

Enforcing Stateful Authorization and Information Flow Policies in FINE

Nikhil Swamy Juan Chen Ravi Chugh
Microsoft Research, Redmond and University of California, San Diego
{nswamy, juanchen}@microsoft.com, rchugh@cs.ucsd.edu

January 20, 2010

Abstract

Proving software free of security bugs is hard. Languages that ensure that programs correctly enforce their security policies would help, but, to date, no security-typed language has the ability to verify the enforcement of the kinds of policies used in practice—dynamic, stateful policies which address a range of concerns including forms of access control and information flow tracking.

This paper presents FINE, a new source-level security-typed language that, through the use of a simple module system and dependent, refinement, and affine types, checks the enforcement of dynamic security policies applied to real software. FINE is proven sound. A prototype implementation of the compiler and several example programs are available from <http://research.microsoft.com/fine>.

1 Introduction

The security of a well-designed software system often revolves around the concept of a reference monitor, a security-critical kernel that mediates access to resources while enforcing a suitable policy. Reference monitors are expected to be compact and implemented in a form amenable to review. However, increasingly, reference monitors are tasked with enforcing complex policies that simultaneously address various aspects of security, mixing, for example, role- and history-based access control with information flow tracking. Policies are authored separately from the programs they govern, they are composed in non-trivial ways, and, as policies change over time, authorization decisions require reasoning about state. This makes it difficult to establish that a reference monitor enforces a policy correctly.

To illustrate the kinds of security concerns that arise in practice, consider the policy used by CONTINUE [14], a widely used program for managing academic conferences. CONTINUE’s security policy is defined using Datalog-like rules in XACML. This policy stands separately from the implementation of the server program, making it hard to connect the policy to the program objects it governs. The policy is also particularly complex in that it makes extensive use of stateful features. For example, the conference management process is staged into a number of phases—in each phase, different policy rules apply. During the submission phase of a conference, authors may submit papers, but this right is revoked after the submission deadline is passed. In the bidding phase, papers are assigned to reviewers after accounting for conflicts of interest. During the rebuttal phase, reviews are disclosed to authors, but care must be taken to ensure that PC-confidential remarks and scores are not revealed. With such a complex policy to enforce, it is not surprising that the developers of CONTINUE report that almost all the interesting bugs they encountered were related to authorization in some form [7]. Policies used with other kinds of software, such as systems that manage medical records, applications that control the outsourcing of software development, and military systems, arguably have even more complex authorization requirements. Formally verifying that the reference monitors of such systems correctly enforce their policies would help alleviate concerns of security vulnerabilities.

This paper presents FINE, a new source-level security-typed programming language that can be used to implement programs like reference monitors and to check that these programs correctly enforce their security policies. FINE distinguishes itself from prior languages in this line, including FlowCaml [18], Jif [5], Fable [21], Aura [13], and

RCF [1], primarily in its ability to express a combination of stateful authorization (none of the prior languages model state) and information flow (which is the focus of FlowCaml and Jif, and can be encoded in Fable and Aura, but not, as far as we are aware, in RCF). The technical contribution of FINE is a new type system (§3) that uses dependent and refinement types to express authorization policies by including first-order logical formulas in the types of program expressions. FINE uses affine types, a weakening of linear types [24], to model changes to the state of an authorization policy. (Variables with an affine type can be used at most once.) The combination of affine and dependent types is subtle and can require tracking uses of affine assumptions in both types and terms. Our formulation keeps the metatheory simple by ensuring that affine variables never appear in types, while still allowing the state of a program to be refined by logical formulas. We also formalize a module system for FINE that provides a simple but strong information-hiding property—we exploit this property to model information flow.

Programming with these advanced typing constructs can impose a significant burden on the programmer. For this reason, languages like Fable and Aura position themselves as intermediate languages because verification depends on intricate security proofs too cumbersome for programmers to write down. Indeed, checking the 2000 lines of code in our benchmark programs produces nearly 200 proof obligations, a proof burden that would overwhelm most programmers. To alleviate this concern, FINE draws on the experience of languages like F7 (an implementation of RCF) and uses Z3 [6], an SMT solver, to automatically discharge proof obligations. The careful combination of refinement and affine types in FINE allows us to use a mature classical prover like Z3. Refinement formulas in FINE only involve the standard logical connectives, avoiding the need for still-experimental linear-logic provers.

We describe our experience using FINE to build several example programs (§4), including a model of the reference monitor of CONTINUE. The complete semantics of FINE, proofs of theorems, and additional examples appear in a technical report [20].

