathsf{SymKey}(\gamma) \to \gamma \to \mathsf{Un}$
val *sdecrypt* $: \mathsf{SymKey}(\gamma) \to \mathsf{Un} \to \gamma$

type $\mathsf{SigKey}(\gamma) = \mathit{SK}$ **of** $\mathsf{Seal}(\gamma)$
type $\mathsf{VerKey}(\gamma) = \mathit{VK}$ **of** $(\mathsf{Un} \to \gamma)$
val *mksigkey* $:$ unit $\to \mathsf{SigKey}(\gamma)$
val *mkverkey* $: \mathsf{SigKey}(\gamma) \to \mathsf{VerKey}(\gamma)$
val *sign* $: \mathsf{SigKey}(\gamma) \to \gamma \to \mathsf{Un}$
val *verify* $: \mathsf{VerKey}(\gamma) \to \mathsf{Un} \to \gamma$

type $\mathsf{DecKey}(\gamma) = \mathit{DK}$ **of** $\mathsf{Seal}(\gamma)$
type $\mathsf{EncKey}(\gamma) = \mathit{EK}$ **of** $(\gamma \to \mathsf{Un})$
val *mkdeckey* $:$ unit $\to \mathsf{DecKey}(\gamma)$
val *mkenckey* $: \mathsf{DecKey}(\gamma) \to \mathsf{EncKey}(\gamma)$
val *encrypt* $: \mathsf{EncKey}(\gamma) \to \gamma \to \mathsf{Un}$
val *decrypt* $: \mathsf{DecKey}(\gamma) \to \mathsf{Un} \to \gamma$

8. EXAMPLE: EPMO

We are finally ready to see our type system at work. We consider a variant of *EPMO*, a nonce-based e-payment protocol proposed by Guttman et al. [Guttman et al. 2004].

8.1. Protocol description

The protocol narration is informally represented in Table XIV (the meaning of the security annotations is explained below).

Initially, a customer C contacts a merchant M to buy some goods g for a given price p; the request is encrypted under the public key of the merchant, $\mathsf{ek}(k_M)$ (which we use as shorthand for "*mkenckey* k_M" throughout the example), and includes a fresh nonce, n_C. If M agrees to proceed in the transaction by providing a response signed with the signing key sk_M, C informs her bank B to authorize the payment. The bank replies by providing C a receipt of authorization, called the *money order*, which is then forwarded to M. Now M can verify that C is entitled to pay for the goods and complete the transaction by sending a signed request to B to cash the money order. At the end of the run, the bank transfers the funds and the merchant ships the goods to the customer.

8.2. Protocol analysis and challenges

A peculiarity of the protocol is that the identifier n_C is employed by C to authenticate *two* different messages, namely the replies by M and B. This pattern cannot be validated by most existing type systems, since the mechanisms hardcoded therein to deal with nonce-handshakes enforce the freshness of each nonce to be checked only once. Our framework, instead, allows for a very natural treatment of such authentication pattern, whose implementation can be written mostly oblivious of the security verifi-

Table XIV A variant of the *EPMO* protocol
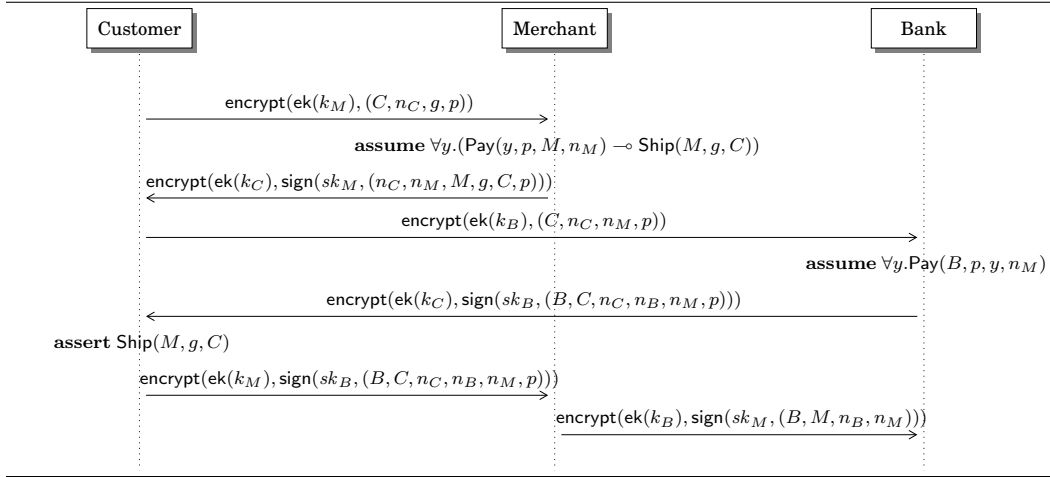
cation process, based on lightweight logical annotations. For the sake of simplicity, we focus only on the aspects of the verification connected to the guarantees provided to C, which are the most interesting ones.

We define two predicates used in the analysis: $\mathsf{Pay}(B, p, M, n_M)$ states that B authorizes the payment p to M in reference to the order identified by n_M, while $\mathsf{Ship}(M, g, C)$ formalizes that M will ship the goods g to C. After the first step of the protocol, we let the merchant M state the formula $\forall y.(\mathsf{Pay}(y, p, M, n_M) \multimap \mathsf{Ship}(M, g, C))$, to signify that she does not care about which bank is going to authorize the payment, but, as long as there is one authorizing bank, she will ship the good g to the client C at the end of the transaction. Conversely, after the appropriate checks on the client's account, we let the bank B assume the formula $\forall y.\mathsf{Pay}(B, p, y, n_M)$, to model that she authorizes the payment for the transaction n_M to any merchant chosen by the client. These two credentials allow the customer C to assert the formula $\mathsf{Ship}(M, g, C)$, a formal assurance on the validity of the transaction.

8.3. Type-checking the customer

The protocol code for the customer, enriched with the most relevant type annotations and the necessary serializers, is shown below. For the sake of readability, we use again F#-like syntax and some standard syntactic sugar like tuples, refined tuple types, algebraic types, and pattern matchings: all these can be encoded in RCF and our type system using standard techniques [Bengtson et al. 2011].

```
(∗ Serializer for M, needed to type-check M ∗)
assume !∀xp, xM, xnM, xg, xC, xnC.
        (∀y.(Pay(y, xp, xM, xnM) ⊸ Ship(xM, xg, xC)) ⊸
        !(N1(xnC) ⊸ (∀y.(Pay(y, xp, xM, xnM) ⊸ Ship(xM, xg, xC)))))

(∗ Serializer for B, needed to type-check B ∗)
assume !∀yB, yp, ynC, ynM.
        (∀y.(Pay(yB, yp, y, ynM)) ⊸ !(N2(ynC) ⊸ (∀y.(Pay(yB, yp, y, ynM)))))

(∗ Typing the message from M to C ∗)
type MsgMC = MsgMC of (xnC : bytes ∗ xnM : bytes ∗ xM : string ∗ xg : string
```

$*\; xC$: string $*\; xp$: int)
$\{!(\mathsf{N1}(xnC) \multimap \forall y.(\mathsf{Pay}(y, xp, xM, xnM) \multimap \mathsf{Ship}(xM, xg, xC))\}$

(∗ Typing the message from B to C ∗)
type MsgBC = *MsgBC* **of** $(yB$: string $*\; yC$: string $*\; ynC$: bytes $*\; ynB$: bytes
$*\; ynM$: bytes $*\; yp$: int)$\{!(\mathsf{N2}(ynC) \multimap \forall y.(\mathsf{Pay}(yB, yp, y, ynM))\}$

(∗ Generate transaction identifiers ∗)
let *mktid :* unit $\to \{x$: bytes $|\; \mathsf{N1}(x) \otimes \mathsf{N2}(x)\}$ = **fun** *() →*
  let *xf = mkfresh ()* **in assume** *(*$\mathsf{N1}(xf) \otimes \mathsf{N2}(xf)$*); xf*

(∗ Customer code ∗)
let *cust C addC M addM B addB g p kC ekM ekB*
          (vkM: VerKey(MsgMC + MsgMB) *(vkB:* VerKey(MsgBC)*) =*
  let *nC = mktid ()* **in**
  (∗ $\mathsf{N1}(nC)$ *and* $\mathsf{N2}(nC)$ *hold true ∗)*
  let *msgCM1 = encrypt ekM (C, nC, g, p)* **in** *send addM msgCM1;*
  let *signMC = decrypt kC (receive addC)* **in**
  let *plainMC = verify vkM signMC* **in**
  match *plainMC* **with** *MsgMC (=nC, xnM, =M, =g, =C, =p) →*
      (∗ $!(\mathsf{N1}(nC) \multimap \forall y.(\mathsf{Pay}(y, p, M, xnM) \multimap \mathsf{Ship}(M, g, C))$ *holds true ∗)*
      let *msgCB = encrypt ekB (C, nC, xnM, p)* **in** *send addB msgCB;*
      let *signBC = decrypt kC (receive addC)* **in**
      let *plainBC = verify vkB signBC* **in**
        match *plainBC* **with** *MsgBC (=B, =C, =nC, xnB, =xnM, =p) →*
            (∗ $!(\mathsf{N2}(nC) \multimap \forall y.(\mathsf{Pay}(B, p, y, xnM))$ *holds true ∗)*
            assert $\mathsf{Ship}(M, g, C)$*;*
            let *msgCM2 = encrypt ekM signBC* **in** *send addM msgCM2*

Initially, we let the customer call the library function *mktid*, which generates a fresh transaction identifier, corresponding to n_C in the protocol specification, and provides via its return type two distinct capabilities $\mathsf{N1}(nC)$ and $\mathsf{N2}(nC)$, later employed to authenticate the two different messages received by C.

Since the signing key of M is used to certify messages of two different types (at steps 2 and 6 of the protocol), the corresponding verification key available to the customer through the variable *vkM* refers to a sum type. We present only the MsgMC component of this type, since it is the one needed to type-check the code of C: the refined formula in the corresponding type definition is retrieved upon verification of *signMC* and describes the promise by M to ship the goods as soon as the requested payment has been authorized by any bank chosen by the client.

We finally use *vkB* to convey the other formula which is needed to type-check C, namely a statement that B authorizes the payment to any merchant to whom C wishes to transfer the money order: this statement is available after verifying the message *signBC*. The hypotheses collected by C are enough to prove her assertion, i.e., to be sure that the request by M has been fulfilled and the goods will be shipped, hence the implementation is well-typed.

Notice that, to conclude that the code actually respects the authorization policy despite the introduction of the serializers, we also have to show that the program ensures the invariant that each control formula is assumed at most once, corresponding to the sufficient condition for control dictated by Proposition 4.5. This is an easy task to carry out, since we can just observe that control formulas are only assumed in the body of

the *mktid* function, which in turn only performs these assumptions over the results of the *mkfresh* function for the generation of fresh bitstrings.

9. EXAMPLE: KERBEROS

In the *EPMO* protocol presented before, the nonce n_C is checked twice by the customer C and plays the role of a transaction identifier. Interestingly, there are protocols where these identifiers are not just checked multiple times, but also by different parties. This is exactly the case for the mutual authentication step of the Kerberos protocol [Steiner et al. 1988].

9.1. Protocol description

An informal narration of the protocol is shown in Table XV (the meaning of the security annotations is explained below).

Table XV The Kerberos protocol (mutual authentication)
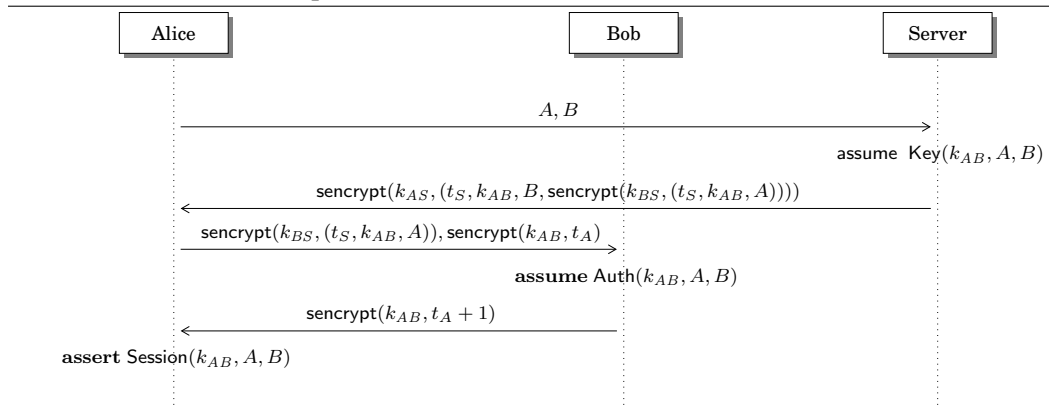
The goal of the protocol is to establish a fresh session key k_{AB} between principals A and B through a trusted server S, which shares a symmetric key with both A and B. Kerberos employs timestamps like t_S and t_A to prove session recentness and protect against replay attacks. Initially, A contacts the server S, providing the identities of the two agents A and B who want to establish a session. The server generates a fresh timestamp t_S and a new session key k_{AB}, then it packages all this information into a message for A and a message for B, which are combined by a nested encryption at step 2 of the protocol. Later, A removes the outer layer of the encryption, checks t_S and retrieves k_{AB}. If the timestamp is fresh, she forwards the inner encrypted message to B; additionally, A includes a fresh timestamp t_A encrypted under k_{AB}. Now B can decrypt the message encrypted by S, check its freshness, and retrieve the session key k_{AB}. Using this key, B can disclose the timestamp t_A and reply to A with $t_A + 1$, thus authenticating herself.

9.2. Protocol analysis and challenges

An intriguing point for our static verification technique is that the timestamp t_S generated by the server is checked by both A and B to ensure that the session key k_{AB} is fresh. As anticipated, this pattern is more sophisticated than the one we discussed for *EPMO*, but the expressiveness of our underlying affine logic framework allows for a simple encoding, discussed in the next section. For the sake of simplicity, in the following we will just focus on the verification of the initiator A.

We start by defining two predicates used in the analysis: $\mathsf{Key}(k_{AB}, A, B)$ states that k_{AB} is a fresh symmetric key intended to establish a session between A and B, while $\mathsf{Auth}(k_{AB}, A, B)$ formalizes that B wishes to communicate with A using key k_{AB}. Intuitively, these are the guarantees available to A after steps 2 and 4 of the protocol, respectively: by combining these two assurances, A can conclude that k_{AB} is a fresh session key which can be safely used to communicate with B. We model this last information through the predicate $\mathsf{Session}(k_{AB}, A, B)$ and we formalize the previous deduction by assuming the authorization policy:

$$!\forall x, y, z.(\mathsf{Key}(x, y, z) \otimes \mathsf{Auth}(x, y, z) \multimap \mathsf{Session}(x, y, z)).$$

We next discuss how we can show the compliance of the protocol against the previous policy by refinement type-checking.

9.3. Implementing and typing timestamps

We turn our attention to the implementation. We build on a very simple library for timestamp management, that we allow the principals to access. We note that timestamps are modeled as monotonic counters. To guarantee the freshness of a timestamp in the case that the opponent executes the protocol function multiple times, we pair the counter with a global, instance-dependent, fresh random bitstring $rand$ that is created at the beginning of the protocol specification using the function $mkfresh$. This usage of a random bitstring models the assumption that different sessions of Kerberos running in parallel will use different timestamps. Of course, we could consider more realistic and complicated implementations, but the following one suffices to convey the intuition about our methodology:

```
(∗ Typing a timestamp ∗)
type TStamp = TS of (bytes ∗ int)

(∗ Increment a timestamp by 1 ∗)
let inc_ts t =
  match t with TS (rt, tt) →
  TS (rt,  tt + 1)

(∗ Pick a fresh timestamp, based on the value stored in r ∗)
let get_ts r =  fun () →
  r := inc_ts !r; !r

(∗ Check a timestamp t for freshness, based on the value stored in r ∗)
let check_ts r id t' =
  match !r with TS (rt, tt) →
  match t' with TS (=rt, tt') →
  if (tt' > tt) then
    r := t'; assume F(id, t')
  else
    failwith "not_a_fresh_timestamp"

(∗ The handle to access the two functions above ∗)
let init_ts  rand glob id =
  let tss = !glob in
  let res = search tss id in
  match res with
  | Some(r) −> (get_ts r, check_ts r id)
```

> | *None*   −> **let** *newref = mkref TS (rand, 0)* **in**
> *glob := (id,newref)::tss; (get_ts newref, check_ts newref id)*

Each principal stores the last received timestamp in a reference, created by an invocation to the function *init_ts*, described below. The function *inc_ts*: $\mathsf{TStamp} \to \mathsf{TStamp}$ is used to increment a timestamp by 1, the function *get_ts*: $\mathsf{Ref}(\mathsf{TStamp}) \to \mathsf{unit} \to \mathsf{TStamp}$ is used to create fresh timestamps, and the dependent function *check_ts*: $\mathsf{Ref}(\mathsf{TStamp}) \to x : \mathsf{string} \to y : \mathsf{TStamp} \to \{\_ : \mathsf{unit} \mid \mathsf{F}(x,y)\}$ is used to check whether a received timestamp y is fresh and can be used to deem timely a communication with the principal x. The code of the function performs a conditional branch: if the timestamp is fresh, it *assumes* the logical formula encoding such a fact; otherwise, it fails. The function **failwith** throws an exception, so it can be safely given the polymorphic type $\mathsf{string} \to \alpha$; as a consequence, *check_ts* can be given the previous dependent function type, whose refined return type provides the freshness assumption.

The function *init_ts* is more complicated. It takes three parameters: the global instantiation-specific nonce *rand*, the identity of a principal *id* and a global reference *glob*, containing a list of pairs *(id',r')*, where *id'* is the identity of a principal and *r'* is a reference containing the last timestamp presented by *id'* (*TS (rand, 0)* if none). The function starts by retrieving this list from *glob*, bounding it to *tss*, and then uses an auxiliary function *search* to detect if there exists an entry of the form *(id,r)* in *tss*. If this is the case, *init_ts* returns a pair of functions *(get_ts r, check_ts r id)*, which will allow the caller to get a fresh timestamp and to check the freshness of the timestamps received by id. If *id* has never presented a timestamp when *init_ts* is invoked, the function creates a fresh reference containing *TS (rand, 0)* and updates the list stored in the reference *glob* to preserve the expected invariant, then it returns again a pair of functions for timestamp management. This implementation ensures that different instances of a protocol participant with the same identity will share the same counter for timestamps, which is important to protect the protocol against replay attacks. The *init_ts* function has type:

$$\mathsf{bytes} \to \mathsf{Ref}(\mathsf{List}(\mathsf{string} * \mathsf{TStamp})) \to x : \mathsf{string} \to$$
$$((\mathsf{unit} \to \mathsf{TStamp}) * (y : \mathsf{TStamp} \to \{\_ : \mathsf{unit} \mid \mathsf{F}(x,y)\})).$$

9.4. Typing the session key using self-dependent key types

Before discussing the implementation of the principal A, we must first consider a subtle issue related to verification. We pointed out that, at step 4 of the protocol, A must be able to infer that k_{AB} has been previously authenticated by B. The problem for verification is that the formula $\mathsf{Auth}(k_{AB}, A, B)$ modeling this fact must be conveyed by the type of the key k_{AB} itself, but neither the key k_{AB} nor the two identifiers A and B occur in the *payload* of the last protocol message, hence we cannot predicate on them using dependent typing. While the problem of letting the payload of the key refer to the identifiers A and B can be solved quite easily, since the referred to identities are globally and publicly known, the problem of letting the payload of a key predicate over the key itself is more involved due to lexical scoping. We show how to devise an encoding to solve the problem of self-dependent key types, which is close in spirit to the session key treatment advocated in previous work [Bugliesi et al. 2012].

Here we rely on a sealing-based encoding, where the *self-dependent* key k_{AB} consists of a *key identifier* i_{AB} and a pair k'_{AB} composed of the sealing and unsealing functions, thus having the form $k_{AB} = (i_{AB}, k'_{AB})$. The predicate Auth of the protocol refers to the identifier i_{AB} of the key k_{AB}, i.e., we actually assume $\mathsf{Auth}(i_{AB}, A, B)$ rather than $\mathsf{Auth}(k_{AB}, A, B)$ as we were discussing in the previous informal overview. The link between each self-dependent key k and its respective key identifier i is logically modeled

by the predicate $\mathsf{KeyIdent}(k,i)$, which holds true for all valid key-identifier pairs. The adapted authorization policy then looks as follows:

$$!\forall w, x, y, z.(\mathsf{Key}(x, y, z) \otimes \mathsf{KeyIdent}(x, w) \otimes \mathsf{Auth}(w, y, z) \multimap \mathsf{Session}(x, y, z)).$$

In the following we present the definition of our sealing-based library for the self-dependent session key k_{AB}. For presentation convenience, we make use of the following notation:

type MsgAB<*x,y,z*> = $(\{t : \mathsf{TStamp} \mid !(\mathsf{F}(y, t) \multimap \mathsf{Auth}(x, y, z))\})$ + TStamp

to denote the (open) type MsgAB of the session key payload. Here, $x \in fv(\mathsf{MsgAB}$<*x,y,z*>$)$ refers to the key identifier, while $y, z \in fv(\mathsf{MsgAB}$<*x,y,z*>$)$ refer to the globally available public identifiers A and B respectively. Note that this type is a sum type, since the key k_{AB} will be used by B to encrypt a timestamp of type $(t : \mathsf{TStamp})\{!(\mathsf{F}(B, t) \multimap \mathsf{Auth}(x, y, z))\}$ and by A to encrypt a non-refined timestamp of type TStamp (since we do not focus on the verification of B here). The sealing-based library for the dependent key k_{AB} shared between A and B is given below:

(∗ Closed type of the session key established by Kerberos.
*   Here, w stands for the key identifier discussed above ∗)*
type DSymKey = *DSym* **of** $(w : \mathsf{string} * ((\mathsf{MsgAB}$<*w,A,B*>$ \to \mathsf{Un}) *$
$(\mathsf{Un} \to \mathsf{MsgAB}$<*w,A,B*>$)))$

(∗ Generate a fresh identifier ∗)
val *new_fresh_id:* $\mathsf{unit} \to \mathsf{string}$

(∗ Create a new self-dependent key ∗)
let *mkdepkey:* $\mathsf{unit} \to \mathsf{DSymKey}$ = **fun** *()* →
  let *id* = *new_fresh_id ()* **in**
    let *s* = *mkseal "dsymkey"* **in**
      DSym (id,s)

(∗ Get the key identifier corresponding to a self-dependent key ∗)
let *get_key_ident k:* $(k : \mathsf{DSymKey} \to \{x : \mathsf{string} \mid !\mathsf{KeyIdent}(k, x)\})$ =
    match *k* **with** *DSym (x, _)* → **assume** $!\mathsf{KeyIdent}(k, x)$*; x*

(∗ Self-dependent symmetric encryption function ∗)
let *depencrypt x k m:* $(x : \mathsf{string} \to \mathsf{DSymKey} \to \mathsf{MsgAB}$<*x,A,B*>$ \to \mathsf{Un})$ =
    match *k* **with** *DSym (=x, (seal, _))* → *seal m*

(∗ Self-dependent symmetric decryption function ∗)
let *depdecrypt x k c:* $(x : \mathsf{string} \to \mathsf{DSymKey} \to \mathsf{Un} \to \mathsf{MsgAB}$<*x,A,B*>$)$ =
    match *k* **with** *DSym (=x, (_, unseal))* → *unseal c*

In the function *mkdepkey* we call the existing seal creation function *mkseal*, which is used to generate a new seal that is paired with a fresh key identifier. Specifically, recall that we have:

type Seal(α) = $(\alpha \to \mathsf{Un}) * (\mathsf{Un} \to \alpha)$
val *mkseal:* $\mathsf{string} \to \mathsf{Seal}(\alpha)$

In the case of the key generation function *mkdepkey*, the placeholder α is replaced by the monomorphic type MsgAB<*id,A,B*>. Hence we must ensure that *id* is in scope when specializing the *mkseal* function.

Finally, we can briefly comment the other functions of our small library. The function *get_key_ident* extracts the identifier i from a dependent key k and tracks the logical dependence $\mathsf{KeyIdent}(k, i)$ through its refined return type. Contrary to standard sealing-based encryption and decryption, the functions *depencrypt* and *depdecrypt* take the key identifier as an additional argument and perform a pattern-matching operation to bridge the dependent typing allowed by pair splitting and the dependent typing enabled by the definition of these functions. In the syntax of types, the need for this pattern matching operation is made apparent by the occurrence of the same variable x in both the function type of *depencrypt*/*depdecrypt* and the data type MsgAB*<x,A,B>*.