2 FINE, by example

We begin by presenting FINE using several examples. Our first example is a simple form of password-based authentication. Next, we discuss permission-based access control enriched with information flow tracking. Finally, we show how to enforce stateful authorization policies by presenting code examples from our main case-study, a model of the CONTINUE conference management server.

2.1 Authentication, access control, and information flow

FINE’s syntax is similar to languages in the ML family. In order to specify and enforce security policies, FINE programmers define modules that provide mediated access to security-sensitive resources. The module Authentication shown below mediates access to authentication routines.

Simple password authentication

```

1 module Authentication
2 type prin = U: string → prin | Admin: prin
3 private type cred :: prin → * = Auth: p:prin → cred p
4 val login: p:prin → string → option (cred p)
5 let login p pw = if (check_pwd_db p pw) then Some (Auth p) else None

```

The type `prin` is a standard variant type that represents principal names as either a `string` for the user’s name, or the distinguished constant `Admin`. The type `cred` (line 3) is a dependent-type constructor with *kind* `prin → *`, e.g., `(cred Admin)` is a legal type of kind `*` (the kind of normal types, distinguished from the kind of affine types, introduced in §2.2) and represents a credential for the `Admin` user. Values of the `cred p` type are constructed using the `Auth` data constructor. This constructor is given a dependent function type—the argument `p` is the name of the principal and is in scope to the right of the function arrow. By declaring `cred private`, the `Authentication` module indicates that its clients cannot directly use the `Auth` constructor. Instead, the only way a client module can obtain a credential is by calling the `login` function (given a dependent function type on line 4). The implementation of `login` (line 5) calls an external function (not shown) to check the password, and, if the password check succeeds, returns a credential for the user `p`. By indexing `cred` with the name of the principal which it authenticates, we can statically detect common security errors. For example, a client

cannot use `login` to obtain a credential for U “Alice” and later pass it off as a credential for Admin—the type of the former, `cred (U “Alice”)`, distinguishes it from the latter, which has type `cred Admin`.

We use `Authentication` to implement the `FileRM` module (shown on the next page), a reference monitor that mediates access to a file system. The policies implemented by reference monitors in FINE have two components: the types given to values exposed in the module’s interface (e.g., the type of `fread` on line 7), and policy axioms introduced by the `assume` construct (e.g., `assume AdminRW` on line 6). A security review of a FINE module must confirm that the types and assumptions adequately capture the intent of a high-level policy. Importantly, client code need not be reviewed—typing ensures that clients comply with the reference monitor’s security policy.

The `FileRM` module aims to provide a basic level of access protection on files by ensuring that principals that read and write to files have the requisite permissions. This basic protection is implemented by lines 1-7 of `FileRM`. The remainder of the module enriches the access control mechanism to track information flows so that, for example, users cannot reveal secrets by copying data from a secret file into a public file.

Permission-based access control and information flow on files

```

1 module FileRM
2 open Authentication (* Use non -private symbols from Authentication's namespace *)
3 (* Propositions and assumptions for file permissions *)
4 type CanRead:: prin → Sys.file → *
5 type CanWrite:: prin → Sys.file → *
6 assume AdminRW: forall f:Sys.file. CanRead Admin f && CanWrite Admin f
7 val fread_simple: p:prin → cred p → {f:Sys.file | CanRead p f} → string
8 (* Types and operators to track information flow *)
9 type label = F : Sys.file → label | J : label → label → label
10 private type tracked :: * → label → * = L : α → p:label → tracked α p
11 val fmap: (α → β) → l:label → tracked α l → tracked β l
12 val tensor: l:label → m:label → tracked (α → β) l → tracked α m → tracked β (J l m)
13 (* Types and axioms for a partial order on labels *)
14 type CanFlow:: label → label → *
15 assume Lattice: forall l:label, m1:label, m2:label. (CanFlow l l) &&
16   ((CanFlow l m1 && CanFlow l m2) ⇒ CanFlow l (J m1 m2)) &&
17   ((CanFlow m1 l && CanFlow m2 l) ⇒ CanFlow (J m1 m2) l)
18 assume Atomicflow: forall f:Sys.file, g:Sys.file.
19   (forall p:prin. CanRead p g ⇒ CanRead p f) ⇒ CanFlow (F f) (F g)
20 (* Secure wrappers for system calls *)
21 val fread: p:prin → cred p → f:{x:Sys.file | CanRead p x} → tracked string (F f)
22 let fread p c f = L (Sys.fread f) (F f)
23 val fwrite: p:prin → cred p → f:{x:Sys.file | CanWrite p x} →
24   l:{y:label | CanFlow y (F f)} → tracked string l → unit
25 let fwrite p c f l (L s x) = Sys.fwrite f s