9.5. Type-checking the initiator

We finally have all the ingredients to discuss how the initiator A is type-checked. The code of the principal looks as follows:

```
(∗ Authorization policy ∗)
assume !∀w, x, y, z.(Key(x, y, z) ⊗ KeyIdent(x, w) ⊗ Auth(w, y, z) ⊸ Session(x, y, z))

(∗ Typing the message from A to B, where MsgAB<x,y,z> will be
   closed by instantiating it in the definition of the session key type ∗)
type MsgAB<x,y,z> = MsgAB of {t : TStamp | !(F(y, t) ⊸ Auth(x, y, z))} + TStamp

(∗ Typing the session key established by Kerberos ∗)
type DSymKey = DSym of (w : string ∗ ((MsgAB<w,A,B> → Un)∗
                                      (Un → MsgAB<w,A,B>)))

(∗ Typing the message from S to A, where MsgSA<x> will be
   closed by instantiating it in the initiator function ∗)
type MsgSA<x> = MsgSA of (xts : TStamp ∗ xkAB : DSymKey ∗ xB : string ∗ y : Un)
  {!(F(xB, xts) ⊸ Key(xkAB, x, xB))}

(∗ Initiator code, where rand is a fresh global bitstring and
glob denotes a global reference, which are both provided in the
protocol specification and are not under the control of the opponent ∗)
let initiator rand glob A addA B addB S addS (kAS: SymKey(MsgSA<A>)) =
    let (get_tsB, check_tsB) = init_ts rand glob B in
      send addS (A,B);
      let msgSA = receive addA in
      let plainSA = sdecrypt kAS msgSA in
      match plainSA with
        MsgSA (xts, xkAB, =B, y) →
          (∗ !(F(B, xts) ⊸ Key(xkAB, A, B)) holds true ∗)
          let _ = check_tsB xts in
          (∗ F(B, xts) holds true ∗)
          let tA = get_tsB () in
          let iAB = get_key_ident xkAB in
          (∗ !KeyIdent(xkAB, iAB) holds true ∗)
          let msgAB = depencrypt iAB xkAB tA in
          send addB (y,msgAB);
          let msgBA = receive addA in
          (∗ tA' = tA + 1 ∗)
          let tA' = inc_ts tA in
          let (=tA') = depdecrypt iAB xkAB msgBA in
```

(∗ $!(\mathsf{F}(B, tA') \multimap \mathsf{Auth}(iAB, A, B))$ *holds true* ∗*)*
let _ = *check_tsB tA'* **in**
(∗ $\mathsf{F}(B, tA')$ *holds true* ∗*)*
assert $\mathsf{Session}(xkAB, A, B)$

Decryption and pattern-matching introduce the guarded formulas needed to type-check the initiator, while invocations to the timestamp library extend the typing environment with the control formulas needed to retrieve the payload formulas of interest. Specifically, the initiator starts by creating the handle to the timestamp library through the call *init_ts B*, which returns the two functions *get_tsB* and *check_tsB*. The interesting point here is the type of *check_tsB*, i.e., $y : \mathsf{TStamp} \rightarrow \{\_ : \mathsf{unit} \mid \mathsf{F}(B, y)\}$, hence a successful call to this function allows for deeming a communication with B as timely. To understand why the function is given that type, recall that *init_ts* has the following type:

$$\mathsf{bytes} \rightarrow \mathsf{Ref}(\mathsf{List}(\mathsf{string} * \mathsf{TStamp})) \rightarrow x : \mathsf{string} \rightarrow$$
$$((\mathsf{unit} \rightarrow \mathsf{TStamp}) * (y : \mathsf{TStamp} \rightarrow \{\_ : \mathsf{unit} \mid \mathsf{F}(x, y)\})).$$

and observe that *check_tsB* is obtained by projecting the second component of the pair returned by the call *init_ts rand glob B*. Now, the logical environment is populated as follows:

(i) when A decrypts the message from S and performs pattern matching, we introduce the formula $!(\mathsf{F}(B, xts) \multimap \mathsf{Key}(xkAB, A, B))$, based on the type of the symmetric key $kAS : \mathsf{SymKey}(\mathsf{MsgSA<A>})$;

(ii) when A calls the *check_tsB* function on the timestamp *xts* received by S, we introduce the formula $\mathsf{F}(B, xts)$, based on the typing discussed above;

(iii) when A calls the *get_key_ident* function on the self-dependent key *xkAB* shared with B, we introduce the formula $!\mathsf{KeyIdent}(xkAB, iAB)$, where *iAB* is the key identifier associated to *xkAB*;

(iv) when A decrypts the message from B using the self-dependent key *xkAB* identified by *iAB*, we introduce the formula $!(\mathsf{F}(B, tA') \multimap \mathsf{Auth}(iAB, A, B))$, based on the type of the *depdecrypt* function associated to *xkAB*, where tA' corresponds to tA incremented by 1;

(v) finally, when A calls the *check_tsB* function on the timestamp tA' received by B, we introduce the formula $\mathsf{F}(B, tA')$, similarly to what we do at point (ii).

Using (i) and (ii), we can prove $\mathsf{Key}(xkAB, A, B)$, while using (iv) and (v) we can prove $\mathsf{Auth}(iAB, A, B)$. These two formulas, along with $!\mathsf{KeyIdent}(xkAB, iAB)$ at point (iii), allow to derive the assertion $\mathsf{Session}(xkAB, A, B)$ based on the underlying authorization policy, hence the initiator is well-typed.

To conclude that the protocol actually respects the authorization policy despite the introduction of the serializers, it is enough to ensure that $\mathsf{F}(B, t)$ is assumed at most once for any possible choice of t. To prove it, we must guarantee that at the beginning of the protocol specification function:

(1) the global fresh value *rand* is freshly generated using the function *mkfresh* that never generates the same value twice;

(2) the global reference *glob* storing the received timestamps is correctly instantiated to the empty list and is not provided by the opponent as an argument to the protocol specification function. We thus note that different participants running with identity A share the same counter for timestamps management by construction of our library (cf. Section 9.3) and that each invocation to *check_tsB* always returns an assumption predicating over increasing values of t.

Table XVI Selected algorithmic rules for typing values and expressions

(VAL VAR ALG)
$$\frac{\Gamma \vdash_{\sf alg} \diamond \qquad (x : T) \in \Gamma}{\Gamma \vdash_{\sf alg} x : T; \mathbf{1}}$$

(VAL FUN ALG)
$$\frac{\Gamma, x : \psi(T) \vdash_{\sf alg} \overline{E} : U; F' \qquad fnfv(T) \subseteq dom(\Gamma) \cup \{x\}}{\Gamma \vdash_{\sf alg} \lambda x : T.\, \overline{E} : (x : T \to U); !\forall x.(forms(x : T) \multimap F')}$$

(VAL REF ALG)
$$\frac{\Gamma \vdash_{\sf alg} \overline{M} : T; F' \qquad fnfv(F) \subseteq dom(\Gamma) \cup \{x\}}{\Gamma \vdash_{\sf alg} \overline{M}_{\{x:\_ \mid F\}} : \{x : T \mid F\}; F' \otimes F\{M/x\}}$$

(VAL PAIR ALG)
$$\frac{\Gamma \vdash_{\sf alg} \overline{M} : T; F_1 \qquad \Gamma \vdash_{\sf alg} \overline{N} : U\{M/x\}; F_2}{\Gamma \vdash_{\sf alg} (\overline{M}, \overline{N}) : x : T * U; !F_1 \otimes !F_2}$$

(EXP LET ALG)
$$\frac{\overline{E} \rightsquigarrow^{\emptyset} [\Delta' \mid \overline{E'}]}{\Gamma \vdash_{\sf alg} \overline{E'} : T; F_1 \qquad \Gamma, x : \psi(T) \vdash_{\sf alg} \overline{D} : U; F_2 \qquad x \notin fv(U) \qquad fnfv(\Delta') \subseteq dom(\Gamma)}{\Gamma \vdash_{\sf alg} \mathsf{let}\ x = \overline{E}\ \mathsf{in}\ \overline{D} : U; \Delta' \multimap (F_1 \otimes \forall x.(forms(x : T) \multimap F_2))}$$

Notation: Here $E := \langle \overline{E} \rangle$ denotes the expression obtained from $\overline{E}$ by erasing all its typing annotations.

10. ALGORITHMIC TYPE-CHECKING

The type system presented in Section 6 includes several non-deterministic rules, which make it hard to implement an efficient decision procedure for typing. In this section we outline an algorithmic variant of the type system, which we prove sound and complete. We first focus on presenting the main intuitions behind the algorithmic type system design and then show the complete formalization.

10.1. Overview

While standard sources of non-determinism (like subtyping or refining value types) can be eliminated using type annotations, the rewriting of logical environments, which is the distinctive source of non-determinism of our system, is harder to deal with. The core idea underlying the algorithmic version of the type system is to dispense with the logical environment Δ and to construct bottom-up a single logical formula that characterizes all the proof obligations that would normally be introduced along the type derivation. In such a way, all the burden due to resource management can be shifted to an external affine logic theorem prover, which has to deal with this issue anyway.

More in detail, every typing judgement of the form $\Gamma; \Delta \vdash \mathcal{J}$ is matched by an algorithmic counterpart of the form $\Gamma \vdash_{\sf alg} \mathcal{J}; F$. Intuitively, typing an expression algorithmically constitutes of two steps:

(1) The expression (decorated with type annotations whenever needed) is type-checked using the algorithmic type system. This process is syntax-directed and fully deterministic, and in case of success yields *one* proof obligation F.
(2) The proof obligation is verified, e.g., using an external theorem prover.

If both steps succeed, then the expression is well-typed.

10.2. Key ideas

We illustrate the main ideas behind our algorithmic type system on some representative rules, shown in Table XVI. The algorithmic rules for kinding (cf. Section 10.4), subtyping (cf. Section 10.5), and typing the remaining values and expressions (cf. Section 10.6) follow along the same lines. For the sake of readability we often abuse notation and we let the multiset $F_1, \ldots, F_n$ stand for the formula $F_1 \otimes \ldots \otimes F_n$.

We first notice that, according to standard practice, we rely on typing annotations to deal with non-structural rules. Annotated terms and expressions are denoted by $\overline{M}$ and $\overline{E}$, respectively. Their syntax is given in Table XXI. The explicit erasure of all typing annotations of an expression is denoted by $\langle \overline{E} \rangle$. For instance, we explicitly annotate values that are expected to be given a refinement type (cf. (VAL REF ALG)) with the expected refinement F and use annotations to assign an explicit argument type T to functional values (cf. (VAL FUN ALG)). In this way, every possible syntactic form for expressions is matched by a single type rule and the selection of appropriate types and refinements does not rely on non-determinism.

We now exemplify the general concepts underlying our technique by contrasting the standard typing rule (VAL FUN) with its algorithmic counterpart (VAL FUN ALG). The main source of non-determinism in (VAL FUN) is the rewriting of Δ to $!\Delta'$. As previously mentioned, our goal is to dispense with logical environments and their rewriting, by collecting a single proof obligation that accounts for the proof obligations generated in the original type system. In the algorithmic version, the proof obligation obtained by giving $\lambda x : T. \overline{E}$ type $V = x : T \to U$ in the environment Γ is:

$$!\forall x.(forms(x : T) \multimap F'),$$

where F' is the proof obligation collected by giving $\overline{E}$ type U in $\Gamma, x : \psi(T)$.

In the following, we briefly justify why this approach is sound, i.e., we argue why $\Gamma \vdash_{\mathsf{alg}} \lambda x : T. \overline{E} : V; !\forall x.(forms(x : T) \multimap F')$ implies that $\Gamma; \Delta \vdash \lambda x. \langle \overline{E} \rangle : V$ for any Δ such that $\Gamma; \Delta \vdash !\forall x.(forms(x : T) \multimap F')$. Notice that the latter judgement is equivalent to assuming that Δ entails $!\forall x.(forms(x : T) \multimap F')$ and both the multiset and the formula are well-formed with respect to Γ. Using the rules of the logic, we can show that a proof of $\Gamma; \Delta \vdash !\forall x.(forms(x : T) \multimap F')$ implies that there exists Δ' such that $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta'$ and:

$$\Gamma; !\Delta' \vdash \forall x.forms(x : T) \multimap F'.$$

Intuitively, this means that we can eliminate the exponential modality by rewriting the logical environment in exponential form. Furthermore, the well-formedness of the (algorithmic) environment $\Gamma, x : \psi(T)$ and the (non-algorithmic) environment $\Gamma; !\Delta'$ ensures that $x \notin dom(\Gamma)$ and thus $x \notin fv(!\Delta')$: in this case, the logic allows us to further eliminate the universal quantification, adding a type binding for x in order to keep the logical environment well-formed (the actual type is not relevant from the logic point of view). Thus, we have:

$$\Gamma, x : \psi(T); !\Delta' \vdash forms(x : T) \multimap F'.$$

Using rule ($\multimap$-LEFT), we can finally prove:

$$\Gamma, x : \psi(T); !\Delta', forms(x : T) \vdash F'.$$

By inductive reasoning, $\Gamma, x : \psi(T); !\Delta', forms(x : T) \vdash \langle \overline{E} \rangle : U$, hence (VAL FUN) allows us to derive $\Gamma; \Delta \vdash \lambda x. \langle \overline{E} \rangle : V$. The proof of completeness is similar.

The other algorithmic (typing) rules are constructed along the same lines, using the following additional observations:

— If a typing rule contains no kinding, subtyping, or typing premise (e.g., (VAL VAR)), the proof obligation of the corresponding algorithmic rule is set to $\mathbf{1}$ (cf. (VAL VAR ALG)) and thus trivially fulfilled.
— If a typing rule contains multiple premises (e.g., (VAL PAIR)), then we combine the proof obligations obtained from the premises conjunctively (cf. (VAL PAIR ALG)).
— If a typing rule relies on extraction (e.g., (EXP LET)) and adds the extracted environment Δ' to the environment before rewriting, the algorithmic variant of the rule (EXP

LET ALG) creates a proof obligation of the form $\Delta' \multimap F$, where F is the proof obligation obtained by combining the proof obligations of the premises using the techniques described above.

With these insights in mind, we now show the complete formalization of the algorithmic type system.

10.3. Base judgements

The base judgements of the algorithmic type system are reported in Table XVII.

Table XVII Algorithmic well-formedness rules

$$
\text{(ENV EMPTY ALG)} \quad \dfrac{\text{(TYPE ENV ENTRY ALG)}}{\Gamma \vdash_{\mathsf{alg}} \diamond \quad dom(\mu) \cap dom(\Gamma) = \emptyset} \quad \dfrac{\text{(TYPE ALG)}}{\Gamma \vdash_{\mathsf{alg}} \diamond \quad fnfv(T) \subseteq dom(\Gamma)}
$$

$$
\dfrac{\varepsilon \vdash_{\mathsf{alg}} \diamond}{\qquad} \quad \dfrac{\mu = x : T \Rightarrow T = \psi(T) \wedge fnfv(T) \subseteq dom(\Gamma)}{\Gamma, \mu \vdash_{\mathsf{alg}} \diamond} \quad \dfrac{}{\Gamma \vdash_{\mathsf{alg}} T}
$$

The only remarkable point here is that we do not have any algorithmic counterpart of rule (DERIVE). In fact, we never need to prove a formula in the algorithmic formulation of the type system, but we just collect the proof obligation for the external affine logic theorem prover.

10.4. Kinding

Table XVIII presents the algorithmic kinding rules. The non-inductive standard kinding rules (KIND VAR) and (KIND UNIT), which just check well-formedness of the environment (or environment membership) and which do not contain a proof obligation of the form $\Gamma; \Delta \vdash F$ amongst their hypotheses, are translated into algorithmic rules that generate the proof obligation 1. All other (recursive) rules (e.g., (KIND FUN)) strongly resemble their algorithmic counterparts (e.g., (KIND FUN ALG)). The proof obligation that is generated in the algorithmic variant consists of a conjunction of the proof obligations that are recursively generated by the premises of that rule, following the same principles of the algorithmic typing rules for values and expressions that we discussed in Section 10.2. Note that a premise that checks the well-formedness of an environment or type does not generate a proof obligation (cf. (KIND REFINE PUBLIC ALG)).

Table XVIII Algorithmic kinding rules

$$
\text{(KIND FUN ALG)}
$$
$$
\text{(KIND VAR ALG)} \qquad \text{(KIND UNIT ALG)} \qquad \dfrac{\Gamma \vdash_{\mathsf{alg}} T :: \overline{k}; F_1}{\Gamma, x : \psi(T) \vdash_{\mathsf{alg}} U :: k; F_2}
$$
$$
\dfrac{\Gamma \vdash_{\mathsf{alg}} \diamond \quad (\alpha :: k) \in \Gamma}{\Gamma \vdash_{\mathsf{alg}} \alpha :: k; \mathbf{1}} \qquad \dfrac{\Gamma \vdash_{\mathsf{alg}} \diamond}{\Gamma \vdash_{\mathsf{alg}} \mathsf{unit} :: k; \mathbf{1}} \qquad \dfrac{}{\Gamma \vdash_{\mathsf{alg}} x : T \to U :: k; !F_1 \otimes !F_2}
$$

$$
\text{(KIND PAIR ALG)} \qquad \text{(KIND SUM ALG)}
$$
$$
\dfrac{\Gamma \vdash_{\mathsf{alg}} T :: k; F_1}{\Gamma, x : \psi(T) \vdash_{\mathsf{alg}} U :: k; F_2} \qquad \dfrac{\Gamma \vdash_{\mathsf{alg}} T :: k; F_1}{\Gamma \vdash_{\mathsf{alg}} U :: k; F_2} \qquad \text{(KIND REC ALG)}
$$
$$
\dfrac{}{\Gamma \vdash_{\mathsf{alg}} x : T * U :: k; !F_1 \otimes !F_2} \qquad \dfrac{}{\Gamma \vdash_{\mathsf{alg}} T + U :: k; !F_1 \otimes !F_2} \qquad \dfrac{\Gamma, \alpha :: k \vdash_{\mathsf{alg}} T :: k; F}{\Gamma \vdash_{\mathsf{alg}} \mu\alpha. T :: k; !F}
$$

$$
\text{(KIND REFINE PUBLIC ALG)} \qquad \text{(KIND REFINE TAINTED ALG)}
$$
$$
\dfrac{\Gamma \vdash_{\mathsf{alg}} \{x : T \mid F\} \quad \Gamma \vdash_{\mathsf{alg}} T :: \mathsf{pub}; F'}{\Gamma \vdash_{\mathsf{alg}} \{x : T \mid F\} :: \mathsf{pub}; F'} \qquad \dfrac{\Gamma \vdash_{\mathsf{alg}} \psi(T) :: \mathsf{tnt}; F' \quad \Gamma, x : \psi(T) \vdash_{\mathsf{alg}} \diamond \quad T \text{ refined}}{\Gamma \vdash_{\mathsf{alg}} \{x : T \mid F\} :: \mathsf{tnt}; (\forall x. forms(x : T)) \otimes F'}
$$

10.5. Subtyping

The algorithmic subtyping rules are presented in Table XX. They resolve the non-determinism related to the environment splitting by following the key insights of algorithmic typing presented in Section 10.2.

Furthermore, the algorithmic subtyping rules resolve the non-determinism that arises due to the fact that standard subtyping is not syntax-driven as described in the following. We use $T \neq_\top U$ to denote that T and U are not refined and do not share the same top-level constructor.

The algorithmic type system makes use of the following observation: for all non-refined types T, U there are at most three standard subtyping rules applicable, namely (SUB REFL), (SUB PUB TNT), and in the case that T and U share the same top-level constructor one corresponding structural subtyping rule, e.g., (SUB FUN) or (SUB PAIR). In the case that T or U are refined, the three standard subtyping rules (SUB REFL), (SUB PUB TNT), or (SUB REFINE) might be applicable.

To reduce this level of non-determinism the algorithmic subtyping rules allow the reflexivity rule (SUB REFL ALG) to be applied only to the non-inductive type unit and type variables α. Furthermore, we restrict the application of the kinding based rule (SUB PUB TNT ALG) to types T, U that are structurally different and not refined, i.e., $T \neq_\top U$. Therefore, we can determine the appropriate subtyping rule by simple syntactic checks. Note that two types T, U, which share the same top-level constructor can still be subtyped using reflexivity or kinding by recursively applying the corresponding structural subtyping rule until one of the subgoals matches the premise of either the (SUB REFL ALG) or (SUB PUB TNT ALG) rule. Similarly, if either T or U or both are refined they can be typed using reflexivity or kinding by first applying the refinement rule (SUB REFINE ALG) and then applying either the (SUB REFL ALG) or (SUB PUB TNT ALG) rule to the subgoal.

This approach is sound and complete for all but the subtyping of two iso-recursive types. This is related to our choice of adapting the iso-recursive subtyping proposed by Backes et al. [Backes et al. 2011], which requires the recursive variable to occur only positively in the iso-recursive type, instead of the Amber rule (cf. (SUB POS REC) in Section 6.4). For instance, given the above constraints, subtyping $\Gamma \vdash_{\mathsf{alg}} \mu\alpha.\,(x : \alpha \to T) <: \mu\alpha.\,(x : \alpha \to T); F$ or $\Gamma \vdash_{\mathsf{alg}} \mu\alpha.\,(x : \alpha \to \mathsf{unit}) <: \mu\alpha.\,(x : \alpha \to \mathsf{unit} + \mathsf{unit}); F$ would not be possible, thus lacking reflexivity and kinding based algorithmic subtyping for iso-recursive types. Therefore, our algorithmic type system contains three rules for subtyping two iso-recursive types: (SUB REFL REC ALG), (SUB PUB TNT REC ALG), and (SUB POS REC ALG), respectively. While checking whether or not to apply rule (SUB REFL REC ALG) can be done by performing a simple equality check on the types, the decision between (SUB PUB TNT REC ALG) and (SUB POS REC ALG) requires some guidance, leading to the introduction of manual annotations of the form SPT to denote that the rule (SUB PUB TNT REC ALG) should be applied. This annotation appears in the subtyping rule for expressions (cf. (EXP SUBSUM ALG)), which we explain in Section 10.6.

The syntax of annotated types $\overline{T}$ is introduced in Table XIX. Intuitively, we allow type annotations SPT only on iso-recursive types and require them to not be nested. We let $\langle \overline{T} \rangle$ denote the explicit erasure of all annotations SPT from an annotated type $\overline{T}$. To facilitate readability we often write T to denote the non-annotated counterpart $\langle \overline{T} \rangle$ of the annotated type $\overline{T}$. We can easily extend the definition of the function ψ (used for the removal of top-level refinements) to annotated types.

Table XIX Syntax of annotated types and annotation erasure

$$
\begin{array}{lll}
\overline{T}, \overline{U}, \overline{V} ::= & & \textit{annotated types} \\
& \text{unit} & \text{unit} \\
& \alpha & \text{type variable} \\
& x : \overline{T} \to \overline{U} & \text{dependent function type (scope of } x \text{ is } \overline{U}) \\
& x : \overline{T} * \overline{U} & \text{dependent pair type (scope of } x \text{ is } \overline{U}) \\
& \overline{T} + \overline{U} & \text{sum type} \\
& \mu\alpha.\,\overline{T} & \text{iso-recursive type without top-level annotation} \\
& (\mu\alpha.\,T)_{\mathsf{SPT}} & \text{iso-recursive type with top-level annotation} \\
& \{x : \overline{T} \mid F\} & \text{refinement type (scope of } x \text{ is } F)
\end{array}
$$

$$
\psi(\overline{U}) = \begin{cases} \psi(\overline{T}) & \text{if } \overline{U} = \{x : \overline{T} \mid F\} \\ \overline{U} & \text{otherwise} \end{cases}
$$

$$
\langle \overline{U} \rangle = \begin{cases}
\text{unit} & \text{if } \overline{U} = \text{unit} \\
\alpha & \text{if } \overline{U} = \alpha \\
x : \langle \overline{U_1} \rangle \to \langle \overline{U_2} \rangle & \text{if } \overline{U} = x : \overline{U_1} \to \overline{U_2} \\
x : \langle \overline{U_1} \rangle * \langle \overline{U_2} \rangle & \text{if } \overline{U} = x : \overline{U_1} * \overline{U_2} \\
\langle \overline{U_1} \rangle + \langle \overline{U_2} \rangle & \text{if } \overline{U} = \overline{U_1} + \overline{U_2} \\
\mu\alpha.\,\langle \overline{T} \rangle & \text{if } \overline{U} = \mu\alpha.\,\overline{T} \\
\mu\alpha.\,T & \text{if } \overline{U} = (\mu\alpha.\,T)_{\mathsf{SPT}} \\
\{x : \langle \overline{T} \rangle \mid F\} & \text{if } \overline{U} = \{x : \overline{T} \mid F\}
\end{cases}
$$

10.6. Typing values and expressions

The algorithmic typing rules for values and expressions are given in Table XXII and Table XXIII, respectively. The rules follow according to the intuition described in Section 10.2. We furthermore rely on type annotations to guide the selection of applicable typing rules and appropriate types. The syntax of annotated values and expressions is given in Table XXI. Here "_" is used to denote a type that is derived by the typing rules and thus does not need to be specified by the annotator. We denote the recursive erasure of all typing annotations by $\langle \overline{E} \rangle$ and often use E to denote the expression $\langle \overline{E} \rangle$ obtained from the annotated expression $\overline{E}$ by erasing all its typing annotations. The extraction relation $\overline{E} \rightsquigarrow^{\emptyset} [\Delta \mid \overline{D}]$ for annotated expressions (cf. Table XXIV) extracts formulas as in the non-annotated case while keeping annotations on the expressions intact but for the case of assumptions, where it changes the type annotation in the original assumption to a subtyping annotation in the extracted assumption for all types different from unit. The notions of free names and free variables correspond to the non-annotated case.