```

`FileRM` defines dependent-type constructors `CanRead` and `CanWrite` to describe access permissions. Permissions are granted using assumptions like `AdminRW`, which states that the `Admin` user has read- and write-permissions on all files. Client programs can use axioms like `AdminRW` to produce evidence of the propositions required to call functions like `fread_simple`, which wrap the underlying system calls. Client programs are assumed to not have direct access to these system calls—this can be established using standard systems techniques like sandboxing [25]. The type of `fread_simple` is used to enforce an access control policy. A caller of `fread_simple` is required to pass in a credential for a user `p` and a file handle `f`, where `f` has the refined type `{x:Sys.file | CanRead p x}` indicating that `p` has permission to read `f`.

We used `fread_simple` mainly to illustrate how refinement types can express simple authorization policies. When leaks due to information flows are a concern, `FileRM` would not include `fread_simple` in the API exposed to client programs. Clients would have to use `fread` instead, which augments `fread_simple` with information flow controls.

The encoding of information flow shown in `FileRM` is based on a model developed with the Fable calculus [21]. Information flow policies are specified and enforced by tagging sensitive data with *security labels* that record provenance. The type `label` (line 9) represents the provenance of data derived from one or more files, `F x` for data from file `x`, and `J l1 l2` for data derived from the files in both `l1` and `l2`. The dependent-type constructor `tracked` associates labels with data. For example, `tracked string (F x)` represents a string that originated from the file `x`. Importantly, `tracked` is defined as a private type. Client programs can only manipulate `tracked` values using functions that appear in the interface of

FileRM, e.g., `fmap`, a functor that allows functions to be lifted into the tracked type and `tensor`, a combinator that treats the tracked type as an indexed applicative functor. Prior work on Fable showed that encodings of this style can be proved to correctly enforce security properties like noninterference.

Next, we define a type `CanFlow` and assumptions to describe a partial order on labels. The `Lattice` assumption states that the `J` constructor behaves as the least-upper-bound relation on a join semi-lattice and that flows are permissible from lower labels to higher ones. The `Atomicflow` assumption states that data can flow from a file `f` to a file `g` only if all principals that can read `g` can also read `f`. The types of `fread` and `fwrite` use these constructs to track information flow. The type of `fread` shows that the content of `f` is returned as a string tagged with its provenance, i.e., tracked string $(F\ f)$. The type of `fwrite` requires that the string written to a file `f` has provenance `l`, where the refinement `CanFlow y (F f)` on the type of `l` requires it to only contain data visible to the readers of `f`.

Specific file permissions and a client program

```

1 open Authentication, FileRM
2 assume R_a: CanRead (U "Alice") "a.txt" &&
3   (forall p:prin.CanRead p "a.txt" => p=U "Alice" || p=Admin)
4 assume R_ab: CanRead (U "Alice") "ab.txt" && CanRead (U "Bob") "ab.txt" &&
5   (forall p:prin.CanRead p "ab.txt" => p=U "Alice" || p=U "Bob" || p=Admin)
6 val strcat: string -> string -> string
7 let sudo (c:cred Admin) =
8   let a, ab = fread Admin c "a.txt", fread Admin c "ab.txt" in
9   let a_ab = tensor (F "a.txt") (F "ab.txt") (fmap strcat (F "a.txt") a) ab in
10  fwrite Admin c "a.txt" (J (F "a.txt") (F "ab.txt")) a_ab

```

Additional policy assumptions and client code. The code sample above includes axioms `R_a` and `R_ab` to define access permissions for some files. (We assume here that `Sys.file` and `string` are synonyms.) We also show a client program, `sudo`, which runs with the credentials of `Admin`, concatenates data from files `a.txt` and `ab.txt`, and writes the result to the file `a.txt`. In addition to `Admin`, the file `a.txt` is readable only by the user `Alice` and `ab.txt` only by `Alice` and `Bob`. Thus, `sudo` is secure since it writes to `a.txt` data that can be read by `Alice` and `Admin`. In contrast, if `sudo` were to write the result to `ab.txt`, the contents of `a.txt` are leaked to `Bob`, and this program should be detected as insecure.

At each call to `fread`, the solver appeals to `AdminRW` to show that `Admin` has read permission on the files. To concatenate tracked strings, we use the `fmap` and `tensor` operators from the `FileRM` API.¹ Here, the type computed for `a_ab` is tracked string $(J (F "a.txt") (F "ab.txt"))$. At line 10, we need to prove `CanFlow (J (F "a.txt") (F "ab.txt")) (F "a.txt")`, which is discharged automatically by `Z3`. Trying to write `a_ab` to `ab.txt` instead results in a type error.

2.2 Stateful authorization in the CONTINUE conference manager

We now present a more substantial example in `FINE`: a model of the `CONTINUE` conference management server. We first present a reference monitor `ConfRM` which mediates access to a database of paper submissions and reviews. Next, we show `ConfPolicy`, a set of policy axioms used to configure the reference monitor. Finally, we discuss `ConfWeb`, a web server processing requests and accessing the database via the reference monitor.

A model of stateful authorization. The design of the `ConfRM` reference monitor is based on a framework due to Dougherty et al. [7] for reasoning about the correctness of Datalog-style dynamic policies. This model specifies policies as inference rules that derive permissions from basic authorization attributes. For example, attributes may include assertions about a principal’s role membership or the phase of the conference, and inference rules could grant permissions to principals depending on the current phase and role activations. Over time, whether due to a program’s actions or due to external events, the set of authorization attributes can change. For example, to access a resource, a principal may alter the state of the authorization policy by activating a role; or, the PC chair can change the phase of the conference. In this state, the policy may grant a specific privilege to the principal, but a subsequent role deactivation revokes the privilege. Dougherty et al. show that this model captures many common policies and can be used to reason about policy correctness.

¹Our implementation currently lacks support for implicit parameters in function calls. Defining all label parameters to be implicit would produce more terse programs. For example, concatenation of tracked strings would read `tensor (fmap strcat a) ab`.

This model of stateful authorization can be represented directly in FINE. The type st represents the set of basic authorization attributes (line 10 in the listing on the next page). Attributes include values like $Role$ (U “Alice”) $Author$ to represent a role activation, or values like $Assigned\ r\ p$ to indicate that a paper p has been assigned to a reviewer r . The type $perm$ represents permissions (the relations derived using inference rules from the basic authorization attributes). For example, $Permit$ (U “Alice”) ($Submit\ p$) represents a permission granted to an author. $ConfRM$ also defines two propositions for stating invariants about the current state of the policy. Line 12 shows the type In , a proposition about list membership, e.g., $In\ a\ s$ states that a is a member of the list s . We elide standard assumptions that axiomatize list membership, but show a simple recursive function $check$ that decides list membership (line 13-15). The proposition $Derivable\ s\ p$ (line 16) asserts that a permission p is derivable from the collection of authorization attributes s . We define two type abbreviations for refinements of the st type: $rst<p>$ are those states in which p is derivable, and $inst<a>$ are those states that include a .

For a flavor of refinement type checking, consider the $check$ function. The essence of typing this function is proving that the $true$ sub-expression can be given the type $\{b:bool \mid In\ a\ l\}\}$. We accomplish this by typing the value $true$ in a context that records equalities between l and $hd::tl$ (induced by the pattern match); an assumption that the expression $(equals\ a\ hd)$ has the type $\{b:bool \mid b=true \Leftrightarrow a=hd\}$ (by a type given to the built-in $equals$ operator); an assumption that $(equals\ a\ hd)$ evaluates to $true$ (since we are typing the **then**-branch); and the axioms for list membership. We determine if the goal $(In\ a\ l)$ is deducible from the assumptions by including the negation of the goal among the assumptions and requiring the solver to prove the resulting theory unsatisfiable.

Modeling state updates with affine types. The type constructor $StateIs$ (line 19) addresses two concerns. A value of type $StateIs\ s$ represents an assertion that s contains the current state of authorization facts. $ConfRM$ uses this assertion to ensure the integrity of its authorization facts. $StateIs$ is declared **private**, so untrusted clients cannot use the $Sign$ constructor to forge $StateIs$ assertions. Moreover, since the authorization state can change over time, FINE’s type system provides a way to revoke $StateIs$ assertions about stale states. For example, after a reviewer r has submitted a review for a paper p , we may add the fact $Reviewed\ r\ p$ to the set of authorization facts s , revoke the assertion $StateIs\ s$, and use $StateIs\ ((Reviewed\ r\ p)::s)$ instead.