Since in the typing rule (VAL FUN) for functions the type of the input is chosen non-deterministically, we use the annotation $\lambda x : T.\,E$ to guide the algorithmic type system (VAL FUN ALG) in the selection of a suitable input type T. The annotation $M_{\{x:\_ \mid F\}}$ explicitly triggers the rule (VAL REF ALG) and expects M to type-check with refinement F, while the annotations $(\text{inl } M)_{\_+U}$ and $(\text{inl } M)_{T+\_}$ are used to provide the respective missing type in the sum type $T + U$ that will be assigned to inl M and inr M (cf. (VAL INL ALG) and (VAL INR ALG)). Furthermore, the rule (EXP SUBSUM) is highly non-deterministic, since its application can be tried at any time using any combination of possible sub- and supertypes. In the algorithmic version of the type system we prevent the unnecessary application of subtyping and help the choice of an appropriate supertype T' by annotating an expression $\overline{E}$ as $\overline{E}_{\_<:\overline{T'}}$ whenever subtyping is necessary

Table XX Algorithmic subtyping rules

(SUB REFL ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} T \qquad T \in \{\mathsf{unit}, \alpha\}}{\Gamma \vdash_{\mathsf{alg}} T <: T; \mathbf{1}}$$

(SUB PUB TNT ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} T :: \mathsf{pub}; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} U :: \mathsf{tnt}; F_2 \qquad T \neq_\top U}{\Gamma \vdash_{\mathsf{alg}} T <: U; F_1 \otimes F_2}$$

(SUB FUN ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{T'} <: \overline{T}; F_1 \qquad \Gamma, x : \psi(T') \vdash_{\mathsf{alg}} \overline{U} <: \overline{U'}; F_2}{\Gamma \vdash_{\mathsf{alg}} x : \overline{T} \to \overline{U} <: x : \overline{T'} \to \overline{U'}; !F_1 \otimes !F_2}$$

(SUB PAIR ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{T'}; F_1 \qquad \Gamma, x : \psi(T) \vdash_{\mathsf{alg}} \overline{U} <: \overline{U'}; F_2}{\Gamma \vdash_{\mathsf{alg}} x : \overline{T} * \overline{U} <: x : \overline{T'} * \overline{U'}; !F_1 \otimes !F_2}$$

(SUB SUM ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{T'}; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} \overline{U} <: \overline{U'}; F_2}{\Gamma \vdash_{\mathsf{alg}} \overline{T} + \overline{U} <: \overline{T'} + \overline{U'}; !F_1 \otimes !F_2}$$

(SUB POS REC ALG)
$$\frac{\Gamma, \alpha \vdash_{\mathsf{alg}} T <: T'; F \qquad \overline{T} \neq \overline{T'} \qquad \alpha \text{ occurs only positively in } \overline{T} \text{ and } \overline{T'}}{\Gamma \vdash_{\mathsf{alg}} \mu\alpha. \overline{T} <: \mu\alpha. \overline{T'}; !F}$$

(SUB REFL REC ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \mu\alpha. T}{\Gamma \vdash_{\mathsf{alg}} \mu\alpha. T <: \mu\alpha. T; \mathbf{1}}$$

(SUB PUB TNT REC ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \mu\alpha. T :: \mathsf{pub}; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} \mu\alpha. U :: \mathsf{tnt}; F_2 \qquad s = \mathsf{SPT} \oplus s' = \mathsf{SPT}}{\Gamma \vdash_{\mathsf{alg}} (\mu\alpha. T)_s <: (\mu\alpha. T')_{s'}; F_1 \otimes F_2}$$

(SUB REFINE ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \psi(\overline{T}) <: \psi(\overline{U}); F \qquad \overline{T} \text{ and/or } \overline{U} \text{ refined} \qquad \Gamma \vdash_{\mathsf{alg}} T \qquad \Gamma \vdash_{\mathsf{alg}} U}{\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{U}; F \otimes \forall y.(forms(y : T) \multimap forms(y : U))}$$

Notation: We write $T \neq_\top U$ to denote that T and U are not refined and do not share the same top-level constructor. $\oplus$ denotes the exclusive or. We use T to denote the non-annotated counterpart $\langle \overline{T} \rangle$ of the annotated type $\overline{T}$.

(cf. (EXP SUBSUM ALG)). Note that the type T' will additionally be annotated with SPT in case that the subtyping should make use of rule (SUB PUB TNT REC ALG), resulting in the annotated type $\overline{T'}$. Since the typing rule (EXP ASSUME) non-deterministically chooses a type T, its algorithmic counterpart (EXP ASSUME ALG) requires an explicit annotation of the form $(\mathsf{assume}\ F)_T$ to provide the expected type T.

10.7. Formal results

We can state and prove the following formal results, which highlight the correctness and the accuracy of the algorithmic type system.

THEOREM 10.1 (SOUNDNESS OF ALGORITHMIC TYPING). *If* $\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F$ *and* $\Gamma; \Delta \vdash F$*, then* $\Gamma; \Delta \vdash \langle \overline{E} \rangle : T$.

PROOF. See Appendix C. □

THEOREM 10.2 (COMPLETENESS OF ALGORITHMIC TYPING). *If* $\Gamma; \Delta \vdash E : T$*, then there exist* $\overline{E}, F$ *such that* $\langle \overline{E} \rangle = E$ *and* $\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F$ *and* $\Gamma; \Delta \vdash F$.

PROOF. See Appendix C. □

10.8. Example

The proof obligation assigned to the *cust* function in Section 8 by the algorithmic formulation of our type system is shown below:

$$\forall C. \forall M. \forall B. \forall g. \forall p.$$

Table XXI Syntax of annotated RCF expressions

$\overline{M}, \overline{N} ::=$		*values*
	x	variable
	$()$	unit
	$(\overline{M}, \overline{N})$	pair
	$\lambda x : T. \overline{E}$	annotated function with input of type T
	$(\text{inl } \overline{M})_{\_+T}$	annotated left constructor
	$(\text{inr } \overline{M})_{T+\_}$	annotated right constructor
	$\text{fold } \overline{M}$	fold constructor
	$\overline{M}_{\{x:\_ \mid F\}}$	value to be refined with F
	$\overline{M}_{\_<:\overline{T}}$	value to be subtyped to $\overline{T}$
$\overline{D}, \overline{E} ::=$		*expressions*
	$\overline{M}$	value
	$\overline{M} \ \overline{N}$	application
	$\overline{M} = \overline{N}$	syntactic equality
	$\text{let } x = \overline{E} \text{ in } \overline{E}'$	let (scope of x is $\overline{E}'$)
	$\text{let } (x, y) = \overline{M} \text{ in } \overline{E}$	pair split (scope of x, y is $\overline{E}$)
	$\text{match } \overline{M} \text{ with } h \ x \text{ then } \overline{E} \text{ else } \overline{E}'$	match (scope of x is $\overline{E}$)
	$(\nu a)\overline{E}$	restriction (scope of a is $\overline{E}$)
	$\overline{E} \rhd \overline{E}'$	fork
	$a!\overline{M}$	message send
	$a?$	message receive
	$\text{assume } 1$	non-annotated truth assumption
	$(\text{assume } F)_T$	annotated assumption with expected type T
	$\text{assert } F$	assertion
	$\overline{E}_{\_<:\overline{T}}$	expression to be subtyped to $\overline{T}$

Table XXII Algorithmic typing rules for values

(VAL VAR ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \diamond \qquad (x : T) \in \Gamma}{\Gamma \vdash_{\mathsf{alg}} x : T; \mathbf{1}}$$

(VAL UNIT ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \diamond}{\Gamma \vdash_{\mathsf{alg}} () : \mathsf{unit}; \mathbf{1}}$$

(VAL FUN ALG)
$$\frac{\Gamma, x : \psi(T) \vdash_{\mathsf{alg}} \overline{E} : U; F' \qquad fnfv(T) \subseteq dom(\Gamma) \cup \{x\}}{\Gamma \vdash_{\mathsf{alg}} \lambda x : T. \overline{E} : x : T \to U; !\forall x.(forms(x : T) \multimap F')}$$

(VAL PAIR ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} \overline{N} : U\{M/x\}; F_2}{\Gamma \vdash_{\mathsf{alg}} (\overline{M}, \overline{N}) : x : T * U; !F_1 \otimes !F_2}$$

(VAL REF ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T; F' \qquad fnfv(F) \subseteq dom(\Gamma) \cup \{x\}}{\Gamma \vdash_{\mathsf{alg}} \overline{M}_{\{x:\_ \mid F\}} : \{x : T \mid F\}; F' \otimes F\{M/x\}}$$

(VAL INL ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T; F' \qquad \Gamma \vdash_{\mathsf{alg}} U}{\Gamma \vdash_{\mathsf{alg}} (\text{inl } \overline{M})_{\_+U} : T + U; !F'}$$

(VAL INR ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : U; F' \qquad \Gamma \vdash_{\mathsf{alg}} T}{\Gamma \vdash_{\mathsf{alg}} (\text{inr } \overline{M})_{T+\_} : T + U; !F'}$$

(VAL FOLD ALG)
$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T\{\mu a. T/\alpha\}; F'}{\Gamma \vdash_{\mathsf{alg}} \text{fold } \overline{M} : \mu\alpha. T; !F'}$$

Notation: Here $M = \langle \overline{M} \rangle$ denotes the value obtained from $\overline{M}$ by erasing all its typing annotations.

$$\forall nC.((\mathsf{N1}(nC) \otimes \mathsf{N2}(nC)) \multimap$$
$$\forall xnM.(!(\mathsf{N1}(nC) \multimap (\forall y.\mathsf{Pay}(y, p, M, xnM) \multimap \mathsf{Ship}(M, g, C))) \multimap$$
$$!(\mathsf{N2}(nC) \multimap (\forall z.\mathsf{Pay}(B, p, z, xnM))) \multimap$$
$$\mathsf{Ship}(M, g, C)))$$

For the sake of readability we removed all unnecessary occurrences of $\mathbf{1}$ and all unused quantified variables. In this example, as well as in the other protocol we considered, the problem of solving equalities is reduced to the unification of variables. This allows us to use the `llprover` [Tomura 1995] theorem prover, which at the time of writing does not support equality theories. The above formula is discharged in less than 20 ms.

Table XXIII Algorithmic typing rules for expressions

(EXP SUBSUM ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} T <: \overline{T'}; F_2}{\Gamma \vdash_{\mathsf{alg}} \overline{E}_{\_<:\overline{T'}} : T'; F_1 \otimes F_2}$$

(EXP APPL ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : x : T \to U; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} \overline{N} : T; F_2}{\Gamma \vdash_{\mathsf{alg}} \overline{M}\ \overline{N} : U\{N/x\}; F_1 \otimes F_2}$$

(EXP LET ALG)

$$\frac{\overline{E} \rightsquigarrow^{\emptyset} [\Delta' \mid \overline{E'}] \qquad \Gamma \vdash_{\mathsf{alg}} \overline{E'} : T; F_1 \qquad \Gamma, x : \psi(T) \vdash_{\mathsf{alg}} \overline{D} : U; F_2 \qquad x \notin fv(U) \qquad fnfv(\Delta') \subseteq dom(\Gamma)}{\Gamma \vdash_{\mathsf{alg}} \mathsf{let}\ x = \overline{E}\ \mathsf{in}\ \overline{D} : U; \Delta' \multimap (F_1 \otimes \forall x.(forms(x : T) \multimap F_2))}$$

(EXP SPLIT ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : x : T * U; F_1 \qquad \Gamma, x : \psi(T), y : \psi(U) \vdash_{\mathsf{alg}} \overline{E} : V; F_2 \qquad \{x, y\} \cap fv(V) = \emptyset}{\Gamma; \Delta \vdash_{\mathsf{alg}} \mathsf{let}\ (x, y) = \overline{M}\ \mathsf{in}\ \overline{E} : V; F_1 \otimes \forall x.\forall y.(forms(x : T) \otimes forms(y : U) \otimes !((x, y) = M) \multimap F_2)}$$

(EXP MATCH ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T; F_1 \qquad \Gamma, x : \psi(H) \vdash_{\mathsf{alg}} \overline{E} : U; F_2 \qquad \Gamma; \Delta_2 \vdash_{\mathsf{alg}} \overline{D} : U; F_3}{(h, H, T) \in \{(\mathsf{inl}, T_1, T_1 + T_2), (\mathsf{inr}, T_2, T_1 + T_2), (\mathsf{fold}, T'\{\mu\alpha. T'/\alpha\}, \mu\alpha. T')\} \qquad fnfv(H) \subseteq dom(\Gamma) \cup \{x\}}$$
$$\frac{}{\Gamma; \Delta \vdash_{\mathsf{alg}} \mathsf{match}\ \overline{M}\ \mathsf{with}\ h\ x\ \mathsf{then}\ \overline{E}\ \mathsf{else}\ \overline{D} : U; F_1 \otimes \forall x.(forms(x : H) \otimes !(h\ x = M) \multimap F_2) \otimes F_3}$$

(EXP EQ ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T; F_1 \qquad \Gamma \vdash_{\mathsf{alg}} \overline{N} : U; F_2 \qquad x \notin (fv(\overline{M}) \cup fv(\overline{N}))}{\Gamma \vdash_{\mathsf{alg}} \overline{M} = \overline{N} : \{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = N)\}; F_1 \otimes F_2}$$

(EXP ASSUME ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} (\mathsf{assume}\ \mathbf{1})_{\_<:T} : T; F' \qquad F \neq \mathbf{1} \qquad fnfv(F) \subseteq dom(\Gamma)}{\Gamma \vdash_{\mathsf{alg}} (\mathsf{assume}\ F)_T : T; F \multimap F'}$$

(EXP TRUE ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \diamond}{\Gamma \vdash \mathsf{assume}\ \mathbf{1} : \mathsf{unit}; \mathbf{1}}$$

(EXP ASSERT ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \diamond \qquad fnfv(F) \subseteq dom(\Gamma)}{\Gamma \vdash_{\mathsf{alg}} \mathsf{assert}\ F : \mathsf{unit}; F}$$

(EXP RES ALG)

$$\frac{\overline{E} \rightsquigarrow^{a} [\Delta' \mid \overline{E'}] \qquad \Gamma, a \updownarrow T \vdash_{\mathsf{alg}} \overline{E'} : U; F \qquad a \notin fn(U) \qquad fnfv(\Delta') \subseteq dom(\Gamma)}{\Gamma \vdash_{\mathsf{alg}} (\nu a \updownarrow T)\overline{E} : U; \Delta' \multimap F}$$

(EXP SEND ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \overline{M} : T; F \qquad (a \updownarrow T) \in \Gamma}{\Gamma \vdash_{\mathsf{alg}} a!\overline{M} : \mathsf{unit}; F}$$

(EXP RECV ALG)

$$\frac{\Gamma \vdash_{\mathsf{alg}} \diamond \qquad (a \updownarrow T) \in \Gamma}{\Gamma \vdash_{\mathsf{alg}} a? : T; \mathbf{1}}$$

(EXP FORK ALG)

$$\frac{\overline{E_1} \rightsquigarrow^{\emptyset} [\Delta_1 \mid \overline{D_1}]}{\overline{E_2} \rightsquigarrow^{\emptyset} [\Delta_2 \mid \overline{D_2}] \qquad \Gamma \vdash_{\mathsf{alg}} \overline{D_1} : T_1; F_A \qquad \Gamma \vdash_{\mathsf{alg}} \overline{D_2} : T_2; F_B \qquad fnfv(\Delta_1, \Delta_2) \subseteq dom(\Gamma)}{\Gamma \vdash_{\mathsf{alg}} \overline{E_1} \upharpoonright \overline{E_2} : T_2; (\Delta_1, \Delta_2) \multimap (F_A \otimes F_B)}$$

Notation: Here $E = \langle\overline{E}\rangle$ denotes the expression obtained from $\overline{E}$ by erasing all its typing annotations.

11. RELATED WORK

Several papers develop type systems for (variants of) RCF [Bhargavan et al. 2010; Bengtson et al. 2011; Fournet et al. 2011; Backes et al. 2011; Swamy et al. 2011] but, with the exception of F* [Swamy et al. 2011], they do not support resource-aware policies: in fact, even for simple linearity properties like injective agreement they rely on hand-written proofs [Bhargavan et al. 2009].

F* [Swamy et al. 2011] is a dependently typed functional language for secure distributed programming, featuring refinement types to reason about authorization poli-

Table XXIV The extraction relation for annotated expressions

(Extr Fork)	(Extr Let)

$$\frac{\overline{E_1} \leadsto^{\widetilde{a}} [\Delta_1 \mid \overline{D_1}] \qquad \overline{E_2} \leadsto^{\widetilde{a}} [\Delta_2 \mid \overline{D_2}]}{(\overline{E_1} \upharpoonright \overline{E_2})_s \leadsto^{\widetilde{a}} [\Delta_1, \Delta_2 \mid (\overline{D_1} \upharpoonright \overline{D_2})_s]}$$

$$\frac{\overline{E_1} \leadsto^{\widetilde{a}} [\Delta \mid \overline{D_1}]}{(\text{let } x = \overline{E_1} \text{ in } \overline{E_2})_s \leadsto^{\widetilde{a}} [\Delta \mid (\text{let } x = \overline{D_1} \text{ in } \overline{E_2})_s]}$$

(Extr Res)

$$\frac{\overline{E} \leadsto^{a,\widetilde{b}} [\Delta \mid \overline{D}]}{((\nu a)\overline{E})_s \leadsto^{\widetilde{b}} [\Delta \mid ((\nu a)\overline{D})_s]}$$

(Extr Assume Unit)

$$\frac{F \neq \mathbf{1} \qquad fn(F) \cap \{\widetilde{a}\} = \emptyset}{((\text{assume } F)_{\text{unit}})_s \leadsto^{\widetilde{a}} [F \mid (\text{assume } \mathbf{1})_s]}$$

(Extr Assume)

$$\frac{F \neq \mathbf{1} \qquad fn(F) \cap \{\widetilde{a}\} = \emptyset \qquad T \neq \text{unit}}{((\text{assume } F)_T)_s \leadsto^{\widetilde{a}} [F \mid ((\text{assume } \mathbf{1})_{\_<:T})_s]}$$

(Extr Exp)

$$\frac{\text{no other rule applies}}{\overline{E} \leadsto^{\widetilde{a}} [\emptyset \mid \overline{E}]}$$

Remark: Note that here s is either SPT or ϵ (i.e., no annotation).

cies and affine types to reason about stateful computations on affine *values*. Similarly to companion proposals for RCF, however, the type system of F^* assumes the existence of the contraction rule in the underlying logic, hence it does not support authorization policies built over affine *formulas*. While some simple authentication patterns (e.g., basic nonce handshakes) may certainly be expressed by encoding affine predicates in terms of affine values, other more complex authentication mechanisms are much harder to handle in these terms. The *EPMO* protocol we analyze in Section 8 provides one such case, as (i) the nonce it employs may not be construed as an affine value because it is used twice, and (ii) the logical formulas justified by cryptographic message exchanges are more structured than simple predicates. Though it might be possible to come up with sophisticated encodings of these authentication mechanisms in the programming language (by resorting to, e.g., pairs of affine tokens to encode a double usage of the same nonce and special functions to eliminate logical implications), such encodings are hard to formulate in a general manner and, we argue, are much better expressed in terms of policy annotations than in some ad-hoc programming pattern.

Bhargavan et al. [Bhargavan et al. 2008] propose a technique for the verification of F# protocol implementations by automatically extracting ProVerif models [Blanchet 2001], using an extension of the functions-as-processes encoding proposed by Milner [Milner 1992]. Remarkably, the analysis can deal with injective agreement. On the other hand, the analysis carried out with ProVerif is not modular and has been shown less robust and scalable than type-checking [Bhargavan et al. 2010]. Furthermore, the fragment of F# considered is rather restrictive: for instance, it does not include higher-order functions and admits only very limited uses of recursion and state.

A formal account on the integration of refinement types and substructural logics was first proposed by Mandelbaum et al. [Mandelbaum et al. 2003] with a system for local reasoning about program state built around a fragment of intuitionistic linear logic. Later, Bierhoff and Aldrich developed a framework for modular type-state checking of object-oriented programs [Bierhoff and Aldrich 2007; Sunshine et al. 2011; Naden et al. 2012]. However, none of these systems deals with the presence of hostile (or untyped) program components, or attackers, a feature that is instead distinctive of our system: adapting the previous frameworks to take into account interactions with an untyped context would require fundamental changes to their typing rules. The original RCF type-checker [Bengtson et al. 2011], for instance, employs a security-oriented kinding relation to reason about messages sent to and received from the attacker, which we also adopt in our type system. Recent variants of the RCF type-checker dispense with the kinding relation and even with concurrency [Swamy et al. 2011], but they rely on

manually proven logical invariants capturing security properties of the cryptographic library and, in some cases, of the protocol itself.

Tov and Pucella [Tov and Pucella 2010] have recently shown how to use behavioral contracts to link code written in an affine language to code written in a conventionally typed language. The idea is to coerce affine values to non-affine ones that can be shared with the context, but can still be reasoned about safely using dynamic access counts. There are intriguing similarities between this approach and the usage of nonces and session keys to enforce linearity properties in an adversarial setting, which are worth to be investigated in the future. The two type systems are, however, fundamentally different, since our present work deals with an affine refinement logic and an adversarial setting, which makes a precise comparison hard to formulate.

Various techniques have been proposed to statically analyze authenticity properties of cryptographic protocols [Armando et al. 2005; Backes et al. 2007a; Cremers 2008; Meier et al. 2013; Blanchet 2011; Backes et al. 2008], among which several types and effects systems [Gordon and Jeffrey 2003; 2004; Bugliesi et al. 2004b; Maffei 2004; Bugliesi et al. 2004a; 2005; 2007; Backes et al. 2007b; Backes et al. 2009; Focardi and Maffei 2011]. These type systems incorporate ad-hoc mechanisms to deal with nonce handshakes and, thus, to enforce injective agreement properties. Our exponential serialization technique can be seen as a logic-based generalization of such mechanisms, independent of the language and the type system. As a consequence, our type system is similarly able to verify authenticity in terms of injective agreement, while allowing for expressing also a number of more sophisticated properties involving access counts and usage bounds. As a downside, the current formulation of our type system does not allow to validate some specific nonce-handshake idioms, like the SOSH scheme [Gordon and Jeffrey 2004]. Still, this can be recovered by extending our type system with union and intersection types, as shown in [Backes et al. 2011; 2014].

In previous work [Bugliesi et al. 2011; 2012], we made initial steps towards the design of a sound system for resource-sensitive authorization, drawing on techniques from type systems for authentication and an affine extension of existing refinement type systems for the applied pi-calculus [Abadi and Fournet 2001]. That work aims at analyzing cryptographic protocols as opposed to their implementations. Furthermore, such a type system is designed around a specific cryptographic library: the consequence is that extending the analysis to new primitives requires significant changes in the soundness proof of the type system. In contrast, the usage of a λ-calculus in this work allows us to encode cryptography in the language using a standard sealing mechanism (cf. Section 7.2), which makes the analysis technique easily extensible to new cryptographic primitives. Finally, the non-standard nature of our previous type system makes it difficult to devise an efficient algorithmic variant, which in turn can be cleanly designed for the present proposal.

12. CONCLUSIONS

We presented the first type system for statically enforcing the (robust) safety of cryptographic protocol implementations with respect to authorization policies expressed in affine logic. Our type system benefits from the novel concept of exponential serialization to achieve a general and flexible treatment of affine formulas in distributed systems: we showed the effectiveness of this technique on two existing cryptographic protocols. We finally proposed an efficient, sound, and complete algorithmic variant of the type system, which is the key for a practical implementation of our analysis technique.

We are currently working on the mechanization of our theory by implementing a type-checker based on the algorithmic typing rules. We plan to facilitate type-checking

and reduce the need for manual type annotations by taking advantage of recent research on type inference in intuitionistic linear logic [Baillot and Hofmann 2010].

REFERENCES

Martín Abadi and Cédric Fournet. 2001. Mobile values, new names, and secure communication. In *Proc. 28th Symposium on Principles of Programming Languages (POPL)*. ACM Press, 104–115.

Alessandro Armando, David A. Basin, Yohan Boichut, Yannick Chevalier, Luca Compagna, Jorge Cuéllar, Paul Hankes Drielsma, Pierre-Cyrille Héam, Olga Kouchnarenko, Jacopo Mantovani, Sebastian Mödersheim, David von Oheimb, Michaël Rusinowitch, Judson Santiago, Mathieu Turuani, Luca Viganò, and Laurent Vigneron. 2005. The AVISPA tool for the automated validation of internet security protocols and applications. In *Proc. Computer Aided Verification'05 (CAV) (Lecture Notes in Computer Science)*. 281–285.

Michael Backes, Agostino Cortesi, Riccardo Focardi, and Matteo Maffei. 2007b. A Calculus of Challenges and Responses. In *Proc. 5th ACM Workshop on Formal Methods in Security Engineering (FMSE)*. ACM Press, 101–116.

Michael Backes, Agostino Cortesi, and Matteo Maffei. 2007a. Causality-based Abstraction of Multiplicity in Cryptographic Protocols. In *Proc. 20th IEEE Symposium on Computer Security Foundations (CSF)*. IEEE Computer Society Press, 355–369.

Michael Backes, Martin P. Grochulla, Catalin Hritcu, and Matteo Maffei. 2009. Achieving Security Despite Compromise Using Zero-knowledge. In *Proc. 22nd IEEE Symposium on Computer Security Foundations (CSF)*. IEEE Computer Society Press, 308–323.