A fragment of a reference monitor for a conference management server

```

1 module ConfRM
2 open Authentication
3 type role = Author | Reviewer | Chair
4 type phase = Submission | Reviewing | Meeting
5 type paper = {id:int; title:string; author:prin; contents:string}
6 type attr = Role : prin → role → attr | Assigned : prin → paper → attr
7           | Phase : phase → attr | Reviewed : prin → paper → attr
8 type action = Submit: paper → action | Review: paper → action
9           | ReadScore: paper → action | CloseSub: action
10 type st = list attr
11 type perm = Permit : prin → action → perm
12 type In :: attr → st → *
13 val check: a:attr → l:st → {b:bool | b=true ⇒ In a l}
14 let rec check a l = match l with [] → false
15                   | hd::tl → if equals a hd then true else check a tl
16 type Derivable :: st → perm → *
17 type rst<p:perm> = {s:st | Derivable s p}
18 type inst<a:attr> = {s:st | In a s}
19 private type StateIs:: st → A = Sign: s:st → StateIs s
20 val submit: q:prin → cred q → p:paper → s:rst<Permit q (Submit p)> → StateIs s → StateIs s
21 val review: r:prin → cred r → p:paper → q:string → s:rst<Permit r (Review p)> →
22       StateIs s → (s’:inst<Reviewed r p> * StateIs s’)
23 val close_sub: c:prin → cred c → s:rst<Permit c CloseSub> →
24       StateIs s → (s’:inst<Phase Reviewing> * StateIs s’)

```

FINE types are classified into two basic kinds: $*$, the kind of normal types, and A , the kind of affine types. By declaring $StateIs\ ::\ st\ \rightarrow\ A$ we indicate that $StateIs$ constructs an affine type from a value of type st . When the state of the authorization policy changes from s to t , $ConfRM$ constructs a value $Sign\ t$ to assert $StateIs\ t$, while destructing a $StateIs\ s$ value to ensure that the assertion about the stale state s can never be used again.

An external API to the conference DB. Lines 20-24 show the types of functions exposed by $ConfRM$ to clients. Using

the refined state type $\text{rst}\langle p \rangle$, the API ensures that each function is only called in states where the permission p is derivable. The `submit` function requires $\text{Permit } q \text{ (Submit } p \text{)}$ to be derivable in the state s . By returning $\text{StateIs } s$, the type of `submit` indicates that it does not change the authorization state. The `review` function allows a reviewer r to submit a review and then changes the authorization state to record the submission. The return type of `review` is a dependent pair consisting of a new list of authorization attributes s' , and an assertion of type $\text{StateIs } s'$ to indicate that s' is the new authorization state. The `close_sub` function has a similar type and allows the program chair to change the phase of the conference.

An example policy and a main event loop for the server

```

1 module ConfPolicy : ConfRM
2 let init:(s:st * StateIs s) = let a = [Role (U "Andy") Chair; ...] in (a, Sign a)
3 assume C1: forall (q:prin), (p:paper), (s:st).
4   In (Phase Submission) s && In (Role q Author) s => Derivable s (Permit q (Submit p))
5 assume C2: forall (r:prin), (p:paper), (s:st).
6   In (Phase Reviewing) s && In (Assigned r p) s => Derivable s (Permit r (Review p))
7 assume ...
8 (* Main event loop *)
9 module ConfWeb
10 open Authentication, ConfRM, ConfPolicy
11 let rec loop s = match get_request() with
12 | Submit_paper q credq paper -> let (a,tok) = s in
13   if (check (Phase Submission) a) and (check (Role q Author) a) then
14     let s1 = submit q credq paper a tok in
15     let _ = resp "Thanks for your submission!" in loop (a, s1)
16   else let _ = resp "Submissions are closed, or you are not an author." in loop (a,tok)
17 | Submit_review r credr paper review -> ...
18 let _ = loop ConfPolicy.init

```

A sample policy. The module `ConfPolicy` above configures the `ConfRM` reference monitor with policy assumptions. At line 2, we show `init`, an initial collection of authorization attributes a , signed to attest that a is the authorization state. The `Sign` data constructor requires the privilege of `ConfRM`—FINE’s module system grants this privilege to `ConfPolicy` using the notation `module ConfPolicy : ConfRM`, which allows `ConfPolicy` to use the private constructors of `ConfRM`. The assumptions C1-C2 show how permissions can be derived from authorization attributes—different conferences can use the same `ConfRM` but get different enforcement semantics by using different policy files.

An event loop to handle web requests. Finally, we show fragments from `ConfWeb`, a program that handles web requests to the conference management site. The main event loop of `ConfWeb` waits for a request (type elided). If principal q wishes to submit a paper, we check that the conference is in the `Submission` phase, and that q is registered in the role of an `Author`. We give the built-in boolean operator `and` the type $x:\text{bool} \rightarrow y:\text{bool} \rightarrow \{z:\text{bool} \mid z=\text{true} \Leftrightarrow x=\text{true} \ \&\& \ y=\text{true}\}$. We can use this type, the type of `check`, and assumption C1, to refine the type of the current state a in the `then`-branch to $\text{rst}\langle \text{Permit } q \text{ (Submit paper)} \rangle$.