Michael Backes, Cătălin Hriţcu, and Matteo Maffei. 2008. Type-checking Zero-knowledge. In *15th ACM Conference on Computer and Communications Security (CCS)*. ACM Press, 357–370.

Michael Backes, Catalin Hriţcu, and Matteo Maffei. 2011. Union and Intersection Types for Secure Protocol Implementations. In *Proc. Theory of Security and Applications (TOSCA) (Lecture Notes in Computer Science)*. Springer-Verlag, 1–28.

Michael Backes, Catalin Hriţcu, and Matteo Maffei. 2014. Union and Intersection Types for Secure Protocol Implementations. *Journal of Computer Security* 22, 2 (2014), 301–353.

Michael Backes, Matteo Maffei, and Dominique Unruh. 2010. Computationally Sound Verification of Source Code. In *Proc. 17th ACM Conference on Computer and Communications Security (CCS)*. ACM Press, 387–398.

Patrick Baillot and Martin Hofmann. 2010. Type Inference in Intuitionistic Linear Logic. In *Proc. 12th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP)*. ACM Press, 219–230.

Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. 2011. Refinement types for secure implementations. *ACM Transactions on Programming Languages and Systems* 33, 2, Article 8 (2011), 45 pages.

Karthikeyan Bhargavan, Ricardo Corin, Pierre-Malo Deniélou, Cédric Fournet, and James J. Leifer. 2009. Cryptographic Protocol Synthesis and Verification for Multiparty Sessions. In *Proc. 22nd IEEE Symposium on Computer Security Foundations (CSF)*. IEEE Computer Society Press, 124–140.

Karthikeyan Bhargavan, Cédric Fournet, and Andrew D. Gordon. 2010. Modular Verification of Security Protocol Code by Typing. In *Proc. 37th Symposium on Principles of Programming Languages (POPL)*. ACM Press, 445–456.

Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Stephen Tse. 2008. Verified Interoperable Implementations of Security Protocols. *ACM Transactions on Programming Languages and Systems* 31, 1, Article 5 (2008), 61 pages.

Kevin Bierhoff and Jonathan Aldrich. 2007. Modular typestate checking of aliased objects. In *Proc. 22nd Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications (OOPSLA)*. ACM Press, 301–320.

Bruno Blanchet. 2001. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules.. In *Proc. 14th IEEE Computer Security Foundations Workshop (CSFW)*. IEEE Computer Society Press, 82–96.

Bruno Blanchet. 2011. Using Horn Clauses for Analyzing Security Protocols. In *Formal Models and Techniques for Analyzing Security Protocols*. Cryptology and Information Security, Vol. 5. IOS Press, Chapter 7, 86–111.

Michele Bugliesi, Stefano Calzavara, Fabienne Eigner, and Matteo Maffei. 2011. Resource-Aware Authorization Policies for Statically Typed Cryptographic Protocols. In *Proc. 24th IEEE Symposium on Computer Security Foundations (CSF)*. IEEE Computer Society Press, 83–98.

Michele Bugliesi, Stefano Calzavara, Fabienne Eigner, and Matteo Maffei. 2012. Affine Refinement Types for Authentication and Authorization. In *Proc. 7th International Symposium on Trustworthy Global Computing (TGC)*. Springer-Verlag, 19–33.

Michele Bugliesi, Stefano Calzavara, Fabienne Eigner, and Matteo Maffei. 2013. Logical Foundations of Secure Resource Management in Protocol Implementations. In *Proc. 2nd Conference on Principles of Security and Trust (POST)*. Springer-Verlag, 105–125.

Michele Bugliesi, Riccardo Focardi, and Matteo Maffei. 2004a. Authenticity by Tagging and Typing. In *Proc. 2nd ACM Workshop on Formal Methods in Security Engineering (FMSE)*. ACM Press, 1–12.

Michele Bugliesi, Riccardo Focardi, and Matteo Maffei. 2004b. Compositional Analysis of Authentication Protocols. In *Proc. 13th European Symposium on Programming (ESOP) (Lecture Notes in Computer Science)*, Vol. 2986. Springer-Verlag, 140–154.

Michele Bugliesi, Riccardo Focardi, and Matteo Maffei. 2005. Analysis of Typed Analyses of Authentication Protocols. In *Proc. 18th IEEE Computer Security Foundations Workshop (CSFW)*. IEEE Computer Society Press, 112–125.

Michele Bugliesi, Riccardo Focardi, and Matteo Maffei. 2007. Dynamic Types for Authentication. *Journal of Computer Security* 15, 6 (2007), 563–617.

Peter C. Chapin, Christian Skalka, and Xiaoyang Sean Wang. 2008. Authorization in Trust Management: Features and Foundations. *Comput. Surveys* 40, 3, Article 9 (2008), 48 pages.

Cas J. F. Cremers. 2008. The Scyther Tool: Verification, Falsification, and Analysis of Security Protocols. In *Proc. Computer Aided Verification'08 (CAV) (Lecture Notes in Computer Science)*. Springer-Verlag, 414–418.

Nicholaas G. de Bruijn. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* 75, 5 (1972), 381 – 392.

Manuel Fähndrich and Robert DeLine. 2002. Adoption and focus: practical linear types for imperative programming. In *Proc. ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM Press, 13–24.

Riccardo Focardi and Matteo Maffei. 2011. Types for Security Protocols. In *Formal Models and Techniques for Analyzing Security Protocols*. Cryptology and Information Security, Vol. 5. IOS Press, Chapter 7, 143–181.

Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. 2005. A Type Discipline for Authorization Policies. In *Proc. 14th European Symposium on Programming (ESOP) (Lecture Notes in Computer Science)*. Springer-Verlag, 141–156.

Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. 2007. A Type Discipline for Authorization in Distributed Systems. In *Proc. 20th IEEE Symposium on Computer Security Foundations (CSF)*. IEEE Computer Society Press, 31–45.

Cédric Fournet, Markulf Kohlweiss, and Pierre-Yves Strub. 2011. Modular Code-Based Cryptographic Verification. In *Proc. 18th ACM Conference on Computer and Communications Security (CCS)*. ACM Press, 341–350.

Jean-Yves Girard. 1995. Linear Logic: Its Syntax and Semantics. In *Advances in Linear Logic (London Mathematical Society Lecture Note Series)*, Vol. 22. Cambridge University Press, 1–42.

Andrew D. Gordon and Alan Jeffrey. 2003. Authenticity by Typing for Security Protocols. *Journal of Computer Security* 11, 4 (2003), 451–519.

Andrew D. Gordon and Alan Jeffrey. 2004. Types and Effects for Asymmetric Cryptographic Protocols. *Journal of Computer Security* 12, 3 (2004), 435–484.

Joshua D. Guttman, F. Javier Thayer, Jay A. Carlson, Jonathan C. Herzog, John D. Ramsdell, and Brian T. Sniffen. 2004. Trust Management in Strand Spaces: A Rely-Guarantee Method. In *Proc. 13th European Symposium on Programming (ESOP) (Lecture Notes in Computer Science)*. Springer-Verlag, 325–339.

Matteo Maffei. 2004. Tags for Multi-Protocol Authentication. In *Proc. 2nd International Workshop on Security Issues in Coordination Models, Languages, and Systems (SECCO '04) (Electronic Notes on Theoretical Computer Science)*. Elsevier Science Publishers Ltd., 55–63.

Yitzhak Mandelbaum, David Walker, and Robert Harper. 2003. An effective theory of type refinements. In *Proc. 8th ACM SIGPLAN International Conference on Functional Programming (ICFP)*. ACM Press, 213–225.

Simon Meier, Benedikt Schmidt, Cas Cremers, and David A. Basin. 2013. The TAMARIN Prover for the Symbolic Analysis of Security Protocols. In *Proc. Computer Aided Verification'13 (CAV) (Lecture Notes in Computer Science)*. Springer-Verlag, 696–701.

Robin Milner. 1992. Functions as Processes. *Mathematical Structures in Computer Science* 2, 2 (1992), 119–141.

James H. Morris. 1973. Protection in Programming Languages. *Commun. ACM* 16, 1 (1973), 15–21.

Karl Naden, Robert Bocchino, Jonathan Aldrich, and Kevin Bierhoff. 2012. A Type System for Borrowing Permissions. In *Proc. 39th Symposium on Principles of Programming Languages (POPL)*. ACM Press, 557–570.

Jennifer G. Steiner, Clifford Neuman, and Jeffrey I. Schiller. 1988. Kerberos: An Authentication Service for Open Network Systems. In *Proc. USENIX Summer Conference*. USENIX Association, 191–202.

Eijiro Sumii and Benjamin C. Pierce. 2007. A Bisimulation for Dynamic Sealing. *Theoretical Computer Science* 375, 1-3 (2007), 169–192.

Joshua Sunshine, Karl Naden, Sven Stork, Jonathan Aldrich, and Éric Tanter. 2011. First-Class State Change in Plaid. In *Proc. 26th Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications (OOPSLA)*. ACM Press, 713–732.

Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. 2011. Secure Distributed Programming with Value-Dependent Types. In *Proc. 16th ACM SIGPLAN International Conference on Functional Programming (ICFP)*. ACM Press, 266–278.

Alwen Tiu and Alberto Momigliano. 2012. Cut elimination for a logic with induction and co-induction. *J. Applied Logic* 10, 4 (2012), 330–367.

Naoyuki Tomura. 1995. llprover - A Linear Logic Prover. (1995). http://bach.istc.kobe-u.ac.jp/llprover/.

Jesse Tov and Riccardo Pucella. 2010. Stateful Contracts for Affine Types. In *Proc. 19th European Symposium on Programming (ESOP)*. Springer-Verlag, 550–569.

Anne S. Troelstra. 1992. Lectures on Linear Logic. CSLI Stanford, LNS, vol. 29. (1992).

APPENDIX

A. SOUNDNESS OF EXPONENTIAL SERIALIZATION

We detail a soundness proof for the exponential serialization technique. The main idea behind the proof is to extend the notion of *rank* from formulas to multiset of assumptions: the rank of the multiset is computed by taking into account the serializers occurring therein. On the basis of the rank, we can always identify a serializer whose presence in the multiset can be shown to not affect the logical entailment of payload formulas, thus it can be safely removed from the assumptions. Iterating this process for each serializer, we establish that none of them affects derivability. In the proof we heavily rely on the stratification hypothesis.

In Appendix A.1 we introduce definitions and notations, which are needed for the technical development. In Appendix A.2 we present several auxiliary lemmas, which are used in the proof of the main results. In Appendix A.3 we carry out the proof of Theorem 4.4, establishing the soundness of exponential serialization. Finally, in Appendix A.4 we prove Proposition 4.5, which states a syntactic criterion for checking if a multiset of formulas is controlled.

A.1. Preliminaries

We first introduce some notational conventions. We let:

$$\hat{S} \in \{!\forall \widetilde{x}.(P \multimap !(C \multimap P)), \forall \widetilde{x}.(P \multimap !(C \multimap P))\}$$

for some (possibly empty) $\widetilde{x}$ and some payload formula P and some control formula C. We also write Δ, F^n as a short for the multiset $\Delta, F, \ldots, F$ (with n occurrences of F).

We say that a multiset of formulas is *well-formed* when it is stratified and it satisfies further simple syntactic conditions, consistent with the productions given in Section 4.

Definition A.1 (*Well-formation*). A multiset of formulas $\Delta = \Delta_1, \Delta_2, \Delta_3$ is *well-formed* if and only if it is stratified and $\Delta_1 = P_1, \ldots, P_l$, $\Delta_2 = G_1, \ldots, G_m$, $\Delta_3 = \hat{S}_1, \ldots, \hat{S}_n$.

We define a partial function $guard$ from formulas to control formulas, defined in the following cases:

— $guard(C \multimap P) = C$;
— $guard(P \multimap G) = guard(G)$;
— $guard(\forall x.F) = guard(F)$ whenever $guard(F)$ is defined;
— $guard(!F) = guard(F)$ whenever $guard(F)$ is defined.

We extend the notion of *rank* to a multiset of formulas Δ as follows:

$$rk(\Delta) = min\ \{rk(C) \mid \exists F \in \Delta : guard(F) = C\}$$

If the previous set is empty, we stipulate $rk(\Delta) = +\infty$.

A control formula C is *active* in Δ if and only if $rk(C) \leq rk(\Delta)$; we simply say that C is active whenever Δ is clear from the context. The previous notion is useful to relax the definition of controlled multiset to a weaker variant.

Definition A.2 (*Weak Control*). A well-formed multiset Δ is *weakly controlled* if and only if, for every active control formula C, we have that $\Delta \vdash C^k$ implies $k \leq 1$.

We note as expected that any controlled multiset is also weakly controlled.

PROPOSITION A.3. *If Δ is controlled, then it is weakly controlled.*

In the next results we focus without loss of generality on cut-free proofs.

A.2. Auxiliary results

The first two lemmas are needed to show that the induction hypothesis can indeed be applied in the proof of a number of subsequent results.

LEMMA A.4. *Let Δ be well-formed. The following implications hold:*

(1) *if $\Delta = \Delta', F$, then Δ' is well-formed and $rk(\Delta) \leq rk(\Delta')$;*
(2) *if $\Delta = \Delta', !F$, then $\Delta, !F$ is well-formed and $rk(\Delta) = rk(\Delta, !F)$;*
(3) *if $\Delta = \Delta', F_1 \otimes F_2$, then Δ', F_1, F_2 is well-formed and $rk(\Delta) = rk(\Delta', F_1, F_2)$;*
(4) *if $\Delta = \Delta_1, \Delta_2, F_1 \multimap F_2$, then Δ_2, F_2 is well-formed and $rk(\Delta) \leq rk(\Delta_2, F_2)$;*
(5) *if $\Delta = \Delta', \forall x.F$, then for every t we have that $\Delta', F\{t/x\}$ is well-formed and $rk(\Delta) = rk(\Delta', F\{t/x\})$;*
(6) *if $\Delta = \Delta', !F$, then Δ', F is well-formed and $rk(\Delta) = rk(\Delta', F)$.*

LEMMA A.5. *Let Δ be weakly controlled. The following implications hold:*

(1) *if $\Delta = \Delta', F$, then Δ' is weakly controlled;*
(2) *if $\Delta = \Delta', !F$, then $\Delta, !F$ is weakly controlled;*
(3) *if $\Delta = \Delta', F_1 \otimes F_2$, then Δ', F_1, F_2 is weakly controlled;*
(4) *if $\Delta = \Delta_1, \Delta_2, F_1 \multimap F_2$ and $\Delta_1 \vdash F_1$, then Δ_2, F_2 is weakly controlled;*
(5) *if $\Delta = \Delta', \forall x.F$, then $\Delta', F\{t/x\}$ is weakly controlled for every t;*
(6) *if $\Delta = \Delta', !F$, then Δ', F is weakly controlled.*

The next lemma formalizes the intuition behind stratification: formulas with a given rank are never needed in the proof of a control formula with a lower rank. This observation plays a prominent role in many of the later results.

LEMMA A.6 (STRATIFICATION). *Let $\Delta = \Delta', F_1, \ldots, F_m$ be well-formed and let C be active in Δ. If $\Delta \vdash C^n$ with $n \geq 1$ and $\forall i \in [1, m] : rk(F_i) > rk(C)$, then $\Delta' \vdash C^n$.*

The next result is a strengthening lemma: if a multiset does not entail a given control formula C, then any implication of the form $C \multimap P$ occurring therein can be removed without affecting derivability.

LEMMA A.7 (STRENGTHENING). *Let $\Delta = \Delta', C \multimap P$ be well-formed and let C be active in Δ. If $\Delta \vdash P'$ and $\Delta' \nvdash C$, then $\Delta' \vdash P'$.*

The next technical lemma allows to apply the induction hypothesis in the proof of the subsequent results.

LEMMA A.8. *Let Δ be weakly controlled, then Δ, B is weakly controlled. Moreover, we have $rk(\Delta) = rk(\Delta, B)$.*

The next lemma formalizes an important intuition: since any active control formula C can be proved at most once in a weakly controlled multiset, all the implications of the form $C \multimap P$ occurring therein can be replaced by a single implication of the same form without affecting derivability. This is needed in the proof of Lemma A.11 (Dereliction).

LEMMA A.9 (BOUNDED REACTION). *Let $\Delta = \Delta', (C \multimap P)^n$ be weakly controlled. If $\Delta \vdash P'$ and C is active in Δ, then $\Delta', C \multimap P \vdash P'$.*

The next technical corollary is used in the proof of Lemma A.11 (Dereliction) below.

COROLLARY A.10. *Let Δ be well-formed. If $\Delta \vdash C$ and C is active in Δ, then Δ contains at least an affine formula.*

The next lemma is reminiscent of the idea behind Lemma A.9 (Bounded Reaction): since any active control formula C can be proved at most once in a weakly controlled

multiset, an exponential implication of the form $!(C \multimap P)$ occurring therein can be replaced by an affine implication $C \multimap P$ without affecting derivability.

LEMMA A.11 (DERELICTION). *Let $\Delta = \Delta', !(C \multimap P)$ be weakly controlled. If $\Delta \vdash P'$ and C is active in Δ, then $\Delta', C \multimap P \vdash P'$.*

The next lemma states that an implication of the form $C \multimap P$ is useless whenever the payload formula P is already available in the context. This is needed to prove Corollary A.13 (Bounded Usage), which is the real result of interest and formalizes a similar idea.

LEMMA A.12. *Let $\Delta = \Delta', C \multimap P$ be well-formed. If $\Delta \vdash P'$, then $\Delta', P \vdash P'$.*

COROLLARY A.13 (BOUNDED USAGE). *Let $\Delta = \Delta', !(C \multimap P)$ be weakly controlled. If $\Delta \vdash P'$ and C is active in Δ, then $\Delta', P \vdash P'$.*

A.3. Proof of Theorem 4.4

We are finally ready to show that serializers do not affect the derivability of payload formulas. However, since a serializer $S = !\forall \widetilde{x}.(P \multimap !(C \multimap P))$ has a non-trivial structure, in the proof of the main theorem we must take into account the possibility of decomposing S through the left rules of the logic, by (i) removing the exponential modality, and (ii) instantiating the quantifiers. Lemma A.14 below accounts for such scenario: its proof relies on Corollary A.13 (Bounded Usage), the central result of the previous section.

LEMMA A.14. *Let $\Delta = \Delta', P \multimap !(C \multimap P)$ be weakly controlled and let C be active in Δ. If $\Delta \vdash P'$, then $\Delta' \vdash P'$.*

Lemma A.15 strongly resembles Lemma A.14 and serves a similar purpose: a formula of the form $\forall \widetilde{x}.(P \multimap !(C \multimap P))$ is obtained from a serializer by removing the exponential modality from it. Since formulas of this form can arise in the proof of the main result, we must first deal with them.

LEMMA A.15. *Let $\Delta = \Delta', \forall \widetilde{x}.(P \multimap !(C \multimap P))$ be weakly controlled and let C be active in Δ. If $\Delta \vdash P'$, then $\Delta' \vdash P'$.*

The next lemma is the key to proving our main theorem: it states that serializers whose guard is active can be safely removed from a weakly controlled multiset without affecting the derivability of payload formulas.

LEMMA A.16 (WEAK SOUNDNESS). *Let $\Delta = \Delta', !\forall \widetilde{x}.(P \multimap !(C \multimap P))$ be weakly controlled and let C be active in Δ. If $\Delta \vdash P'$, then $\Delta' \vdash P'$.*

The next proposition is needed to identify a candidate serializer to remove in the proof of the main theorem, according to the explained proof strategy.

PROPOSITION A.17. *Let $\Delta = P_1, \dots, P_m, S_1, \dots, S_n$. If $n > 0$, then there exists $S_i \in \Delta$ such that $guard(S_i)$ is active in Δ.*

In the next lemma we make explicit our proof strategy. The lemma immediately entails our real result of interest, Theorem 4.4 below.

LEMMA A.18. *Let $\Delta' = P_1, \dots, P_m$. If $\Delta = \Delta', S_1, \dots, S_n$ is weakly controlled and $\Delta \vdash P'$, then $\Delta' \vdash P'$.*

RESTATEMENT 1 (OF THEOREM 4.4). *Let $\Delta' = P_1, \dots, P_m$. If $\Delta = \Delta', S_1, \dots, S_n$ is controlled and $\Delta \vdash P'$, then $\Delta' \vdash P'$.*

PROOF. Immediate by Lemma A.18, since any controlled multiset is also weakly controlled by Proposition A.3. □

A.4. Proof of Proposition 4.5

The next proposition is a simple observation on the derivability of control formulas. It is needed in the proof of Proposition 4.5 below.

PROPOSITION A.19. *Let* $\Delta = B_1, \ldots, B_l, C_1, \ldots, C_m$. *If* $\Delta \vdash C^k$*, then* C *occurs at least* k *times in* Δ.

Finally, we are ready to prove the correctness of our syntactic criterion for checking if a multiset is controlled.

RESTATEMENT 2 (OF PROPOSITION 4.5). *If* $\Delta = B_1, \ldots, B_l, C_1, \ldots, C_m, S_1, \ldots, S_n$ *is stratified and the control formulas in* Δ *are pairwise distinct, then* Δ *is controlled.*

PROOF. We first show that Δ is weakly controlled. Let $\Delta_1 = B_1, \ldots, B_l$, $\Delta_2 = C_1, \ldots, C_m$ and $\Delta_3 = S_1, \ldots, S_n$. Let us assume by contradiction that $\Delta \vdash C^h$ with $h \geq 2$ for some active control formula C. By Lemma A.6 (Stratification) we have $\Delta_2 \vdash C^h$, since any formula in Δ_1 and Δ_3 has an infinite rank, while the rank of C is finite. By Proposition A.19, C must occur at least h times in Δ_2, but this is contradictory with respect to the initial hypotheses.

Now we know that Δ is weakly controlled and we can show that it is, in fact, controlled. Let us assume by contradiction that $\Delta \vdash C^h$ with $h \geq 2$ for some arbitrary control formula C, then by Lemma A.18 we have $\Delta_1, \Delta_2 \vdash C^h$. By Proposition A.19, C must occur at least h times in Δ_1, Δ_2, but this is contradictory with respect to the initial hypotheses. □

B. SOUNDNESS OF THE TYPE SYSTEM

We present a complete soundness proof for our type system. The structure of the proof is standard: we first establish a Subject Reduction theorem, which shows that types are preserved upon reduction, and then we prove that well-typed programs are *statically* safe. By combining these two guarantees, we establish that the type system enforces our safety notion. Finally, we prove an Opponent Typability lemma, which states that any opponent is trivially well-typed: this allows us to carry out a simple proof of robust safety, based on our safety theorem.

The present appendix is organized as follows:

— Appendix B.1 develops basic properties of affine logic, which are needed in the soundness proof of the type system;
— Appendix B.2 establishes some basic results about the type system and the environment rewriting relation;
— Appendix B.3 presents the main properties of kinding and subtyping, most notably the transitivity of the subtyping relation;
— Appendix B.4 establishes a standard substitution lemma;
— Appendix B.5 provides the inversion lemmas for the constructed values of our framework;
— Appendix B.6 presents the fundamental properties of the extraction relation;
— Appendix B.7 details the proof of the Subject Reduction theorem, building upon the results of the previous sections;
— Appendix B.8 presents the proof of robust safety.

B.1. Properties of the logic

We first show that affine logic is closed under substitution of variables with closed terms. This is important to prove the substitution lemma of our type system.

LEMMA B.1 (SUBSTITUTION FOR THE LOGIC). *For all Δ, F and all substitutions σ of variables with closed terms, it holds that $\Delta \vdash F$ implies $\Delta\sigma \vdash F\sigma$.*

In the next result we recall that we write $\Delta \vdash \Delta'$ to stand for $\Delta \vdash F_1 \otimes \ldots \otimes F_n$ whenever $\Delta' = F_1, \ldots, F_n$. If Δ' is empty, we let $\Delta \vdash \Delta'$ stand for $\Delta \vdash \mathbf{1}$.

LEMMA B.2 (PROPERTIES OF CONJUNCTION). *The following properties hold:*

(1) For all $n \geq 0$, we have $\Delta, F_1, \ldots, F_n \vdash F$ iff $\Delta, F_1 \otimes \ldots \otimes F_n \vdash F$.
(2) For all Δ, Δ' it holds that $\Delta' \subseteq \Delta$ implies $\Delta \vdash \Delta'$.

The next result is a generalization to multisets of formulas of the standard Cut rule, characteristic of sequent calculi presentation of formal logic.

LEMMA B.3 (MULTICUT). *If $\Delta \vdash \Delta'$ and $\Delta', \Delta'' \vdash \Delta'''$, then $\Delta, \Delta'' \vdash \Delta'''$.*

The next technical lemma formalizes the intuition that exponential formulas can be proved an arbitrary number of times.

LEMMA B.4 (PROPERTIES OF CONTRACTION). *The following properties hold:*

(1) For all Δ it holds that $!\Delta \vdash !\Delta, !\Delta$.
(2) For all Δ, Δ' it holds that if $\Delta \vdash !\Delta'$, then $\Delta \vdash !\Delta', !\Delta'$.

B.2. Basic results

The next results are completely standard. In the following we typically let $\mathcal{J}$ range over the judgements $\{\diamond, T, F, T :: k, T <: U, E : T\}$.

LEMMA B.5 (DERIVED JUDGEMENTS). *It holds that:*

(1) If $\Gamma; \Delta \vdash \diamond$, then $fnfv(\Delta) \subseteq dom(\Gamma)$ and $\forall \Delta' \subseteq \Delta : \Gamma; \Delta' \vdash \diamond$.
(2) If $\Gamma; \Delta \vdash \diamond$ and $(x : T) \in \Gamma$, then $T = \psi(T)$.
(3) If $\Gamma; \Delta \vdash T$, then $\Gamma; \emptyset \vdash \psi(T)$.
(4) If $\Gamma; \Delta \vdash T$, then $\Gamma; \Delta \vdash \diamond$ and $fnfv(T) \subseteq dom(\Gamma)$.
(5) If $\Gamma; \Delta \vdash F$, then $\Gamma; \Delta \vdash \diamond$ and $fnfv(F) \subseteq dom(\Gamma)$.
(6) If $\Gamma; \Delta \hookrightarrow \Gamma; \Delta'$, then $\Gamma; \Delta \vdash \diamond$ and $\Gamma; \Delta' \vdash \diamond$.
(7) If $\Gamma; \Delta \vdash T :: k$, then $\Gamma; \Delta \vdash T$.
(8) If $\Gamma; \Delta \vdash T <: T'$, then $\Gamma; \Delta \vdash T$ and $\Gamma; \Delta \vdash T'$.
(9) If $\Gamma; \Delta \vdash E : T$, then $\Gamma; \Delta \vdash T$ and $fnfv(E) \subseteq dom(\Gamma)$.

LEMMA B.6 (JOINING ENVS). *If $\Gamma; \Delta \vdash \diamond$ and $\Gamma; \Delta' \vdash \diamond$, then $\Gamma; \Delta, \Delta' \vdash \diamond$.*

Notation 1 (*Environment Entry η*). We define an environment entry η to be either a type environment entry μ or a formula F.

Notation 2 (*Environment Join $\bullet$*). We introduce the following notation for environment join:

$$(\Gamma; \Delta) \bullet \mu \triangleq \begin{cases} \Gamma, x : \psi(T); \Delta, forms(x : T) & \text{if } \mu = x : T \\ \Gamma, \mu; \Delta & \text{otherwise} \end{cases}$$

$$(\Gamma; \Delta) \bullet F \triangleq \Gamma; \Delta, F$$

$$(\Gamma; \Delta) \bullet (\Gamma'; \Delta') \triangleq \Gamma, \Gamma'; \Delta, \Delta'$$

LEMMA B.7 (WEAKENING). *If* $(\Gamma;\Delta) \bullet (\Gamma';\Delta') \vdash \mathcal{J}$ *and* $(\Gamma;\Delta) \bullet \eta \bullet (\Gamma';\Delta') \vdash \diamond$, *then* $(\Gamma;\Delta) \bullet \eta \bullet (\Gamma';\Delta') \vdash \mathcal{J}$.

The next lemma establishes some basic properties of the rewriting relation. Intuitively, we show that this relation is consistent with logical entailment, in that it satisfies some expected properties which hold true for the latter.

LEMMA B.8 (PROPERTIES OF REWRITING). *The following statements hold true:*

(1) If $\Gamma;\Delta \vdash \diamond$ *and* $\Delta' \subseteq \Delta$, *then* $\Gamma;\Delta \hookrightarrow \Gamma;\Delta'$.
(2) If $\Gamma;\Delta_1 \hookrightarrow \Gamma;\Delta'_1$ *and* $\Gamma;\Delta_2 \hookrightarrow \Gamma;\Delta'_2$, *then* $\Gamma;\Delta_1,\Delta_2 \hookrightarrow \Gamma;\Delta'_1,\Delta'_2$.
(3) If $\Gamma;\Delta \hookrightarrow \Gamma;\Delta'$ *and* $\Gamma;\Delta' \hookrightarrow \Gamma;\Delta''$, *then* $\Gamma;\Delta \hookrightarrow \Gamma;\Delta''$.
(4) If $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta'$, *then* $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta',!\Delta'$.

The next result states that, if an environment $\Gamma;\Delta$ can be rewritten to an environment $\Gamma;\Delta'$, then it can derive all the judgements provable by the latter.

LEMMA B.9 (REWRITE WEAK). *If* $\Gamma;\Delta' \vdash \mathcal{J}$ *and* $\Gamma;\Delta \hookrightarrow \Gamma;\Delta'$, *then* $\Gamma;\Delta \vdash \mathcal{J}$.

The next technical lemma states that rewriting does not introduce free variables.

LEMMA B.10 (REWRITING AND VARIABLES). *If* $x \notin dom(\Gamma)$ *and* $\Gamma;\Delta \hookrightarrow \Gamma;\Delta'$, *then* $x \notin fv(\Delta')$.

The next lemma is an expected property of the refinement stripping function ψ, i.e., that it removes all the refinement formulas from a type.

LEMMA B.11 (SOUNDNESS OF ψ). *For every type T, we have* $forms(x : \psi(T)) = \emptyset$.

The next lemma states that the stripping function ψ is idempotent, i.e., there is no purpose in stripping refinements twice from the same type.

LEMMA B.12 (IDEMPOTENT ψ). *For every type T, we have* $\psi(\psi(T)) = \psi(T)$.

B.3. Properties of kinding and subtyping

The next result states that, whenever a typing environment can assign a kind to a type T, then it can be rewritten so as to be split in two distinct components: the first one is exponential and it is needed to kind-check the structural information $\psi(T)$, while the second one can be used to derive the refinement formulas $forms(x : T)$ when T is tainted. This result is extensively used in the proofs, most likely to deal with the subtleties introduced by environment splitting.

LEMMA B.13 (BARE KINDS). *If* $\Gamma;\Delta \vdash T :: k$, *then there exist $!\Delta'$ and Δ'' such that* $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta',\Delta''$ *and* $\Gamma;!\Delta' \vdash \psi(T) :: k$. *Moreover, if $k = $ tnt, we can also require* $\Delta'' \vdash forms(x : T)$ *for any* $x \notin dom(\Gamma)$.

The next technical lemma is needed in the proof of Lemma B.19 below. It states that the assignment of a public kind does not depend on the refinement formulas associated to the type, but only on structural information.

LEMMA B.14 (BARE KINDS REVERSE). *If* $\Gamma;\Delta \vdash T$ *and* $\Gamma;\Delta \vdash \psi(T) :: $ pub, *then* $\Gamma;\Delta \vdash T :: $ pub.

The next result is similar in spirit to Lemma B.13, but it applies to subtyping. Again the goal is to identify a possible rewriting of the typing environment such that the structural subtyping relation and the refinement formulas can be proved separately. This is needed in a number of places to deal with the complications introduced by environment splitting.

LEMMA B.15 (BARE SUBTYPES). *If* $\Gamma; \Delta \vdash T <: U$, *then there exist* $!\Delta'$ *and* Δ'' *such that* $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta', \Delta''$ *and* $\Gamma; !\Delta' \vdash \psi(T) <: \psi(U)$ *and* $\Delta'', forms(x : T) \vdash forms(x : U)$ *for any* $x \notin dom(\Gamma)$.

The next technical lemma is needed in the proof of Lemma B.19 below. It states that only refinement formulas are relevant for many judgements of our type system, so we can always replace a purely structural type $\psi(T)$ with any other (well-formed) purely structural type $\psi(T')$ in the typing environment.

LEMMA B.16 (REPLACING UNREFINED BINDINGS). *For all* $\mathcal{J} \in \{\diamond, U, F, U :: k, U <: U'\}$ *it holds that if* $\Gamma, x : \psi(T), \Gamma'; \Delta \vdash \mathcal{J}$ *and* $\Gamma; \emptyset \vdash \psi(T')$, *then* $\Gamma, x : \psi(T'), \Gamma'; \Delta \vdash \mathcal{J}$. *Moreover, the depth of the two derivations is the same.*

Definition B.17 (*Compartmental Notation for Environments*).   Let $\Gamma[(\mu_i)^{i \in \{1, \ldots, n\}}]$ denote the environment obtained by inserting the entries $\mu_1, \ldots, \mu_n$ at fixed positions between the entries of the environment Γ.

The next technical lemma is needed in the proof of Lemma B.19 below. Intuitively, it states that kinding annotations for type variables do not play any role for many judgements of our type system.

LEMMA B.18 (TYPE VARIABLES AND KINDING). *For all* $\Gamma = \Gamma_0[(\alpha_i)^{i \in \{1, \ldots, n\}}]$ *and* $\hat{\Gamma} = \Gamma_0[(\alpha_i :: k_i)^{i \in \{1, \ldots, n\}}]$ *it holds that:*

(1) $dom(\Gamma) = dom(\hat{\Gamma})$;
(2) $\Gamma; \Delta \vdash \diamond$ *if and only if* $\hat{\Gamma}; \Delta \vdash \diamond$;
(3) $\Gamma; \Delta \hookrightarrow \Gamma; \Delta'$ *if and only if* $\hat{\Gamma}; \Delta \hookrightarrow \hat{\Gamma}; \Delta'$;
(4) $\Gamma; \Delta \vdash T$ *if and only if* $\hat{\Gamma}; \Delta \vdash T$;
(5) $\Gamma; \Delta \vdash F$ *if and only if* $\hat{\Gamma}; \Delta \vdash F$;
(6) If $\Gamma; \Delta \vdash T :: k$, *then* $\hat{\Gamma}; \Delta \vdash T :: k$.

The next lemma states that any subtype of a public type is public, while any supertype of a tainted type is tainted. This is needed to prove Lemma B.20 below.

LEMMA B.19 (PUBLIC DOWN/TAINTED UP). *For all environments* $\Gamma; \Delta$ *and types* T, T' *it holds that:*

(1) If $\Gamma; \Delta \vdash T <: T'$ *and* $\Gamma; \Delta' \vdash T' ::$ pub, *then* $\Gamma; \Delta, \Delta' \vdash T ::$ pub.
(2) If $\Gamma; \Delta \vdash T <: T'$ *and* $\Gamma; \Delta' \vdash T ::$ tnt, *then* $\Gamma; \Delta, \Delta' \vdash T' ::$ tnt.

The next result is central to proving the transitivity of our subtyping relation. It establishes a standard characterization of public and tainted kinds: a type is public iff it is a subtype of Un, while it is tainted iff it is a supertype of Un.

LEMMA B.20 (PUBLIC TAINTED). *For all environments* $\Gamma; \Delta$ *and types* T *we have:*

(1) $\Gamma; \Delta \vdash T ::$ pub *if and only if* $\Gamma; \Delta \vdash T <:$ Un.
(2) $\Gamma; \Delta \vdash T ::$ tnt *if and only if* $\Gamma; \Delta \vdash$ Un $<: T$.

The next technical lemma details a relationship between the stripping function ψ and the subtyping relation. It is invoked only once in the proof of transitivity for the subtyping relation.

LEMMA B.21 (SUBTYPING AND ψ). *The following statements hold true:*

(1) If $\Gamma; \emptyset \vdash T$, *then* $\Gamma; \emptyset \vdash T <: \psi(T)$.
(2) If $\Gamma; \Delta \vdash \psi(T) <: U$ *and* $\Gamma; \emptyset \vdash T$, *then* $\Gamma; \Delta \vdash T <: U$.

We are finally ready to prove the transitivity of the subtyping relation. This is a standard formulation for an affine setting.

LEMMA B.22 (TRANSITIVITY). *If $\Gamma;\Delta \vdash T <: T'$ and $\Gamma;\Delta' \vdash T' <: T''$, then $\Gamma;\Delta,\Delta' \vdash T <: T''$.*

B.4. Properties of substitution

The next result establishes for value typing judgements a property we already showed for kinding and subtyping judgements. Namely, if a typing environment $\Gamma;\Delta$ can assign a type T to a value M, then it can be rewritten into two distinct typing environments: an exponential environment $\Gamma;!\Delta'$ where M is assigned the structural type $\psi(T)$ and a possibly non-exponential environment $\Gamma;\Delta''$ where the refinement formulas $forms(x:T)$ can be proved on the value M.

LEMMA B.23 (BARE TYPES). *Let $fv(M) = \emptyset$. If $\Gamma;\Delta \vdash M : T$, then there exist $!\Delta'$ and Δ'' such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta',\Delta''$ and $\Gamma;!\Delta' \vdash M : \psi(T)$ and $\Delta'' \vdash forms(x:T)\{M/x\}$ for any $x \notin dom(\Gamma)$.*

The next lemma states that multiplicative conjunctions occurring in refinement types can be equivalently broken into their atomic components. This is needed in the proof of Lemma B.25 below.

LEMMA B.24 ($\otimes$ SUB). *If $\Gamma;\Delta \vdash \{x:T \mid F_1 \otimes F_2\}$, then $\Gamma;\emptyset \vdash \{x:T \mid F_1 \otimes F_2\} <:> \{x:\{x:T \mid F_1\} \mid F_2\}$.*

The next result is a very convenient lemma, which is needed in the proof of our substitution lemma. It essentially states that, if a typing environment $\Gamma;\Delta$ can assign a type T to a value M, then it can be rewritten into two distinct typing environments: an exponential environment $\Gamma;!\Delta'$ where M is assigned the structural type $\psi(T)$ and a possibly non-exponential environment $\Gamma;\Delta''$ where M is assigned the original type T. Hence, purely structural typing judgements can be proved arbitrarily many times. This is again needed to deal with the subtleties introduced by environment splitting.

LEMMA B.25 (AFFINE TYPING). *If $\Gamma;\Delta \vdash M : T$, then there exist $!\Delta'$ and Δ'' such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta',\Delta''$ and $\Gamma;!\Delta' \vdash M : \psi(T)$ and $\Gamma;\Delta'' \vdash M : T$.*

The next simple lemma states that, if a value M is assigned a refinement type T, then the refinement formulas $forms(x:T)$ can be proved on M from the formulas in the typing environment.

LEMMA B.26 (FORMULAS). *If $\Gamma;\Delta \vdash M : T$ and $x \notin dom(\Gamma)$, then $\Delta \vdash forms(x:T)\{M/x\}$.*

The next lemma establishes some basic syntactic properties of substitution.

LEMMA B.27 (BASIC SUBSTITUTION). *The following statements hold true:*

(1) For every type T, we have $\psi(T)\{M/x\} = \psi(T\{M/x\})$.
(2) If $x \neq y$, then $forms(y:T)\{M/x\} = forms(y:T\{M/x\})$.

Finally, we can state and prove our substitution lemma, showing that typing is preserved by substitution of closed values for variables with the same type. The statement is complicated by the necessity to join different environments, but the formulation is consistent with standard presentations of substructural type systems.

LEMMA B.28 (SUBSTITUTION). *Suppose that $\Gamma;\Delta \vdash M : U$ and $fv(M) = \emptyset$. The following statements hold true:*

(1) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \vdash \diamond$, *then* $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \vdash \diamond$.

(2) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \vdash F$, *then* $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \vdash F\{M/x\}$.

(3) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \hookrightarrow \Gamma, x : \psi(U), \Gamma'; \Delta^*$, *then*
   $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \hookrightarrow \Gamma, \Gamma'\{M/x\}; \Delta^*\{M/x\}$.

(4) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \vdash T$, *then* $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \vdash T\{M/x\}$.

(5) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \vdash T :: k$, *then* $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \vdash T\{M/x\} :: k$.

(6) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \vdash T <: T'$, *then* $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \vdash T\{M/x\} <: T'\{M/x\}$.

(7) *If* $(\Gamma; \Delta') \bullet x : U \bullet (\Gamma'; \Delta'') \vdash E : T$, *then* $\Gamma, \Gamma'\{M/x\}; \Delta, (\Delta', \Delta'')\{M/x\} \vdash E\{M/x\} : T\{M/x\}$.

B.5. Inversion lemmas

The next result is a standard bound weakening lemma: any occurrence of a type T in the typing environment can be safely replaced with a subtype T'.

LEMMA B.29 (BOUND WEAK). *Let* $\Gamma; \Delta \vdash T' <: T$. *If* $\Gamma, x : \psi(T), \Gamma'; \Delta', forms(x : T) \vdash \mathcal{J}$, *then* $\Gamma, x : \psi(T'), \Gamma'; \Delta, \Delta', forms(x : T') \vdash \mathcal{J}$.

We now present two technical lemmas which are needed to establish the inversion result for iso-recursive type constructors.

LEMMA B.30 (TYPE VARIABLES AND KINDING). *If* $\Gamma, \alpha, \Gamma'; \Delta \vdash T :: k$, *then* $\alpha \notin fnfv(T)$.

LEMMA B.31 (TYPE SUBSTITUTION). *For all* T, T' *such that* $T = \psi(T)$ *and* $T' = \psi(T')$ *it holds that:*

(1) *If* $\Gamma, \alpha, \Gamma'; \Delta \vdash \mathcal{J}$ *and* $\Gamma; \Delta' \vdash T$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); \Delta, \Delta' \vdash \mathcal{J}\{T/\alpha\}$.

(2) *If* $\Gamma, \alpha :: k, \Gamma'; \Delta \vdash \diamond$ *and* $\Gamma; \Delta' \vdash T :: k$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); \Delta, \Delta' \vdash \diamond$.

(3) *If* $\Gamma, \alpha :: k, \Gamma'; \Delta \vdash U$ *and* $\Gamma; \Delta' \vdash T :: k$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); \Delta, \Delta' \vdash U\{T/\alpha\}$.

(4) *If* $\Gamma, \alpha :: k, \Gamma'; \Delta \vdash U :: k'$ *and* $\Gamma; \Delta' \vdash T :: k$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); \Delta, \Delta' \vdash U\{T/\alpha\} :: k'$.

(5) *We have:*
   —*If* $\Gamma, \alpha, \Gamma'; \emptyset \vdash U$ *and* α *only occurs positively in* U *and* $\Gamma; !\Delta \vdash T <: T'$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); !\Delta \vdash U\{T/\alpha\} <: U\{T'/\alpha\}$.
   —*If* $\Gamma, \alpha, \Gamma'; \emptyset \vdash U$ *and* α *only occurs negatively in* U *and* $\Gamma; !\Delta \vdash T <: T'$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); !\Delta \vdash U\{T'/\alpha\} <: U\{T/\alpha\}$.

(6) *If* $\Gamma, \alpha, \Gamma'; \Delta \vdash U <: U'$ *and* α *only occurs positively in* U, U' *and* $\Gamma; \Delta' \vdash T <: T'$, *then* $\Gamma, (\Gamma'\{T/\alpha\}); \Delta, \Delta' \vdash U\{T/\alpha\} <: U'\{T'/\alpha\}$.

We can finally state and prove a number of inversion results for the constructed values of our framework. The goal is showing that the elementary components of these constructed values have indeed the expected types. There is a substantial amount of work to do, but the technical details are mostly standard.

LEMMA B.32 (INVERSION FOR FUNCTIONS). *The following statements hold:*

(1) *If* $\Gamma; \Delta \vdash \lambda x. E : V$, *then there exist* Δ_1, Δ_2, T, U *such that* $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_1, \Delta_2$ *and* $\Gamma; \Delta_1 \vdash \lambda x. E : x : T \to U$ *(by a top-level application of* VAL FUN*) and* $\Gamma; \Delta_2 \vdash x : T \to U <: \psi(V)$.

(2) *If* $\Gamma; \Delta \vdash x : T \to U <: x : T' \to U'$, *then there exist* Δ_1, Δ_2 *such that* $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, !\Delta_2$ *and* $\Gamma; !\Delta_1 \vdash T' <: T$ *and* $\Gamma, x : \psi(T'); !\Delta_2 \vdash U <: U'$.

(3) *If* $\Gamma; \Delta \vdash \lambda x. E : x : T \to U$, *then there exists a* Δ' *such that* $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta'$ *and* $(\Gamma; !\Delta') \bullet x : T \vdash E : U$.

(4) *If* $\Gamma; \Delta \vdash \lambda x. E : x : T \to U$, *then* $(\Gamma; \Delta) \bullet x : T \vdash E : U$.

LEMMA B.33 (INVERSION FOR PAIRS). *The following statements hold:*

(1) *If $\Gamma;\Delta \vdash (M,N) : V$, then there exist Δ_1, Δ_2, T, U such that $\Gamma;\Delta \hookrightarrow \Gamma;\Delta_1, \Delta_2$ and $\Gamma;\Delta_1 \vdash (M,N) : x : T * U$ (by a top-level application of* VAL PAIR*) and $\Gamma;\Delta_2 \vdash x : T * U <: \psi(V)$.*
(2) *If $\Gamma;\Delta \vdash x : T * U <: x : T' * U'$, then there exist Δ_1, Δ_2 such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1, !\Delta_2$ and $\Gamma;!\Delta_1 \vdash T <: T'$ and $\Gamma, x : \psi(T);!\Delta_2 \vdash U <: U'$.*
(3) *If $\Gamma;\Delta \vdash (M,N) : x : T * U$, then there exist Δ_1, Δ_2 such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1, !\Delta_2$ and $\Gamma;!\Delta_1 \vdash M : T$ and $\Gamma;!\Delta_2 \vdash N : U\{M/x\}$.*

LEMMA B.34 (INVERSION FOR SUM CONSTRUCTORS). *The following statements hold:*

(1) *Let $h \in \{\mathsf{inl}, \mathsf{inr}\}$. If $\Gamma;\Delta \vdash h\ M : V$, then there exist Δ_1, Δ_2, T, U such that $\Gamma;\Delta \hookrightarrow \Gamma;\Delta_1, \Delta_2$ and $\Gamma;\Delta_1 \vdash h\ M : T + U$ (by a top-level application of* VAL H*) and $\Gamma;\Delta_2 \vdash T + U <: \psi(V)$.*
(2) *If $\Gamma;\Delta \vdash T + U <: T' + U'$, then there exist Δ_1, Δ_2 such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1, !\Delta_2$ and $\Gamma;!\Delta_1 \vdash T <: T'$ and*
(3) *If $\Gamma;\Delta \vdash \mathsf{inl}\ M : T + U$, then there exist $!\Delta$ such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta$ and $\Gamma;!\Delta \vdash M : T$ and $\Gamma;!\Delta \vdash U$.*
(4) *If $\Gamma;\Delta \vdash \mathsf{inr}\ M : T + U$, then there exist Δ' such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta'$ and $\Gamma;!\Delta' \vdash M : U$ and $\Gamma;!\Delta' \vdash T$.*
(5) *If $\Gamma;\Delta \vdash \mathsf{inl}\ M : T + U$, then $\Gamma;\Delta \vdash M : T$.*
(6) *If $\Gamma;\Delta \vdash \mathsf{inr}\ M : T + U$, then $\Gamma;\Delta \vdash M : U$.*

LEMMA B.35 (INVERSION FOR RECURSIVE CONSTRUCTORS). *The following statements hold:*

(1) *If $\Gamma;\Delta \vdash \mathsf{fold}\ M : V$, then there exist Δ_1, Δ_2, T such that $\Gamma;\Delta \hookrightarrow \Gamma;\Delta_1, \Delta_2$ and $\Gamma;\Delta_1 \vdash \mathsf{fold}\ M : \mu\alpha.T$ (by a top-level application of* VAL FOLD*) and $\Gamma;\Delta_2 \vdash \mu\alpha.T <: \psi(V)$.*
(2) *If $\Gamma;\Delta \vdash \mu\alpha.T <: \mu\alpha.T'$, then there exists Δ' such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta'$ and $\Gamma;!\Delta' \vdash T\{\mu\alpha.T/\alpha\} <: T'\{\mu\alpha.T'/\alpha\}$.*
(3) *If $\Gamma;\Delta \vdash \mathsf{fold}\ M : \mu\alpha.T$, then there exist Δ' such that $\Gamma;\Delta \hookrightarrow \Gamma;!\Delta'$ and $\Gamma;!\Delta' \vdash M : T\{\mu\alpha.T/\alpha\}$.*
(4) *If $\Gamma;\Delta \vdash \mathsf{fold}\ M : \mu\alpha.T$, then $\Gamma;\Delta \vdash M : T\{\mu\alpha.T/\alpha\}$.*

B.6. Properties of extraction

We first present some simple, but useful properties of the extraction relation.

LEMMA B.36 (EXTRACTION AND FREE VALUES). *If $E \leadsto^{\widetilde{a}} [\Delta \mid D]$, then $fnfv(\Delta) \cup fnfv(D) \subseteq fnfv(E)$.*

LEMMA B.37 (EXTENDING EXTRACTION). *If $E \leadsto^{\widetilde{b}} [\Delta \mid D]$ and $a \notin fn(E)$, then $E \leadsto^{a,\widetilde{b}} [\Delta \mid D]$.*

LEMMA B.38 (RESTRICTING EXTRACTION). *If $E \leadsto^{\widetilde{a}} [\Delta \mid D]$ and $E \leadsto^{\widetilde{b}} [\Delta' \mid D']$ with $\{\widetilde{b}\} \subseteq \{\widetilde{a}\}$, then $D \leadsto^{\widetilde{b}} [\Delta'' \mid D']$, where $\Delta' = \Delta, \Delta''$.*

LEMMA B.39 (TRANSITIVITY OF EXTRACTION). *Let $E \leadsto^{\widetilde{b}} [\Delta' \mid E']$ and $E' \leadsto^{\widetilde{c}} [\Delta'' \mid E'']$, where $\{\widetilde{c}\} \subseteq \{\widetilde{b}\}$, then $E \leadsto^{\widetilde{c}} [\Delta', \Delta'' \mid E'']$.*

LEMMA B.40 (IDEMPOTENT EXTRACTION). *If $E \leadsto^{\widetilde{a}} [\Delta \mid D]$, then $D \leadsto^{\widetilde{a}} [\emptyset \mid D]$.*

The next result shows that heating preserves logic: if $E \Rightarrow E'$, then the formulas extracted from E are exactly the same of the formulas extracted from E'. Moreover, the purged expressions D and D' obtained after extracting the assumptions from E

and E' respectively are again related by heating. All this information is needed to show that heating preserves typing (Lemma B.46 below). In the following proofs we often write $E \rightsquigarrow [\Delta \mid D]$ whenever $E \rightsquigarrow^{\widetilde{a}} [\Delta \mid D]$ for some $\widetilde{a}$ clear from the context.