2.3 Elements of FINE that enable stateful programming

Before proceeding to a formal semantics for FINE, we discuss a number of elements in the design of FINE that facilitate, and in some cases simplify, stateful programming.

Non-affine state simplifies programming. Programming with affine types can be difficult, since affine variables can never be used more than once. Our approach of using an affine assertion $\text{StateIs } s$ to track the current authorization state minimizes the difficulty. Importantly, the collection of authorization facts s is itself not affine and can be freely used several times, e.g., s is used in several calls to `check`. Non-affine state also enables writing functions like `check`, which, if s was affine, would destroy the state of the program. Only the affine token, $\text{tok}:\text{StateIs } s$, must be used with care, to ensure that it is not duplicated.

Non-affine refinements simplify automated proofs. Even ignoring the inability of prior languages to handle stateful policies, the proof terms required for our examples in languages like `Fable` or `Aura` would be extremely unwieldy. By ensuring that refinement formulas always apply to non-affine values, our proof system is kept tractable, allowing us to use `Z3` to automatically discharge proof obligations. A naïve combination of dependent and affine types would

allow refinements to apply to affine values, necessitating an embedding of linear logic in Z3. Our approach avoids this complication, while retaining the ability to refine the changing state of a program with logical formulas.

Affine types enable flexible mixing of stateful and pure code. Another approach to working with stateful policies could be to use an abstract monad. FINE’s module system certainly supports programming in this style. However, affine types afford greater flexibility. For example, rather than monadically threading a monolithic store through the program, FINE programs can partition the state and pass only the relevant parts of the store to functions that need it. We use this idiom to good effect in one of our benchmark programs (FileAutomaton in §4), in which a bit of state representing the current state of a file is associated with the file handle rather than using a monolithic store to maintain the state of all file handles. Another benchmark, a model of an email client, uses affine types to model capabilities [15] that grant programs restricted access to certain sensitive stateful operations, such as sending emails.

3 Formalizing FINE

Our compiler translates FINE programs in type-preserving manner to .NET bytecode (CIL) [8]. Although we do not report on our type-preservation results in this paper, this design plays a significant role in various aspects of FINE’s type system. This section formalizes FINE, presents a soundness result for the type system, and an information-hiding property for the module system. We begin by presenting a core syntax for FINE.

3.1 Core syntax

Our formulation of FINE’s module system is based on Grossman et al’s [11] syntactic approach to type abstraction. In this formulation, module names correspond to “principals” and are ranged over by the meta-variables p , q , and r . Source expressions are annotated with the names of the modules to which they belong—in the form $\langle e \rangle_p$, the expression e delimited within brackets is privileged to use p ’s private types concretely. A principal constant is denoted \mathbf{p} , and we include two distinguished principals: \top includes the privileges of all other principals, and \perp has no privileges. Values are partitioned into families corresponding to principals. A pre-value for code with p -privilege, u_p , is a variable or a fully-applied data constructor D . Values for p are either its pre-values, abstractions, or pre-values u_q for some other principal q , enclosed within brackets to denote that u_q carries q -privilege. The dynamic semantics of FINE (§3.3) tracks the privilege associated with an expression using these brackets and allows us to prove (§3.4) that programs without p -privilege treat p -values abstractly.

Core syntax of FINE

$p, q, r ::= \mathbf{p} \mid \top \mid \perp$	principals
$u_p ::= x \mid D \bar{\tau} \bar{v}_p$	pre p -values
$v_p ::= u_p \mid \lambda x:\tau.e \mid \Lambda \alpha::\kappa.e \mid \langle u_q \rangle_q$	p -values
$e ::= v_p \mid \text{let } x = e_1 \text{ in } e_2 \mid \text{fix } f:\tau.e \mid v_p v_q \mid v_p \tau \mid \langle e \rangle_p$ $\quad \text{match } v_p \text{ with } D \bar{\tau} \bar{x} \rightarrow e_1 \text{ else } e_2$	terms
$\tau, \phi ::= \alpha \mid x:\tau \rightarrow \tau' \mid \forall \alpha::\kappa.\tau \mid \{x:\tau \mid \phi\} \mid !\tau \mid T \mid \tau \tau' \mid \tau v_p$	types
$\kappa ::= \star \mid \mathbf{A} \mid \star \rightarrow \kappa \mid \mathbf{A} \rightarrow \kappa \mid \tau \rightarrow \kappa$	kinds
$S ::= T::\kappa \mid D:(p, \tau) \mid p \sqsubseteq q \mid S, S' \mid \cdot$	signature
$\Gamma ::= \alpha::\kappa \mid x:(p, \cdot, \tau) \mid x:(p, \text{fix}, \tau) \mid v_p \doteq v'_p \mid \Gamma, \Gamma' \mid \cdot$	type env.