LEMMA B.41 (HEATING PRESERVES LOGIC).   *If $E \Rightarrow E'$ and $E \rightsquigarrow^{\widetilde{a}} [\Delta \mid D]$, then $E' \rightsquigarrow^{\widetilde{a}} [\Delta \mid D']$ for some D' such that $D \Rightarrow D'$. Moreover, the depth of the derivation of $D \Rightarrow D'$ equals that of $E \Rightarrow E'$.*

The next lemma is in the same spirit of Lemma B.41, but it predicates over the reduction relation rather than on heating and it is slightly more complicated. This is needed in the proof of Subject Reduction (Theorem B.48 below).

LEMMA B.42 (REDUCTION PRESERVES LOGIC).   *If $E \rightarrow E'$ and $E \rightsquigarrow^{\widetilde{a}} [\Delta \mid D]$, then $D \rightarrow D'$ and $E' \rightsquigarrow^{\widetilde{a}} [\Delta, \Delta' \mid D'']$ for some D', D'', Δ' such that $D' \rightsquigarrow^{\widetilde{a}} [\Delta' \mid D^*]$ with $D^* \Rightarrow D''$. Moreover, the depth of the derivation of $D \rightarrow D'$ equals that of $E \rightarrow E'$.*

B.7. Proof of subject reduction

In the proof of Lemmas B.44, B.45, B.46 and Theorem B.48 below we rely on an observation about the structure of the type derivations to simplify the formal reasoning and carry out the proofs. First, we consider an alternative formulation of typing for values, presented in Table XXV, which removes the non-structural rule (VAL REFINE). We also assume to keep the original typing rules for expressions.

Table XXV Alternative rules for typing values

VAL VAR REFINE
$$\frac{(x : T) \in \Gamma \qquad \Gamma; \Delta \vdash F\{x/y\}}{\Gamma; \Delta \vdash^{\mathsf{alt}} x : \{y : T \mid F\}}$$

VAL UNIT REFINE
$$\frac{\Gamma; \Delta \vdash F\{()/y\}}{\Gamma; \Delta \vdash^{\mathsf{alt}} () : \{y : \mathsf{unit} \mid F\}}$$

VAL FUN REFINE
$$\frac{(\Gamma; !\Delta_1) \bullet x : T \vdash^{\mathsf{alt}} E : U \qquad \Gamma; \Delta_2 \vdash F\{\lambda x. E/y\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, \Delta_2}{\Gamma; \Delta \vdash^{\mathsf{alt}} \lambda x. E : \{y : x : T \rightarrow U \mid F\}}$$

VAL PAIR REFINE
$$\frac{\Gamma; !\Delta_1 \vdash^{\mathsf{alt}} M : T \qquad \Gamma; !\Delta_2 \vdash^{\mathsf{alt}} N : U\{M/x\} \qquad \Gamma; \Delta_3 \vdash F\{(M, N)/y\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, !\Delta_2, \Delta_3}{\Gamma; \Delta \vdash^{\mathsf{alt}} (M, N) : \{y : x : T * U \mid F\}}$$

VAL INL REFINE
$$\frac{\Gamma; !\Delta_1 \vdash^{\mathsf{alt}} M : T \qquad \Gamma; !\Delta_1 \vdash U \qquad \Gamma; \Delta_2 \vdash F\{\mathsf{inl}\ M/y\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, \Delta_2}{\Gamma; \Delta \vdash^{\mathsf{alt}} \mathsf{inl}\ M : \{y : T + U \mid F\}}$$

VAL INR REFINE
$$\frac{\Gamma; !\Delta_1 \vdash^{\mathsf{alt}} M : U \qquad \Gamma; !\Delta_1 \vdash T \qquad \Gamma; \Delta_2 \vdash F\{\mathsf{inr}\ M/y\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, \Delta_2}{\Gamma; \Delta \vdash^{\mathsf{alt}} \mathsf{inr}\ M : \{y : T + U \mid F\}}$$

VAL FOLD REFINE
$$\frac{\Gamma; !\Delta_1 \vdash^{\mathsf{alt}} M : T\{\mu a. T/\alpha\} \qquad \Gamma; \Delta_2 \vdash F\{\mathsf{fold}\ M/y\} \qquad \Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, \Delta_2}{\Gamma; \Delta \vdash^{\mathsf{alt}} \mathsf{fold}\ M : \{y : \mu\alpha. T \mid F\}}$$

We can show that the original and the alternative formulation coincide.

LEMMA B.43 (ALTERNATIVE TYPING). $\Gamma; \Delta \vdash E : T$ *if and only if* $\Gamma; \Delta \vdash^{\mathsf{alt}} E : T$.

Now the idea is to appeal to the transitivity of both the subtyping relation (Lemma B.22) and the environment rewriting relation (Lemma B.8) to rearrange the structure of any type derivation constructed under the alternative typing rules.

Namely, we observe that for any expression E the general form of such a type derivation is as follows:

$$\frac{\Gamma;\Delta_1 \vdash^{\mathsf{alt}} E : T_1 \quad \Gamma;\Delta_2 \vdash T_1 <: T_2 \quad \Gamma;\Delta_3 \hookrightarrow \Gamma;\Delta_1, \Delta_2}{\frac{\vdots}{\frac{\Gamma;\Delta_{2n-1} \vdash^{\mathsf{alt}} E : T_{2n-1} \quad\quad \Gamma;\Delta_{2n} \vdash T_{2n-1} <: T \quad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_{2n-1}, \Delta_{2n}}{\Gamma;\Delta \vdash^{\mathsf{alt}} E : T}}}$$

where the last rule applied to derive $\Gamma;\Delta_1 \vdash^{\mathsf{alt}} E : T_1$ is not (EXP SUBSUM). Without loss of generality, we reorganize the derivation as follows:

$$\frac{\Gamma;\Delta_1 \vdash^{\mathsf{alt}} E : T_1 \quad \Gamma;\Delta^* \vdash T_1 <: T \quad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1, \Delta^*}{\Gamma;\Delta \vdash^{\mathsf{alt}} E : T}$$

with $\Delta^* = \Delta_2, \Delta_4, \ldots, \Delta_{2n}$. Notice that also derivations which do not use rule (EXP SUBSUM) can be rearranged as detailed, since the subtyping relation is reflexive. Moreover, given that original typing and alternative typing coincide by Lemma B.43, we note that the previous transformation can be applied to any type derivation.

Now we can show that extraction preserves typing: this is needed to show that heating preserves typing (Lemma B.46 below).

LEMMA B.44 (EXTRACTION PRESERVES TYPING). *If* $\Gamma;\Delta \vdash E : T$ *and* $E \rightsquigarrow^{\widetilde{a}} [\Delta' \mid E']$, *then* $\Gamma;\Delta, \Delta' \vdash E' : T$.

Similarly to the previous result, we can also show that *inverting* an extraction preserves typing: again, this is needed to prove that heating preserves typing (Lemma B.46 below).

LEMMA B.45 (INVERTING EXTRACTION PRESERVES TYPING). *Let* $E \rightsquigarrow^{\widetilde{b}} [\Delta' \mid E']$. *If* $\Gamma;\Delta, \Delta' \vdash E' : T$, *then* $\Gamma;\Delta \vdash E : T$.

The next result, sometimes called Subject Heating, shows that typing is preserved by heating. This is needed in the proof of the Subject Reduction theorem, since the reduction relation is closed under heating.

LEMMA B.46 (HEATING PRESERVES TYPING). *If* $\Gamma;\Delta \vdash E : T$ *and* $E \Rightarrow E'$, *then* $\Gamma;\Delta \vdash E' : T$.

The next simple lemma states that tautologies can be safely removed from any typing environment. This is used in some cases of the Subject Reduction proof, to deal with the logical formulas we explicitly introduce in the typing environment to make type-checking more precise (cf. EXP SPLIT).

LEMMA B.47 (REMOVING TAUTOLOGIES). *If* $\Gamma;\Delta, F \vdash E : T$ *and* $\emptyset \vdash F$, *then* $\Gamma;\Delta \vdash E : T$.

We can finally prove the Subject Reduction theorem. Its statement is remarkably simple: this is mainly due to our type system design, which discharges to the underlying affine logical framework all the complicated issues related to resource consumption. Thus, we do not need to explicitly track in the semantics which resources are consumed upon reduction, unlike to many other substructural type systems.

THEOREM B.48 (SUBJECT REDUCTION). *Let* $fv(E) = \emptyset$. *If* $\Gamma;\Delta \vdash E : T$ *and* $E \rightarrow E'$, *then* $\Gamma;\Delta \vdash E' : T$.

PROOF. By induction on the derivation of $E \rightarrow E'$. In the proof we implicitly appeal to Lemma B.5 and Lemma B.8 several times:

Case (RED FUN): assume $(\lambda x. E)\ N \to E\{N/x\}$ and $\Gamma; \Delta \vdash (\lambda x. E)\ N : T$. The typing judgement must follow by an instance of (EXP APPL) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash (\lambda x. E)\ N : U'\{N/x\}$ and $\Gamma; \Delta_B \vdash U'\{N/x\} <: T$ with:
—$\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
—$\Gamma; \Delta_A \hookrightarrow \Gamma; \Delta_1, \Delta_2$
—$\Gamma; \Delta_1 \vdash \lambda x. E : x : T' \to U'$
—$\Gamma; \Delta_2 \vdash N : T'$
By Lemma B.32 we know that $\Gamma; \Delta_1 \vdash \lambda x. E : x : T' \to U'$ implies $(\Gamma; \Delta_1) \bullet x : T' \vdash E : U'$. Now notice that $x \notin dom(\Gamma)$ by Lemma B.5, hence $x \notin fv(\Delta_1)$ by Lemma B.10. By applying Lemma B.28, we then get $\Gamma; \Delta_1, \Delta_2 \vdash E\{N/x\} : U'\{N/x\}$. Since $\Gamma; \Delta \hookrightarrow \Gamma; (\Delta_1, \Delta_2), \Delta_B$, the conclusion $\Gamma; \Delta \vdash E\{N/x\} : T$ follows by an application of (EXP SUBSUM).

Case (RED SPLIT): assume let $(x, y) = (M, N)$ in $E \to E\{M/x\}\{N/y\}$ and $\Gamma; \Delta \vdash$ let $(x, y) = (M, N)$ in $E : T$. The typing judgement must follow by an instance of (EXP SPLIT) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash$ let $(x, y) = (M, N)$ in $E : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:
—$\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
—$\Gamma; \Delta_A \hookrightarrow \Gamma; \Delta_1, \Delta_2$
—$\Gamma; \Delta_1 \vdash (M, N) : x : T' * U'$
—$(\Gamma; \Delta_2) \bullet x : T' \bullet y : U' \bullet !((x, y) = (M, N)) \vdash E : V$
—$\{x, y\} \cap fv(V) = \emptyset$
By Lemma B.33 we know that $\Gamma; \Delta_1 \vdash (M, N) : x : T' * U'$ implies:
—$\Gamma; \Delta_1 \hookrightarrow \Gamma; \Delta_{11}, \Delta_{12}$
—$\Gamma; \Delta_{11} \vdash M : T'$
—$\Gamma; \Delta_{12} \vdash N : U'\{M/x\}$
Now notice that $x \notin dom(\Gamma)$ by Lemma B.5, hence $x \notin fv(\Delta_2)$ by Lemma B.10. By applying Lemma B.28 twice and noting that $\{x, y\} \cap fv(V) = \emptyset$, we then get $\Gamma; \Delta_{11}, \Delta_{12}, \Delta_2, !((M, N) = (M, N)) \vdash E : V$. Since $\emptyset \vdash !((M, N) = (M, N))$, the latter judgement implies $\Gamma; \Delta_{11}, \Delta_{12}, \Delta_2 \vdash E : V$ by Lemma B.47. Since $\Gamma; \Delta \hookrightarrow \Gamma; (\Delta_{11}, \Delta_{12}, \Delta_2), \Delta_B$, the conclusion $\Gamma; \Delta \vdash E : T$ follows by (EXP SUBSUM).

Case (RED MATCH): assume match $h\ N$ with $h\ x$ then E else $E' \to E\{N/x\}$ and $\Gamma; \Delta \vdash$ match $h\ N$ with $h\ x$ then E else $E' : T$. The typing judgement must follow by an instance of (EXP MATCH) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash$ match $h\ N$ with $h\ x$ then E else $E' : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:
—$\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
—$\Gamma; \Delta_A \hookrightarrow \Gamma; \Delta_1, \Delta_2$
—$\Gamma; \Delta_1 \vdash h\ N : T'$
—$(\Gamma; \Delta_2) \bullet x : U' \bullet !(h\ x = h\ N) \vdash E : V$
—$\Gamma; \Delta_2 \vdash E' : V$
—$(h, T', U') \in \{(\mathsf{inl}, T_1 + T_2, T_1), (\mathsf{inr}, T_1 + T_2, T_2), (\mathsf{fold}, \mu\alpha. T_1, T_1\{\mu\alpha. T_1/\alpha\})\}$
According to the form of h, we invoke either Lemma B.34 or Lemma B.35. and we get $\Gamma; \Delta_1 \vdash N : U'$. Now we notice that $\Gamma; \Delta_2 \vdash E' : V$ implies $fnfv(V) \subseteq dom(\Gamma)$ by Lemma B.5, hence the fact that $x \notin dom(\Gamma)$ implies $x \notin fv(V)$. Moreover, $x \notin dom(\Gamma)$ implies $x \notin fv(\Delta_2)$ by Lemma B.10. By applying Lemma B.28 we then get $\Gamma; \Delta_1, \Delta_2, !(h\ N = h\ N) \vdash E\{N/x\} : V$. Since $\emptyset \vdash !(h\ N = h\ N)$, the latter judgement implies $\Gamma; \Delta_1, \Delta_2 \vdash E\{N/x\} : V$ by Lemma B.47. Since $\Gamma; \Delta \hookrightarrow \Gamma; (\Delta_1, \Delta_2), \Delta_B$, the conclusion $\Gamma; \Delta \vdash E\{N/x\} : T$ follows by (EXP SUBSUM).

Assume now match M with $h\ x$ then E else $E' \to E'$ with $M \neq h\ N$ for all N. The type derivation has the same structure as before, but for the obvious changes. Since $\Gamma; \Delta_2 \vdash E' : V$ and $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_2, \Delta_B$, the conclusion $\Gamma; \Delta \vdash E' : T$ follows by (EXP SUBSUM).

Case (RED EQ): assume we have $M = M \to$ true and $\Gamma; \Delta \vdash M = M : T$. The typing judgement must follow by an instance of (EXP EQ) after an instance of (EXP SUBSUM), hence it must be the case that:

$$\Gamma; \Delta_A \vdash M = M : \{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = M)\}$$

and:

$$\Gamma; \Delta_B \vdash \{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = M)\} <: T.$$

with $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$. Recall now that true $\triangleq$ inl( ) and bool $\triangleq$ unit + unit, so it is easy to show that we have $\Gamma; \Delta_A \vdash$ true : bool. Now we note that:

$$\Gamma; \emptyset \vdash !(\mathsf{true} = \mathsf{true} \multimap M = M),$$

thus we get $\Gamma; \Delta_A \vdash$ true : $\{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = M)\}$ by (VAL REFINE) and the conclusion $\Gamma; \Delta \vdash$ true : T follows by an application of (EXP SUBSUM).
Assume, instead, that $M = N \to$ false with $M \neq N$ and $\Gamma; \Delta \vdash M = N : T$. The typing judgement must follow by an instance of (EXP EQ) after an instance of (EXP SUBSUM), hence it must be the case that:

$$\Gamma; \Delta_A \vdash M = N : \{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = N)\}$$

and:

$$\Gamma; \Delta_B \vdash \{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = N)\} <: T.$$

with $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$. Now we note that:

$$\Gamma; \emptyset \vdash !(\mathsf{false} = \mathsf{true} \multimap M = N),$$

thus we get $\Gamma; \Delta_A \vdash$ false : $\{x : \mathsf{bool} \mid !(x = \mathsf{true} \multimap M = N)\}$ by (VAL REFINE) and the conclusion $\Gamma; \Delta \vdash$ false : T follows by an application of (EXP SUBSUM).
Case (RED COMM): assume $a!M \stackrel{\cdot}{\shortmid} a? \to M$ and $\Gamma; \Delta \vdash a!M \stackrel{\cdot}{\shortmid} a? : T$. The typing judgement must follow by an instance of (EXP FORK) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash a!M \stackrel{\cdot}{\shortmid} a? : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:
— $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
— $a!M \rightsquigarrow [\emptyset \mid a!M]$
— $a? \rightsquigarrow [\emptyset \mid a?]$
— $\Gamma; \Delta_A \hookrightarrow \Gamma; \Delta_1, \Delta_2$
— $\Gamma; \Delta_1 \vdash a!M : U$
— $\Gamma; \Delta_2 \vdash a? : V$
We notice that $\Gamma; \Delta_1 \vdash a!M : U$ must follow by an instance of (EXP SEND) after an instance of (EXP SUBSUM), hence:
— $\Gamma; \Delta_1 \hookrightarrow \Gamma; \Delta_{11}, \Delta_{12}$
— $\Gamma; \Delta_{11} \vdash a!M : \mathsf{unit}$
— $\Gamma; \Delta_{12} \vdash \mathsf{unit} <: U$
— $(a \updownarrow T') \in \Gamma$
— $\Gamma; \Delta_{11} \vdash M : T'$
We also notice that $\Gamma; \Delta_2 \vdash a? : V$ must follow by an instance of (EXP RECV) after an instance of (EXP SUBSUM), hence:
— $\Gamma; \Delta_2 \hookrightarrow \Gamma; \Delta_{21}, \Delta_{22}$
— $\Gamma; \Delta_{21} \vdash a? : T'$, since $(a \updownarrow T') \in \Gamma$
— $\Gamma; \Delta_{22} \vdash T' <: V$
Thus we get $\Gamma; \Delta_{11}, \Delta_{22} \vdash M : V$ by (EXP SUBSUM). Since $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_{11}, \Delta_{22}, \Delta_B$, the conclusion $\Gamma; \Delta \vdash M : T$ follows by an application of (EXP SUBSUM).
Case (RED LET VAL): assume let $x = M$ in $E \to E\{M/x\}$ and $\Gamma; \Delta \vdash$ let $x = M$ in $E : T$. The typing judgement must follow by an instance of (EXP LET) after an instance of

(EXP SUBSUM). Notice that $M \rightsquigarrow [\emptyset \mid M]$, hence it must be the case that $\Gamma; \Delta_A \vdash$ let $x = M$ in $E : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:

— $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
— $\Gamma; \Delta_A \hookrightarrow \Gamma; \Delta_1, \Delta_2$
— $\Gamma; \Delta_1 \vdash M : U$
— $(\Gamma; \Delta_2) \bullet x : U \vdash E : V$
— $x \notin fv(V)$

Now notice that $x \notin dom(\Gamma)$ by Lemma B.5, hence $x \notin fv(\Delta_2)$ by Lemma B.10. By applying Lemma B.28 and noting that $x \notin fv(V)$, we then get $\Gamma; \Delta_1, \Delta_2 \vdash E\{M/x\} :$ V. Since $\Gamma; \Delta \hookrightarrow \Gamma; (\Delta_1, \Delta_2), \Delta_B$, the conclusion $\Gamma; \Delta \vdash E\{M/x\} : T$ follows by an application of (EXP SUBSUM).

Case (RED LET): assume let $x = E$ in $E'' \rightarrow$ let $x = E'$ in E'' with $E \rightarrow E'$ and $\Gamma; \Delta \vdash$ let $x = E$ in $E'' : T$. The typing judgement must follow by an instance of (EXP LET) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash$ let $x = E$ in $E'' : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:

— $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
— $E \rightsquigarrow [\Delta' \mid D]$
— $\Gamma; \Delta_A, \Delta' \hookrightarrow \Gamma; \Delta_1, \Delta_2$
— $\Gamma; \Delta_1 \vdash D : U$
— $(\Gamma; \Delta_2) \bullet x : U \vdash E'' : V$

By Lemma B.42 we know that $E \rightarrow E'$ and $E \rightsquigarrow [\Delta' \mid D]$ imply that there exist D', Δ'', D'', D^* such that $D \rightarrow D'$ and $E' \rightsquigarrow [\Delta', \Delta'' \mid D'']$ with $D' \rightsquigarrow [\Delta'' \mid D^*]$ and $D^* \Rightarrow$ D''. Since Lemma B.42 is depth-preserving, we can apply the inductive hypothesis and get $\Gamma; \Delta_1 \vdash D' : U$. Given that $D' \rightsquigarrow [\Delta'' \mid D^*]$ and $\Gamma; \Delta_1 \vdash D' : U$, we get $\Gamma; \Delta_1, \Delta'' \vdash D^* : U$ by Lemma B.44. Since $D^* \Rightarrow D''$ and $\Gamma; \Delta_1, \Delta'' \vdash D^* : U$, we get $\Gamma; \Delta_1, \Delta'' \vdash D'' : U$ by Lemma B.46. Hence, we have:

— $E' \rightsquigarrow [\Delta', \Delta'' \mid D'']$
— $\Gamma; \Delta_A, \Delta', \Delta'' \hookrightarrow \Gamma; (\Delta_1, \Delta''), \Delta_2$
— $\Gamma; \Delta_1, \Delta'' \vdash D'' : U$
— $(\Gamma; \Delta_2) \bullet x : U \vdash E'' : V$

We can then apply rule (EXP LET) to get $\Gamma; \Delta_A \vdash$ let $x = E'$ in $E'' : V$. The conclusion $\Gamma; \Delta \vdash$ let $x = E'$ in $E'' : T$ follows by (EXP SUBSUM).

Case (RED RES): assume $(\nu a)E \rightarrow (\nu a)E'$ with $E \rightarrow E'$ and $\Gamma; \Delta \vdash (\nu a)E : T$. The typing judgement must follow by an instance of (EXP RES) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash (\nu a)E : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:

— $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
— $E \rightsquigarrow^a [\Delta' \mid D]$
— $\Gamma, a \updownarrow U; \Delta_A, \Delta' \vdash D : V$

By Lemma B.42 we know that $E \rightarrow E'$ and $E \rightsquigarrow^a [\Delta' \mid D]$ imply that there exist D', Δ'', D'', D^* such that $D \rightarrow D'$ and $E' \rightsquigarrow^a [\Delta', \Delta'' \mid D'']$ with $D' \rightsquigarrow^a [\Delta'' \mid D^*]$ and $D^* \Rightarrow D''$. Since Lemma B.42 is depth-preserving, we can apply the inductive hypothesis and get $\Gamma, a \updownarrow U; \Delta_A, \Delta' \vdash D' : V$. Given that $D' \rightsquigarrow^a [\Delta'' \mid D^*]$ and $\Gamma, a \updownarrow U; \Delta_A, \Delta' \vdash D' : V$, we get $\Gamma, a \updownarrow U; \Delta_A, \Delta', \Delta'' \vdash D^* : V$ by Lemma B.44. By Lemma B.46 we get $\Gamma, a \updownarrow U; \Delta_A, \Delta', \Delta'' \vdash D'' : V$. Hence, we have:

— $E' \rightsquigarrow^a [\Delta', \Delta'' \mid D'']$
— $\Gamma, a \updownarrow U; \Delta_A, \Delta', \Delta'' \vdash D'' : V$

We can then apply rule (EXP RES) to get $\Gamma; \Delta_A \vdash (\nu a)E' : V$. The conclusion $\Gamma; \Delta \vdash (\nu a)E' : T$ follows by (EXP SUBSUM).