Expressions e are standard for a polymorphic lambda calculus. Types τ include dependent function types $x:\tau \rightarrow \tau'$, where x names the formal parameter and is bound in τ' . Polymorphic types $\forall \alpha::\kappa.\tau$ decorate the abstracted type variable α with its kind κ . Refinement types are written $\{x:\tau \mid \phi\}$, where ϕ is a type in which x is bound. An affine qualifier can be attached to a type using $!\tau$. Type constructors T can be applied to other types using $\tau \tau'$ or terms using τv_p . Note that type-level terms are always values, not expressions—this restriction explains our use of A-normal form [10] for the expression language. This form allows every intermediate result to be named and for these names to appear, potentially, as type indices. Types are partitioned into normal types (kind \star) and affine types (kind \mathbf{A}). Type constructors T construct types of kind κ from normal types ($\star \rightarrow \kappa$), affine types ($\mathbf{A} \rightarrow \kappa$), or τ -typed terms ($\tau \rightarrow \kappa$). Although included in our implementation, for simplicity, our formalization omits dependent pairs.

Desugaring FINE modules. The type and data constructor declarations in a FINE module are desugared to a signature S . The type constructors of the Authentication module of §2.1, for example, are desugared to $\text{prin}::\star$ and $\text{cred}::\text{prin} \rightarrow \star$. Data constructors D are associated their type, as well as the privilege p required for their use. For example, the constructors of the prin type are $\text{U}:(\perp, \text{string} \rightarrow \text{prin})$ and $\text{Admin}:(\perp, \text{prin})$, indicating that these may be used freely in unprivileged code. In contrast, being declared **private**, the constructor of the cred type is desugared to $\text{Auth} : (\text{Authentication}, p:\text{prin} \rightarrow \text{cred } p)$, indicating that it may only be used in code marked with the privilege of the Authentication module. Additionally, signatures use $p \sqsubseteq q$ to record a partial order among principals, with $\perp \sqsubseteq p \sqsubseteq \top$, for all p . We use this to represent sharing between modules, as achieved by the $\text{ConfPolicy} : \text{ConfRM}$ declaration from §2.2. This is translated to the relation $\text{ConfRM} \sqsubseteq \text{ConfPolicy}$, to indicate that ConfPolicy holds the privileges of ConfRM (and, in particular, can use ConfRM 's private data constructors).

Desugaring formulas and assumptions. Refinement formulas and assumptions are represented using type and data constructors, respectively. For example, we use type constructors like $\text{And}::\star \rightarrow \star \rightarrow \star$ to represent the logical connectives. We model equality by specializing it to each type, e.g., $\text{Eq_bool}::\text{bool} \rightarrow \text{bool} \rightarrow \star$. A polymorphic treatment of equality poses no fundamental difficulty, but we use a monomorphic treatment here for simplicity. Quantification is represented using the binders in dependent functions and pairs. For example, the AdminRW assumption from §2.1 is desugared to $\text{AdminRW} : (\perp, f:\text{file} \rightarrow \text{And}(\text{CanRead Admin } f) (\text{CanWrite Admin } f))$. Note that assumptions are always public—we leave an exploration of private assumptions to future work.

Well-formedness conditions on data constructors. The soundness of FINE's type system relies on some restrictions on the use of data constructors D . We mention these restrictions briefly here, and elaborate upon these in the Appendix. First, we disallow partial application of data constructors as this complicates our translation to CIL. Next, we require the type of each data constructor to be of the form: $\forall \bar{\alpha}::\bar{\kappa}. x_1:\tau_1 \rightarrow \dots \rightarrow x_n:\tau_n \rightarrow \tau$, i.e., we require any type arguments to precede any term arguments, although each term argument $x_i:\tau_i$ may itself contain quantifiers. This restriction is merely a convenience—it simplifies the shape of our pattern matching constructs. Finally, for each data constructor D with a type as shown above, we require $\bar{\alpha} \subseteq \text{Free-type-variables}(\tau)$, i.e., every type argument must appear as an index on the constructed type τ . This is a more significant restriction and is necessary for showing that well-typed programs enjoy a type-erasure property.