Case (RED FORK 1): assume $E \wr E'' \rightarrow E' \wr E''$ with $E \rightarrow E'$ and $\Gamma; \Delta \vdash E \wr E'' : T$. The typing judgement must follow by an instance of (EXP FORK) after an instance of (EXP SUBSUM), hence it must be the case that $\Gamma; \Delta_A \vdash E \wr E'' : V$ and $\Gamma; \Delta_B \vdash V <: T$ with:

— $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_A, \Delta_B$
— $E \rightsquigarrow [\Delta' \mid D_1]$
— $E'' \rightsquigarrow [\Delta'' \mid D_2]$
— $\Gamma; \Delta_A, \Delta', \Delta'' \hookrightarrow \Gamma; \Delta_1, \Delta_2$
— $\Gamma; \Delta_1 \vdash D_1 : U$
— $\Gamma; \Delta_2 \vdash D_2 : V$

By Lemma B.42 we know that $E \rightarrow E'$ and $E \rightsquigarrow [\Delta' \mid D_1]$ imply that there exist D_1', Δ^*, D'', D^* such that $D_1 \rightarrow D_1'$ and $E' \rightsquigarrow [\Delta', \Delta^* \mid D'']$ with $D_1' \rightsquigarrow [\Delta^* \mid D^*]$ and $D^* \Rrightarrow D''$. Since Lemma B.42 is depth-preserving, we can apply the inductive hypothesis and get $\Gamma; \Delta_1 \vdash D_1' : U$. Given that $D_1' \rightsquigarrow [\Delta^* \mid D^*]$ and $\Gamma; \Delta_1 \vdash D_1' : U$, we get $\Gamma; \Delta_1, \Delta^* \vdash D^* : U$ by Lemma B.44. By Lemma B.46 we get $\Gamma; \Delta_1, \Delta^* \vdash D'' : U$. Hence, we have:
— $E' \rightsquigarrow [\Delta', \Delta^* \mid D'']$
— $E'' \rightsquigarrow [\Delta'' \mid D_2]$
— $\Gamma; \Delta_A, \Delta', \Delta^*, \Delta'' \hookrightarrow \Gamma; (\Delta_1, \Delta^*), \Delta_2$
— $\Gamma; \Delta_1, \Delta^* \vdash D'' : U$
— $\Gamma; \Delta_2 \vdash D_2 : V$

We can then apply rule (EXP FORK) to get $\Gamma; \Delta_A \vdash E' \curvearrowright E'' : V$. The conclusion $\Gamma; \Delta \vdash E' \curvearrowright E'' : T$ follows by (EXP SUBSUM).

Case (RED FORK 2): analogous to the previous case.

Case (RED HEAT): assume $E \rightarrow E'$ with $E \Rrightarrow D$, $D \rightarrow D'$ and $D' \rightarrow E'$. Assume further that $\Gamma; \Delta \vdash E : T$. By Lemma B.46 we have $\Gamma; \Delta \vdash D : T$. By inductive hypothesis $\Gamma; \Delta \vdash D' : T$, hence $\Gamma; \Delta \vdash E' : T$ again by Lemma B.46.

□

B.8. Proof of (robust) safety

We first show that well-typed structures are statically safe.

LEMMA B.49 (STATIC SAFETY). *If $\varepsilon; \emptyset \vdash \mathbf{S} : T$, then $\mathbf{S}$ is statically safe.*

The safety theorem below states that any well-typed expression is safe. Its proof is simple and relies on the previous results.

RESTATEMENT 3 (OF THEOREM 6.1). *If $\varepsilon; \emptyset \vdash E : T$, then E is safe.*

PROOF. In order to prove that E is safe it suffices to show that, for all expressions E' and structures $\mathbf{S}$ such that $E \rightarrow^* E'$ and $E' \Rrightarrow \mathbf{S}$, it holds that $\mathbf{S}$ is statically safe.

By Theorem B.48, $\varepsilon; \emptyset \vdash E : T$ implies $\varepsilon; \emptyset \vdash E' : T$. By Lemma B.46, $E' \Rrightarrow \mathbf{S}$ implies $\varepsilon; \emptyset \vdash \mathbf{S} : T$. We can conclude that $\mathbf{S}$ is statically safe by Lemma B.49. □

The next lemma is important to show that any opponent is trivially well-typed: it identifies Un with a number of structural types built around Un itself.

LEMMA B.50 (UNIVERSAL TYPE). *If $\Gamma; \emptyset \vdash \diamond$, then $\Gamma; \emptyset \vdash T <:> \mathsf{Un}$ for all $T \in \{\mathsf{unit}, x : \mathsf{Un} \rightarrow \mathsf{Un}, x : \mathsf{Un} * \mathsf{Un}, \mathsf{Un} + \mathsf{Un}, \mu\alpha. \mathsf{Un}\}$.*

We can now show that any opponent is well-typed. The statement is slightly more general than expected, since we appeal to inductive reasoning in the proof.

LEMMA B.51 (OPPONENT TYPABILITY). *Let $\Gamma; \emptyset \vdash \diamond$. Let O be an expression that does not contain any assumption or assertion such that $(a \updownarrow \mathsf{Un}) \in \Gamma$ for each $a \in fn(O)$ and $(x : \mathsf{Un}) \in \Gamma$ for each $x \in fv(O)$, then $\Gamma; \emptyset \vdash O : \mathsf{Un}$.*

Finally, we can prove our main result of interest: if an expression E is assigned type Un by our type system, then it is robustly safe. The proof is an easy consequence of Theorem 6.1 and Lemma B.51.

RESTATEMENT 4 (OF THEOREM 6.2). *If* $\varepsilon; \emptyset \vdash E :$ Un*, then* E *is robustly safe.*

PROOF. Consider an arbitrary opponent O, we need to show that the application $O\ E$ is safe. Recall that:

$$O\ E \triangleq \mathsf{let}\ f = O\ \mathsf{in}\ \mathsf{let}\ x = E\ \mathsf{in}\ f\ x.$$

Let $\Gamma = a_1 \updownarrow \mathsf{Un}, \ldots, a_n \updownarrow \mathsf{Un}$ with $fn(O) = \{a_1, \ldots, a_n\}$. Since the opponent O is closed by definition, by Lemma B.51 we know that $\Gamma; \emptyset \vdash O :$ Un. We can apply (EXP SUBSUM) and Lemma B.50 to derive:

$$\Gamma; \emptyset \vdash O : \mathsf{Un} \to \mathsf{Un}. \tag{1}$$

We can apply Lemma B.7 to $\varepsilon; \emptyset \vdash E :$ Un and get $\Gamma; \emptyset \vdash E :$ Un. Assume now $E \rightsquigarrow [\Delta \mid D]$, by Lemma B.44 we have $\Gamma; \Delta \vdash D :$ Un. By Lemma B.7 we then get:

$$\Gamma, f : \mathsf{Un} \to \mathsf{Un}; \Delta \vdash D : \mathsf{Un}. \tag{2}$$

Since $O \rightsquigarrow^{\emptyset} [\emptyset \mid O]$, we can construct the following type derivation:

$$\cfrac{(1) \cfrac{\cdots}{\Gamma; \emptyset \vdash O : \mathsf{Un} \to \mathsf{Un}} \qquad (2) \cfrac{\cfrac{\cdots}{\Gamma, f : \mathsf{Un} \to \mathsf{Un}; \Delta \vdash D : \mathsf{Un}} \qquad \cfrac{\cdots}{\Gamma, f : \mathsf{Un} \to \mathsf{Un}, x : \mathsf{Un}; \emptyset \vdash f\ x : \mathsf{Un}}\ \text{EXP APPL}}{\cfrac{\Gamma, f : \mathsf{Un} \to \mathsf{Un}; \emptyset \vdash \mathsf{let}\ x = E\ \mathsf{in}\ f\ x : \mathsf{Un}}{}\ \text{EXP LET}}}{\Gamma; \emptyset \vdash \mathsf{let}\ f = O\ \mathsf{in}\ \mathsf{let}\ x = E\ \mathsf{in}\ f\ x : \mathsf{Un}}\ \text{EXP LET}$$

Since $O\ E \rightsquigarrow^{\widetilde{b}} [\emptyset \mid O\ E]$ for all $\widetilde{b}$, we can get $\varepsilon; \emptyset \vdash (\nu a_1) \ldots (\nu a_n)(O\ E) :$ Un by applying n times rule (EXP RES) to the conclusion of the derivation above. By Theorem 6.1, we then know that $(\nu a_1) \ldots (\nu a_n)(O\ E)$ is safe. Since restrictions do not affect safety, we can conclude. □

C. SOUNDNESS AND COMPLETENESS OF ALGORITHMIC TYPING

In this section we prove the soundness (Theorem 10.1) and completeness (Theorem 10.2) of the algorithmic variant of our type system.

C.1. Logical properties

We begin by showing some important properties of the logic that play a pivotal role in the bottom-up construction of the unique proof obligation in the algorithmic type system and the corresponding proofs of soundness and completeness.

We use the following convenient notation to denote all logical entailment rules that modify the set of premises.

Definition C.1 (*Left Rules* $\vdash^{\mathsf{L}}$). We say $\Delta \vdash^{\mathsf{L}} F$ if the last applied logical entailment rule is a rule of the form (R-LEFT) or (CONTR) or (WEAK).

LEMMA C.2 (IMPLICATION).

(1) *For all* Δ, F, F' *we have that* $\Delta \vdash F \multimap F'$ *iff* $\Delta, F \vdash F'$.
(2) *For all* Γ, Δ, F, F' *we have that* $\Gamma; \Delta \vdash F \multimap F'$ *iff* $\Gamma; \Delta, F \vdash F'$.

LEMMA C.3 (UNIVERSAL QUANTIFICATION). *It holds that:*

(1) *For all* x, Δ, F *such that* $x \notin fv(\Delta)$*, we have that* $\Delta \vdash F$ *iff* $\Delta \vdash \forall x.F$.
(2) *For all* Γ, x, T, Δ, F *such that* $\Gamma, x : \psi(T); \Delta \vdash \diamond$ *and* $x \notin fnfv(\Delta)$*, we have that* $\Gamma, x : \psi(T); \Delta \vdash F$ *iff* $\Gamma; \Delta \vdash \forall x.F$.

Table XXVI Intermediate Subtyping $<:_{\mathsf{alg}}$

(SUB REFL *)
$$\frac{\Gamma;\Delta \vdash T \qquad T \in \{\mathsf{unit}, \alpha\}}{\Gamma;\Delta \vdash T <:_{\mathsf{alg}} T}$$

(SUB PUB TNT *)
$$\frac{\Gamma;\Delta_1 \vdash T :: \mathsf{pub} \qquad \Gamma;\Delta_2 \vdash U :: \mathsf{tnt} \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2 \qquad T \neq_\top U}{\Gamma;\Delta \vdash T <:_{\mathsf{alg}} U}$$

(SUB FUN *)
$$\frac{\begin{array}{c}\Gamma;!\Delta_1 \vdash \overline{T'} <:_{\mathsf{alg}} \overline{T} \\ \Gamma, x : \psi(T');!\Delta_2 \vdash \overline{U} <:_{\mathsf{alg}} \overline{U'} \\ \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2\end{array}}{\Gamma;\Delta \vdash x : \overline{T} \to \overline{U} <:_{\mathsf{alg}} x : \overline{T'} \to \overline{U'}}$$

(SUB PAIR *)
$$\frac{\begin{array}{c}\Gamma;!\Delta_1 \vdash \overline{T} <:_{\mathsf{alg}} \overline{T'} \\ \Gamma, x : \psi(T);!\Delta_2 \vdash \overline{U} <:_{\mathsf{alg}} \overline{U'} \\ \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2\end{array}}{\Gamma;\Delta \vdash x : \overline{T} * \overline{U} <:_{\mathsf{alg}} x : \overline{T'} * \overline{U'}}$$

(SUB SUM *)
$$\frac{\begin{array}{c}\Gamma;!\Delta_1 \vdash \overline{T} <:_{\mathsf{alg}} \overline{T'} \qquad \Gamma;!\Delta_2 \vdash \overline{U} <:_{\mathsf{alg}} \overline{U'} \\ \Gamma;\Delta \hookrightarrow \Gamma;!\Delta_1,!\Delta_2\end{array}}{\Gamma;\Delta \vdash \overline{T} + \overline{U} <:_{\mathsf{alg}} \overline{T'} + \overline{U'}}$$

(SUB POS REC *)
$$\frac{\begin{array}{c}\Gamma,\alpha;!\Delta' \vdash \overline{T} <:_{\mathsf{alg}} \overline{T'} \\ \alpha \text{ occurs only positively in } \overline{T} \text{ and } \overline{T'} \\ T \neq T' \qquad \Gamma;\Delta \hookrightarrow \Gamma;!\Delta'\end{array}}{\Gamma;\Delta \vdash \mu\alpha.\overline{T} <:_{\mathsf{alg}} \mu\alpha.\overline{T'}}$$

(SUB REFL REC *)
$$\frac{\Gamma;\Delta \vdash \mu\alpha.T}{\Gamma;\Delta \vdash \mu\alpha.T <:_{\mathsf{alg}} \mu\alpha.T}$$

(SUB PUB TNT REC *)
$$\frac{\begin{array}{c}\Gamma;\Delta_1 \vdash \mu\alpha.T :: \mathsf{pub} \qquad \Gamma;\Delta_2 \vdash \mu\alpha.U :: \mathsf{tnt} \\ s = \mathsf{SPT} \oplus s' = \mathsf{SPT} \qquad \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2\end{array}}{\Gamma;\Delta \vdash (\mu\alpha.T)_s <:_{\mathsf{alg}} (\mu\alpha.U)_{s'}}$$

(SUB REFINE *)
$$\frac{\begin{array}{c}\Gamma;\Delta_1 \vdash \psi(\overline{T}) <:_{\mathsf{alg}} \psi(\overline{U}) \qquad \Gamma, y : \psi(T);\Delta_2, forms(y : T) \vdash forms(y : U) \\ \Gamma;\Delta \hookrightarrow \Gamma;\Delta_1,\Delta_2 \qquad \overline{T} \text{ and/or } \overline{U} \text{ refined}\end{array}}{\Gamma;\Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}}$$

Notation: We use T to denote the non-annotated counterpart $\langle\overline{T}\rangle$ of the annotated type $\overline{T}$.

C.2. Soundness and completeness of the algorithmic judgements

LEMMA C.4 (SOUNDNESS AND COMPLETENESS OF ALGORITHMIC WELL-FORMEDNESS).
For all Γ, the following holds true:

(1) $\Gamma \vdash_{\mathsf{alg}} \diamond$ *iff* $\Gamma;\emptyset \vdash \diamond$
(2) *for all T,* $\Gamma \vdash_{\mathsf{alg}} T$ *iff* $\Gamma;\emptyset \vdash T$

LEMMA C.5 (SOUNDNESS AND COMPLETENESS OF ALGORITHMIC KINDING).
For all Γ, T, k, the following holds true:

(1) for all F, Δ such that $\Gamma \vdash_{\mathsf{alg}} T :: k; F$ and $\Gamma;\Delta \vdash F$, we have that $\Gamma;\Delta \vdash T :: k$.
(2) for all Δ such that $\Gamma;\Delta \vdash T :: k$, there exists F such that $\Gamma \vdash_{\mathsf{alg}} T :: k; F$ and $\Gamma;\Delta \vdash F$;

In order to prove soundness and completeness of subtyping we will proceed in two steps: we first introduce an intermediate algorithmic variant of the standard subtyping relation $\Gamma;\Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}$ for annotated types that extends the standard one by adding the side conditions in (SUB REFL) and (SUB PUB TNT) present in algorithmic subtyping and provides three disjoint rules for subtyping two iso-recursive types, following the insights given in Section 10.5.

We will then show that we can find annotations to prove the standard subtyping and intermediate subtyping equivalent and show the soundness and completeness of algorithmic subtyping with respect to intermediate subtyping $<:_{\mathsf{alg}}$.

The full definition of the intermediate subtyping rules can be found in Table XXVI. The rules make use of the previously introduced annotated types $\overline{T}$ that might contain type annotation SPT. As in algorithmic subtyping, we assume the function ψ to extend

to annotated types and we write $T = \langle \overline{T} \rangle$ to denote the type that results from erasing all type annotations from $\overline{T}$. We say T and $\overline{T}$ are equal up to type annotations.

The soundness and completeness proof of intermediate subtyping makes use of the following propositions and lemmas. We write $\overline{T} = \{x_m : \ldots \{x_2 : \{x_1 : \overline{U} \mid F_1\} \mid F_2\} \ldots \mid F_m\}$ to denote nested (annotated) refinement types, for $m = 0$ this notation simply denotes the annotated type $\overline{U}$.

The following proposition states that all types are also annotated types by construction, which we will use implicitly throughout the following proofs.

PROPOSITION C.6 (TYPES AND ANNOTATED TYPES). *Let T be a type. Then T is also an annotated type, such that $\langle T \rangle = T$.*

LEMMA C.7 (REFINEMENT ERASURE OF ANNOTATED TYPES). *For all types T and annotated types $\overline{T}$ such that $T = \langle \overline{T} \rangle$ it holds that $\psi(T) = \psi(\langle \overline{T} \rangle) = \langle \psi(\overline{T}) \rangle$.*

LEMMA C.8 (NESTED REFINEMENTS). *For all types T it holds that there exist $m \geq 0$ and T_{min} and $F_1, \ldots, F_m, x_1, \ldots, x_m$ (if $m > 0$) such that $T = \{x_m : \ldots \{x_2 : \{x_1 : T_{min} \mid F_1\} \mid F_2\} \ldots \mid F_m\}$ and $\psi(T) = \psi(T_{min}) = T_{min}$.*

LEMMA C.9 (ANNOTATED REFINEMENT TYPES). *For all types $T = \{x_m : \ldots \{x_2 : \{x_1 : T_{min} \mid F_1\} \mid F_2\} \ldots \mid F_m\}$, where $\psi(T) = \psi(T_{min}) = T_{min}$, and all annotated types $\overline{T_{min}}$ and $\overline{T} = \{x_m : \ldots \{x_2 : \{x_1 : \overline{T_{min}} \mid F_1\} \mid F_2\} \ldots \mid F_m\}$ such that $T_{min} = \langle \overline{T_{min}} \rangle$ it holds that*

— $\langle \overline{T} \rangle = T$ and
— $\psi(\overline{T}) = \overline{T_{min}}$.

LEMMA C.10 (SOUNDNESS AND COMPLETENESS OF INTERMEDIATE SUBTYPING). *For all Γ, Δ, T, U it holds that:*

(1) *If $\Gamma; \Delta \vdash T$ then $\Gamma; \Delta \vdash T <:_{\mathsf{alg}} T$.*
(2) *If there exist Δ_1, Δ_2 such that $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_1, \Delta_2$ and $\Gamma; \Delta_1 \vdash T :: \mathsf{pub}$ and $\Gamma; \Delta_2 \vdash U :: \mathsf{tnt}$ then there exist annotated types $\overline{T}$ and $\overline{U}$ such that $T = \langle \overline{T} \rangle$, $U = \langle \overline{U} \rangle$, and $\Gamma; \Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}$.*
(3) *If $\Gamma; \Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}$ and $T = \langle \overline{T} \rangle$, $U = \langle \overline{U} \rangle$, then $\Gamma; \Delta \vdash T <: U$.*
(4) *If $\Gamma; \Delta \vdash T <: U$ then there exist $\overline{T}, \overline{U}$ such that $T = \langle \overline{T} \rangle$, $U = \langle \overline{U} \rangle$, and $\Gamma; \Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}$.*

LEMMA C.11 (SOUNDNESS AND COMPLETENESS OF ALGORITHMIC SUBTYPING).

(1) *For all $\Gamma, \Delta, \overline{T}, \overline{U}$, the following holds true:*
    (a) *For all F such that $\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{U}; F$ and $\Gamma; \Delta \vdash F$, we have that $\Gamma; \Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}$.*
    (b) *If $\Gamma; \Delta \vdash \overline{T} <:_{\mathsf{alg}} \overline{U}$, then there exists F such that $\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{U}; F$ and $\Gamma; \Delta \vdash F$;*
(2) *For all Γ, Δ, T, U, the following holds true:*
    (a) *For all $\overline{T}, \overline{U}, F$ such that $\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{U}; F$ and $\Gamma; \Delta \vdash F$ and $T = \langle \overline{T} \rangle$, $U = \langle \overline{U} \rangle$, we have that $\Gamma; \Delta \vdash T <: U$.*
    (b) *If $\Gamma; \Delta \vdash T <: U$ then there exist $\overline{T}, \overline{U}, F$ such that $T = \langle \overline{T} \rangle$, $U = \langle \overline{U} \rangle$, and $\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{U}; F$ and $\Gamma; \Delta \vdash F$.*

The following lemma is used in the proof of soundness and completeness of algorithmic typing and states that for algorithmic subtyping type annotations need only occur in *either* the sub- or supertype.

LEMMA C.12 (ONE-SIDED TYPE ANNOTATIONS). *For all* $\Gamma, \overline{T}, \overline{U}, T, UF$ *such that* $\Gamma \vdash_{\mathsf{alg}} \overline{T} <: \overline{U}; F$ *such that* $T = \langle \overline{T} \rangle$ *and* $U = \langle \overline{U} \rangle$ *it holds that:*

(1) *there exist* $\overline{U}'$ *such that* $\langle \overline{U} \rangle = \langle \overline{U}' \rangle = U$ *and* $\Gamma \vdash_{\mathsf{alg}} T <: \overline{U}'; F;$
(2) *there exist* $\overline{T}'$ *such that* $\langle \overline{T} \rangle = \langle \overline{T}' \rangle = T$ *and* $\Gamma \vdash_{\mathsf{alg}} \overline{T}' <: U; F.$

The following proposition is used in the proof of soundness and completeness of algorithmic typing and states that extraction is unaffected by typing annotations.

PROPOSITION C.13 (ANNOTATED EXTRACTION). *For all* $\widetilde{a}, \Delta$ *it holds that:*

(1) *for all* $\overline{E}, \overline{D}$ *such that* $\overline{E} \rightsquigarrow^{\widetilde{a}} [\Delta \mid \overline{D}]$ *it must be the case that* $\langle \overline{E} \rangle \rightsquigarrow^{\widetilde{a}} [\Delta \mid \langle \overline{D} \rangle];$
(2) *for all* $E, D, \overline{D}$ *such that* $D = \langle \overline{D} \rangle$ *and* $E \rightsquigarrow^{\widetilde{a}} [\Delta \mid D]$ *it must be the case that there exists an annotated expression* $\overline{E}$ *such that* $E = \langle \overline{E} \rangle$ *and* $\overline{E} \rightsquigarrow^{\widetilde{a}} [\Delta \mid \overline{D}].$

LEMMA C.14 (TYPING TRUTH ASSUMPTION). *For all* Γ, Δ, T *such that* $\Gamma; \Delta \vdash$ assume $\mathbf{1} : T$ *it holds that* $\Gamma; \emptyset \vdash$ assume $\mathbf{1} : T$ *and* $\Gamma; \Delta \vdash$ unit $<: T.$

RESTATEMENT 5 (OF THEOREMS 10.1 AND 10.2). *For all* Γ, Δ, T, *the following holds true:*

(1) *for all* $\overline{E}, F$ *such that* $\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F$ *and* $\Gamma; \Delta \vdash F$, *we have that* $\Gamma; \Delta \vdash \langle \overline{E} \rangle : T;$
(2) *for all* E *such that* $\Gamma; \Delta \vdash E : T$, *there exist* $\overline{E}, F$ *such that* $\langle \overline{E} \rangle = E$, $\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F$, *and* $\Gamma; \Delta \vdash F.$

PROOF.

(1) The proof proceeds by induction on the length of $\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F$. The base cases are (VAL VAR ALG), (VAL UNIT ALG), (EXP TRUE ALG), (EXP RECV ALG), and (EXP ASSERT ALG): in all of these cases $\overline{E} = \langle \overline{E} \rangle$ (meaning $\overline{E}$ does not contain annotations) and they follow by an inspection of the typing rules and by Lemma C.4. We show the induction cases in the following. Most cases follow a very similar structure so we show detailed examples for standard proof strategies and omit the details for analogous cases.

Case (VAL FUN ALG): $\Gamma \vdash_{\mathsf{alg}} \lambda x : T_1. \overline{D} : x : T_1 \to T_2; !\forall x.(forms(x : T_1) \multimap F')$ is proved by $\Gamma, x : \psi(T_1) \vdash_{\mathsf{alg}} \overline{D} : T_2; F'$. We also know that $\Gamma; \Delta \vdash !\forall x.(forms(x : T_1) \multimap F').$

To show: $\Gamma; \Delta \vdash \langle \lambda x : T. \overline{D} \rangle : x : T_1 \to T_2$. We first note that $\langle \lambda x : T. \overline{D} \rangle$ is equal to $\lambda x. \langle \overline{D} \rangle$. By (REWRITE), (DERIVE), and Lemma B.5 we know that $\Gamma; \Delta \hookrightarrow \Gamma; !\forall x.(forms(x : T_1) \multimap F')$ and $\Gamma; !\forall x.(forms(x : T_1) \multimap F') \vdash \forall x.(forms(x : T_1) \multimap F')$ by (IDENT), (!-LEFT), and (DERIVE).

Without loss of generality, let us assume $x \notin dom(\Gamma)$ and, thus, $x \notin fnfv(!\forall x.(forms(x : T_1) \multimap F'))$. (This assumption can be fulfilled by α-renaming x if necessary.)

By Lemma B.5, we can easily see that $\Gamma, x : \psi(T_1); !\forall x.(forms(x : T_1) \multimap F') \vdash \diamond.$ By Lemma C.3, $\Gamma, x : \psi(T_1); !\forall x.(forms(x : T_1) \multimap F') \vdash forms(x : T_1) \multimap F'$. By Lemma C.2, $\Gamma, x : \psi(T_1); !\forall x.(forms(x : T_1) \multimap F'), forms(x : T_1) \vdash F'.$

By induction hypothesis, $\Gamma, x : \psi(T_1); !\forall x.(forms(x : T_1) \multimap F'), forms(x : T_1) \vdash \langle \overline{D} \rangle : T_2.$

The result follows by an application of (VAL FUN).

Case (VAL PAIR ALG): $\Gamma \vdash_{\mathsf{alg}} (\overline{M}, \overline{N}) : x : T_1 * T_2; !F_1 \otimes !F_2$ is proved by $\Gamma \vdash_{\mathsf{alg}} \overline{M} : T_1; F_1$ and $\Gamma \vdash_{\mathsf{alg}} \overline{N} : T_2\{M/x\}; F_2$, where $M := \langle \overline{M} \rangle$. We also know that $\Gamma; \Delta \vdash !F_1 \otimes !F_2.$

To show: $\Gamma; \Delta \vdash \langle (\overline{M}, \overline{N}) \rangle : x : T_1 * T_2$. We first note that $\langle (\overline{M}, \overline{N}) \rangle$ is equal to $(\langle \overline{M} \rangle, \langle \overline{N} \rangle)$. By (REWRITE), (DERIVE), and Lemma B.5 we know that $\Gamma; \Delta \hookrightarrow \Gamma; !F_1, !F_2$ and $\Gamma; !F_1 \vdash F_1$ and $\Gamma; !F_2 \vdash F_2$ by (IDENT), (!-LEFT), and (DERIVE).

By applying the induction hypothesis twice we know that $\Gamma; !F_1 \vdash \langle \overline{M} \rangle : T_1$ and $\Gamma; !F_2 \vdash \langle \overline{N} \rangle : T_2\{\langle \overline{M} \rangle / x\}$.
The result follows by an application of (VAL PAIR).

Case (VAL INL ALG): $\Gamma \vdash_{\mathsf{alg}} (\mathsf{inl}\ \overline{M})_{\_+T_2} : T_1 + T_2; !F'$ is proved by $\Gamma \vdash_{\mathsf{alg}} \overline{M} : T_1; F_1$ and $\Gamma \vdash_{\mathsf{alg}} T_2$. We also know that $\Gamma; \Delta \vdash !F'$.

To show: $\Gamma; \Delta \vdash \langle (\mathsf{inl}\ \overline{M})_{\_+T_2} \rangle : T_1 + T_2$. We first note that $\langle (\mathsf{inl}\ \overline{M})_{\_+T_2} \rangle$ is equal to $\mathsf{inl}\ \langle \overline{M} \rangle$. By (REWRITE), (DERIVE), and Lemma B.5 we know that $\Gamma; \Delta \hookrightarrow \Gamma; !F'$ and $\Gamma; !F' \vdash F'$ by (IDENT), (!-LEFT), and (DERIVE).
By applying the induction hypothesis we know that $\Gamma; !F' \vdash \langle \overline{M} \rangle : T_1$.
By Lemma C.4 we know that $\Gamma; \emptyset \vdash T_2$ and thus by Lemma B.7 $\Gamma; !F' \vdash T_2$.
The result follows by an application of (VAL INL).

Case (VAL INR ALG): The proof is analogous to the one for (VAL INL ALG).

Case (VAL FOLD ALG): $\Gamma \vdash_{\mathsf{alg}} \mathsf{fold}\ \overline{M} : \mu\alpha.T'; !F'$ is proved by $\Gamma \vdash_{\mathsf{alg}} \overline{M} : T'\{\mu a.T'/\alpha\}; F'$. We also know that $\Gamma; \Delta \vdash !F'$.

To show: $\Gamma; \Delta \vdash \langle \mathsf{fold}\ \overline{M} \rangle : \mu\alpha.T'$. We first note that $\langle \mathsf{fold}\ \overline{M} \rangle$ is equal to $\mathsf{fold}\ \langle \overline{M} \rangle$. By (REWRITE), (DERIVE), and Lemma B.5 we know that $\Gamma; \Delta \hookrightarrow \Gamma; !F'$ and $\Gamma; !F' \vdash F'$ by (IDENT), (!-LEFT), and (DERIVE).
By applying the induction hypothesis we know that $\Gamma; !F' \vdash \langle \overline{M} \rangle : T'\{\mu a.T'/\alpha\}$.
The result follows by an application of (VAL FOLD).

Case (VAL REF ALG): $\Gamma \vdash_{\mathsf{alg}} \overline{M}_{\{x:\_ \mid F\}} : \{x : T' \mid F\}; F' \otimes F\{\langle \overline{M} \rangle / x\}$ is proved by $\Gamma \vdash_{\mathsf{alg}} \overline{M} : T'; F'$ and $fnfv(F) \subseteq dom(\Gamma) \cup \{x\}$. We also know that $\Gamma; \Delta \vdash F' \otimes F\{\langle \overline{M} \rangle / x\}$.

To show: $\Gamma; \Delta \vdash \langle \overline{M}_{\{x:\_ \mid F\}} \rangle : \{x : T' \mid F\}$. We first note that $\langle \overline{M}_{\{x:\_ \mid F\}} \rangle$ is equal to $\langle \overline{M} \rangle$. By (REWRITE), (DERIVE), and Lemma B.5 we know that $\Gamma; \Delta \hookrightarrow \Gamma; F', F\{\langle \overline{M} \rangle / x\}$ and $\Gamma; F' \vdash F'$ and $\Gamma; F\{\langle \overline{M} \rangle / x\} \vdash F\{\langle \overline{M} \rangle / x\}$ by (IDENT). By induction hypothesis, $\Gamma; F' \vdash \langle \overline{M} \rangle : T$. The result follows from (VAL REFINE).

Case (EXP APPL ALG): The proof follows straightforwardly from (REWRITE), (DERIVE), Lemma B.5, (IDENT) by applying the induction hypothesis twice.

Case (EXP LET ALG): $\Gamma \vdash_{\mathsf{alg}} \mathsf{let}\ x = \overline{E_1}\ \mathsf{in}\ \overline{E_2} : T; \Delta' \multimap (F_1 \otimes (\forall x.forms(x : U) \multimap F_2))$ is proved by $\overline{E_1} \rightsquigarrow^\emptyset [\Delta' \mid \overline{E_1'}]$, $\Gamma \vdash_{\mathsf{alg}} \overline{E_1'} : U; F_1$, $\Gamma, x : \psi(U) \vdash_{\mathsf{alg}} \overline{E_2} : T; F_2$, $x \notin fv(T)$. We also know that $\Gamma; \Delta \vdash \Delta' \multimap (F_1 \otimes (\forall x.forms(x : U) \multimap F_2))$.

To show: $\Gamma; \Delta \vdash \langle \mathsf{let}\ x = \overline{E_1}\ \mathsf{in}\ \overline{E_2} \rangle : T$. We first note that $\langle \mathsf{let}\ x = \overline{E_1}\ \mathsf{in}\ \overline{E_2} \rangle$ is equal to $\mathsf{let}\ x = \langle \overline{E_1} \rangle\ \mathsf{in}\ \langle \overline{E_2} \rangle$. By Lemma C.2 and Lemma B.2, $\Gamma; \Delta, \Delta' \vdash F_1 \otimes (\forall x.forms(x : U) \multimap F_2)$.
By (REWRITE), (DERIVE), and Lemma B.5 it holds that $\Gamma; \Delta, \Delta' \hookrightarrow \Gamma; F_1, (\forall x.forms(x : U) \multimap F_2)$ and $\Gamma; F_1 \vdash F_1$ and $\Gamma; \forall x.forms(x : U) \multimap F_2 \vdash \forall x.forms(x : U) \multimap F_2$ by (IDENT) and (DERIVE).
Without loss of generality, let us assume $x \notin dom(\Gamma)$ and, thus, $x \notin fnfv(\forall x.forms(x : U) \multimap F_2)$. (This assumption can be fulfilled by α-renaming x if necessary.)
By Lemma B.5, we can easily see that $\Gamma, x : \psi(U); \forall x.forms(x : U) \multimap F_2 \vdash \diamond$. By Lemma C.3, $\Gamma, x : \psi(U); \forall x.forms(x : U) \multimap F_2 \vdash forms(x : U) \multimap F_2$. By Lemma C.2, $\Gamma, x : \psi(U); \forall x.forms(x : U) \multimap F_2, forms(x : U) \vdash F_2$.
We note that by statement (1) of Proposition C.13 it holds that $\langle \overline{E_1} \rangle \rightsquigarrow^\emptyset [\Delta' \mid \langle \overline{E_1'} \rangle]$.
By induction hypothesis, $\Gamma; F_1 \vdash \langle \overline{E_1'} \rangle : T$ and $\Gamma, x : \psi(U); \forall x.forms(x : U) \multimap F_2, forms(x : U) \vdash \langle \overline{E_2} \rangle : U$.
The result follows from (EXP LET).

Case (EXP SPLIT ALG): The proof follows a similar strategy as the one for (EXP LET ALG).

Case (EXP MATCH ALG): The proof follows a similar strategy as the one for (EXP LET ALG).

Case (EXP EQ ALG): The proof follows straightforwardly from (REWRITE), (DERIVE), (IDENT), Lemma B.5, the induction hypothesis, ($\otimes$-RIGHT), and Lemma B.9.

Case (EXP ASSUME ALG): $\Gamma \vdash_{\mathsf{alg}}$ (assume $F_1)_T : T; F_1 \multimap F_2$ is proved by $\Gamma \vdash_{\mathsf{alg}}$ (assume $\mathbf{1})_{\_<:T} : T; F_2$ and $fnfv(F) \subseteq dom(\Gamma)$, where $F_1 \neq \mathbf{1}$. We also know that $\Gamma; \Delta \vdash F_1 \multimap F_2$.

> *To show:* $\Gamma; \Delta \vdash \langle(\text{assume } F_1)_T \rangle : T$. We first note that $\langle(\text{assume } F_1)_T\rangle$ is equal to assume F_1.
> By Lemma C.2 we know that $\Gamma; \Delta, F_1 \vdash F_2$.
> By applying the induction hypothesis we know that $\Gamma; \Delta, F_1 \vdash$ assume $\mathbf{1} : T$.
> The result follows by an application of (EXP ASSUME) .

Case (EXP RES ALG): The proof follows a similar and slightly simplified strategy as the one for (EXP LET ALG).

Case (EXP SEND ALG): The proof follows straightforwardly from the induction hypothesis using the fact that $\langle a! \overline{M} \rangle$ is equal to $a! \langle \overline{M} \rangle$.

Case (EXP FORK ALG): The proof follows a similar strategy as the one for (EXP LET ALG).

(2) The proof proceeds by induction on the length of $\Gamma; \Delta \vdash E : T$. The base cases are (VAL VAR), (VAL UNIT), (EXP TRUE), (EXP RECV), and (EXP ASSERT): in these cases we choose $\overline{E} := E$ and $F := \mathbf{1}$. The statement follows by an inspection of the typing rules and by Lemma C.4.

We show the induction cases in the following. Most cases follow a very similar structure so we show detailed examples for standard proof strategies and omit the details for analogous cases.

For all cases the proof is split into two parts: we first show that there exists an annotated term $\overline{E}$ and a formula F such that $\Gamma \vdash_{\mathsf{alg}} \overline{E} : T; F$ and $\langle \overline{E} \rangle = E$. We then prove that $\Gamma; \Delta \vdash F$.

Case (VAL FUN): $\Gamma; \Delta \vdash \lambda x. D : x : T_1 \to T_2$ is proved by $\Gamma, x : \psi(T_1); !\Delta', forms(x : T_1) \vdash D : T_2$ and $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta'$.

By induction hypothesis, there exist $\overline{D}, F'$ such that $\langle \overline{D} \rangle = D$, $\Gamma, x : \psi(T_1) \vdash \overline{D} : T_2; F'$ and $\Gamma, x : \psi(T_1); !\Delta', forms(x : T_1) \vdash F'$.

> *To show:* $\Gamma \vdash_{\mathsf{alg}} \lambda x : T_1. \overline{D} : x : T_1 \to T_2; !\forall x.(forms(x : T_1) \multimap F')$. By Lemma B.5, $fnfv(T_1) \subseteq dom(\Gamma) \cup \{x\}$. The result follows from (VAL FUN ALG). We note that $\langle \lambda x : T_1. \overline{D} \rangle = \lambda x. D$.
> *To show:* $\Gamma; \Delta \vdash !\forall x.(forms(x : T_1) \multimap F')$. By ($\multimap$-RIGHT), $\Gamma, x : \psi(T_1); !\Delta' \vdash forms(x : T_1) \multimap F'$.
> By Lemma B.5, $\Gamma, x : \psi(T_1) \vdash \diamond$ and $x \notin fnfv(!\Delta') \subseteq dom(\Gamma)$. By Lemma C.3, $\Gamma; !\Delta' \vdash \forall x.(forms(x : T_1) \multimap F')$. By (!-RIGHT), $\Gamma; !\Delta' \vdash !\forall x.(forms(x : T_1) \multimap F')$.
> The result follows from Lemma B.9.

Case (VAL PAIR): $\Gamma; \Delta \vdash (M, N) : x : T_1 * T_2$ is proved by $\Gamma; !\Delta_1 \vdash M : T_1$ $\Gamma; !\Delta_2 \vdash N : T_2\{M/x\}$ and $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta_1, !\Delta_2$.

By applying the induction hypothesis twice we know that there exist $\overline{M}, \overline{N}, F_1, F_2$ such that $\langle \overline{M} \rangle = M$ and $\langle \overline{N} \rangle = N$ and $\Gamma \vdash \overline{M} : T_1; F_1$ and $\Gamma \vdash \overline{N} : T_2\{M/x\}; F_2$ and $\Gamma; !\Delta_1 \vdash F_1$ and $\Gamma; !\Delta_2 \vdash F_2$.

> *To show:* $\Gamma \vdash_{\mathsf{alg}} (\overline{M}, \overline{N}) : x : T_1 * T_2; !F_1 \otimes !F_2$. The result follows immediately from (VAL PAIR ALG). We note that $\langle (\overline{M}, \overline{N}) \rangle = (M, N)$.
> *To show:* $\Gamma; \Delta \vdash !F_1 \otimes !F_2$. We apply (!-RIGHT) to derive that $\Gamma; !\Delta_1 \vdash !F_1$ and $\Gamma; !\Delta_2 \vdash !F_2$. By ($\otimes$-RIGHT), $\Gamma; !\Delta_1, !\Delta_2 \vdash !F_1 \otimes !F_2$.
> The result follows from Lemma B.9.

Case (VAL INL): $\Gamma; \Delta \vdash$ inl $M : T_1 + T_2$ is proved by $\Gamma; !\Delta' \vdash M : T_1$ and $\Gamma; !\Delta' \vdash T_2$ and $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta'$.

By applying the induction hypothesis we know that there exist $\overline{M}, F'$ such that $\langle \overline{M} \rangle = M$ and $\Gamma \vdash \overline{M} : T_1; F'$ and $\Gamma; !\Delta' \vdash F'$.

> *To show:* $\Gamma \vdash_{\mathsf{alg}} (\text{inl } \overline{M})_{\_+T_2} : T_1 + T_2; !F'$. We know that $\Gamma; \emptyset \vdash_{\mathsf{alg}} T_2$ by Lemma B.5 and thus $\Gamma \vdash_{\mathsf{alg}} T_2$ by Lemma C.4. The result follows immediately from (VAL INL ALG). We note that $\langle (\text{inl } \overline{M})_{\_+T_2} \rangle = \text{inl } M$.
> *To show:* $\Gamma; \Delta \vdash !F'$. By (!-RIGHT) we know that $\Gamma; !\Delta' \vdash !F'$. The result follows from Lemma B.9.

Case (VAL INR): The proof is analogous to the case of (VAL INL).

Case (VAL FOLD): $\Gamma; \Delta \vdash$ fold $M : \mu\alpha. T'$ is proved by $\Gamma; !\Delta' \vdash M : T'\{\mu\alpha. T'/\alpha\}$ and $\Gamma; \Delta \hookrightarrow \Gamma; !\Delta'$.

By applying the induction hypothesis we know that there exist $\overline{M}, F'$ such that $\langle \overline{M} \rangle = M$ and $\Gamma \vdash \overline{M} : T'\{\mu\alpha. T'/\alpha\}; F'$ and $\Gamma; !\Delta' \vdash F'$.

> *To show:* $\Gamma \vdash_{\mathsf{alg}}$ fold $\overline{M} : \mu\alpha. T'; !F'$. We know that $\Gamma; \emptyset \vdash_{\mathsf{alg}} T_2$ by Lemma B.5 and thus $\Gamma \vdash_{\mathsf{alg}} T_2$ by Lemma C.4. The result follows immediately from (VAL INL ALG). We note that $\langle \text{inl } \overline{M} \rangle = \text{inl } M$.
> *To show:* $\Gamma; \Delta \vdash !F'$. By (!-RIGHT) we know that $\Gamma; !\Delta' \vdash !F'$. The result follows from Lemma B.9.

Case (VAL REFINE): $\Gamma; \Delta \vdash M : \{x : T' \mid F'\}$ is proved by $\Gamma; \Delta_1 \vdash M : T'$, $\Gamma; \Delta_2 \vdash F'\{M/x\}$, and $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_1, \Delta_2$.

By induction hypothesis, there exist $\overline{M}, F''$ such that $\langle \overline{M} \rangle = M$, $\Gamma \vdash_{\mathsf{alg}} \overline{M} : T'$, and $\Gamma; \Delta_1 \vdash F''$.

> *To show:* $\Gamma \vdash_{\mathsf{alg}} \overline{M}_{\{x:\_ \mid F'\}} : \{x : T \mid F'\}; F'' \otimes F'\{M/x\}$. By Lemma B.5, $fnfv(F') \subseteq dom(\Gamma) \cup \{x\}$. The result follows from (VAL REF ALG). We note that $\langle \overline{M}_{\{x:\_ \mid F'\}} \rangle = M$.
> *To show:* $\Gamma; \Delta \vdash F'' \otimes F'\{M/x\}$. The result follows from ($\otimes$-RIGHT) and Lemma B.9.

Case (EXP SUBSUM): $\Gamma; \Delta \vdash E : T$ is proved by $\Gamma; \Delta_1 \vdash E : T'$ and $\Gamma; \Delta_2 \vdash T' <: T$ and $\Gamma; \Delta \hookrightarrow \Gamma; \Delta_1, \Delta_2$.

By applying the induction hypothesis we know that there exist $\overline{E}, F'$ such that $\langle \overline{E} \rangle = E$ and $\Gamma \vdash \overline{E} : T'; F'$ and $\Gamma; \Delta_1 \vdash F'$.

Furthermore, by Lemma C.11 we know that there exist $\overline{T'}, \overline{T}, F''$ such that that $T' = \langle \overline{T'} \rangle$, $T = \langle \overline{T} \rangle$, and $\Gamma \vdash_{\mathsf{alg}} \overline{T'} <: \overline{T}; F''$ and $\Gamma; \Delta_2 \vdash F''$. By Lemma C.12 it follows that there exists $\overline{T}^*$ such that that $T = \langle \overline{T}^* \rangle$ and $\Gamma \vdash_{\mathsf{alg}} T' <: \overline{T}^*; F''$

> *To show:* $\Gamma \vdash_{\mathsf{alg}} \overline{E}_{\_<:\overline{T}^*} : T; F_1 \otimes F_2$. The result follows immediately from (EXP SUBSUM ALG). We note that $\langle \overline{E} \rangle = E$ as stated above.
> *To show:* $\Gamma; \Delta \vdash !F'$. By (!-RIGHT) we know that $\Gamma; !\Delta' \vdash !F'$. The result follows from Lemma B.9.

Case (EXP APPL): The proof follows straightforwardly from the induction hypothesis using ($\otimes$-RIGHT) and Lemma B.9.

Case (EXP LET): $\Gamma; \Delta \vdash$ let $x = E_1$ in $E_2 : T$ is proved by $E_1 \rightsquigarrow^\emptyset [\Delta' \mid E_1']$, $\Gamma; \Delta_1 \vdash E_1' : U$, $\Gamma, x : \psi(U); \Delta_2, forms(x : U) \vdash E_2 : T$, $x \notin fv(T)$, and $\Gamma; \Delta, \Delta' \hookrightarrow \Gamma; \Delta_1, \Delta_2$.

By induction hypothesis, there exists $\overline{E_1'}, F_1$ such that $\langle \overline{E_1'} \rangle = E_1'$, $\Gamma \vdash_{\mathsf{alg}} \overline{E_1'} : U; F_1$, and $\Gamma; \Delta_1 \vdash F_1$. By induction hypothesis, there exists $\overline{E_2}, F_2$ such that $\langle \overline{E_2} \rangle = E_2$, $\Gamma, x : \psi(U) \vdash \overline{E_2} : T; F_2$, and $\Gamma, x : \psi(U); \Delta_2, forms(x : U) \vdash F_2$.

We note that by statement (2) of Proposition C.13 it holds that there exists $\overline{E_1}$ such that $\overline{E_1} \rightsquigarrow^\emptyset [\Delta' \mid \overline{E_1'}]$.

To show: $\Gamma \vdash_{\mathsf{alg}}$ let $x = \overline{E_1}$ in $\overline{E_2} : T; \Delta' \multimap (F_1 \otimes \forall x.(forms(x : U) \multimap F_2))$. By Lemma B.5, $fnfv(\Delta_1) \subseteq dom(\Gamma)$. The result follows from (EXP LET ALG). We note that $\langle$let $x = \overline{E_1}$ in $\overline{E_2}\rangle$ is equal to let $x = E_1$ in E_2.

To show: $\Gamma; \Delta \vdash \Delta' \multimap (F_1 \otimes \forall x.(forms(x : U) \multimap F_2))$. By ($\multimap$-RIGHT), $\Gamma, x : \psi(U); \Delta_2 \vdash forms(x : U) \multimap F_2$. By Lemma B.5, $x \notin fnfv(\Delta_2) \subseteq dom(\Gamma)$. By Lemma C.3, $\Gamma, x : \psi(U); \Delta_2 \vdash \forall x.(forms(x : U) \multimap F_2)$. By ($\otimes$-RIGHT), $\Gamma; \Delta_1, \Delta_2 \vdash F_1 \otimes \forall x.(forms(x : U) \multimap F_2)$. By Lemma B.9, $\Gamma; \Delta, \Delta' \vdash F_1 \otimes \forall x.(forms(x : U) \multimap F_2)$. By ($\multimap$-RIGHT), $\Gamma; \Delta \vdash \Delta' \multimap (F_1 \otimes \forall x.(forms(x : U) \multimap F_2))$.

Case (EXP SPLIT): The proof follows a similar strategy as the one for (EXP LET).

Case (EXP MATCH): The proof follows a similar strategy as the one for (EXP LET).

Case (EXP EQ): The proof follows straightforwardly from applying the induction hypothesis twice and using ($\otimes$-RIGHT) and Lemma B.9.

Case (EXP ASSUME): $\Gamma; \Delta \vdash$ assume $F' : T$ is proved by $\Gamma; \Delta, F' \vdash$ assume $\mathbf{1} : T$, where $F' \neq \mathbf{1}$.

We first note that by Lemma C.14 it holds that $\Gamma; \emptyset \vdash$ assume $\mathbf{1} :$ unit and $\Gamma; \Delta, F' \vdash$ unit $<: T$.

By combining Lemma C.11 and Lemma C.12 we know that there exist $\overline{T}, F''$ such that that $T = \langle \overline{T} \rangle$, and $\Gamma \vdash_{\mathsf{alg}}$ unit $<: \overline{T}; F''$ and $\Gamma; \Delta, F' \vdash F''$. By inspection of the algorithmic subtyping rules it follows that $\overline{T}$ must not contain any annotations $(T = \overline{T})$ and thus $\Gamma \vdash_{\mathsf{alg}}$ unit $<: T; F''$.

By applying the induction hypothesis (see proof of base case (EXP TRUE)) to $\Gamma; \emptyset \vdash$ assume $\mathbf{1} :$ unit it follows that $\Gamma \vdash_{\mathsf{alg}}$ assume $\mathbf{1} :$ unit; $\mathbf{1}$ and $\Gamma; \emptyset \vdash \mathbf{1}$.

To show: $\Gamma \vdash_{\mathsf{alg}}$ (assume $F')_T : T; F' \multimap (\mathbf{1} \otimes F'')$. We first apply (EXP SUBSUM ALG) to derive that $\Gamma \vdash_{\mathsf{alg}}$ (assume $\mathbf{1})_{\_<:T} : T; \mathbf{1} \otimes F''$.

We know that $\Gamma; \emptyset \vdash_{\mathsf{alg}} T_2$ by Lemma B.5 and thus $\Gamma \vdash_{\mathsf{alg}} T_2$ by Lemma C.4. The result follows from (EXP ASSUME ALG). We note that $\langle$(assume $F')_T\rangle =$ assume F'.

To show: $\Gamma; \Delta \vdash F' \multimap (\mathbf{1} \otimes F'')$. As stated above we know that $\Gamma; \emptyset \vdash \mathbf{1}$ and $\Gamma; \Delta, F' \vdash F''$. By ($\otimes$-RIGHT) it holds that $\Gamma; \Delta, F' \vdash \mathbf{1} \otimes F''$

Case (EXP RES): The proof follows a similar strategy as the one for (EXP LET).

Case (EXP SEND): The proof follows straightforwardly from the induction hypothesis.

Case (EXP FORK): The proof follows a similar strategy as the one for (EXP LET).

$\square$