3.2 Static semantics

The static semantics makes use of a typing environment Γ , which binds type and term variables, and records the results of pattern matching tests using $v_p \doteq v'_p$. Variables x , like data constructors, are associated with a principal p representing the privilege required for their use. Additionally, variable bindings are tagged with a marker $\mu ::= \cdot \mid \text{fix}$ that serves to distinguish fixpoint variables from other variables—the former cannot be used at the type level. We write $x : (p, \tau)$ as a shorthand for $x : (p, \cdot, \tau)$.

Well-formedness of kinds: $S \vdash_i k$, and kinding of types: $S; \Gamma \vdash \tau :: \kappa$

Where, $i ::= \cdot \mid \perp$, and $\star \leq \star$, $\mathbf{A} \leq \mathbf{A}$, $\star \leq \mathbf{A}$

$$\begin{array}{c}
\frac{}{S \vdash \star} \quad \frac{}{S \vdash_i \mathbf{A}} \quad \frac{S \vdash_i \kappa}{S \vdash_i \star \rightarrow \kappa} \quad \frac{S \vdash_{\perp} \kappa}{S \vdash_i \mathbf{A} \rightarrow \kappa} \quad \frac{S; \cdot \vdash \tau :: \star \quad S \vdash_i \kappa}{S \vdash_i \tau \rightarrow \kappa} \\
\frac{}{S; \Gamma \vdash \alpha :: \Gamma(\alpha)} \text{(K-Var)} \quad \frac{}{S; \Gamma \vdash T :: S(T)} \text{(K-TC)} \quad \frac{S; \Gamma \vdash \tau :: \star}{S; \Gamma \vdash \tau :: \mathbf{A}} \text{(K-Afn)} \\
\frac{S; \Gamma, \alpha:\kappa \vdash \tau :: \kappa' \quad \kappa, \kappa' \in \{\star, \mathbf{A}\}}{S; \Gamma \vdash \forall \alpha::\kappa. \tau :: \star} \text{(K-Univ)} \quad \frac{S; \Gamma \vdash \tau_1 :: \kappa \quad \kappa \leq \kappa' \quad S; \Gamma, x:(p, \tau_1) \vdash \tau_2 :: \kappa'}{S; \Gamma \vdash x:\tau_1 \rightarrow \tau_2 :: \star} \text{(K-Arrow)} \quad \frac{S; \Gamma \vdash \tau_1 :: \kappa' \rightarrow \kappa \quad S; \Gamma \vdash \tau_2 :: \kappa'}{S; \Gamma \vdash \tau_1 \tau_2 :: \kappa} \text{(K-App)} \\
\frac{S; \Gamma \vdash \tau_1 :: \tau \rightarrow \kappa \quad S; \Gamma; \cdot \vdash_{\perp} v_p : \tau \quad x \in FV(v_p) \Rightarrow (x, \tau') \in \Gamma}{S; \Gamma \vdash \tau_1 v_p :: \kappa} \text{(K-Dep)} \quad \frac{S; \Gamma \vdash \tau :: \star \quad S; \Gamma, x:(p, \tau) \vdash \phi :: \star}{S; \Gamma \vdash \{x:\tau \mid \phi\} :: \star} \text{(K-Refine)}
\end{array}$$

The first judgment $S \vdash_i \kappa$, shown above, defines a well-formedness relation on kinds. This judgment establishes two properties. First, types constructed from affine types must themselves be affine—this is standard [24]. Without this restriction, an affine value can be stored in a non-affine value and be used more than once. To enforce this property,

we index the judgment using $i ::= \cdot \mid 1$, and when checking a kind $A \rightarrow \kappa$, we require κ to finally produce an A -kinded type. The second restriction, enforced by the first premise ($S; \cdot \vdash \tau :: \star$) of the last rule, ensures that only non-affine values appear in a dependent type. Note that we omit higher kinds (e.g., $(\star \rightarrow \star) \rightarrow \star$) as these are not easily translated to CIL.

The judgment $S; \Gamma \vdash \tau :: \kappa$ states that τ has kind κ . Types inhabited by terms always have kind \star or A . ($K\text{-Affin}$) rules out “doubly-affine” types ($!!\tau$). ($K\text{-Univ}$) allows abstraction only over \star and A -kinded types. ($K\text{-Arrow}$) requires that the type τ_1 of a function’s parameter always have kind \star or A and that functions with affine arguments produce affine results, both captured by an auxiliary relation on kinds, $\kappa \leq \kappa'$. ($K\text{-Dep}$) checks the well-formedness of dependent types. As in Aura and RCF, we restrict type-level terms to values e.g., $\text{Eq_bool}(\text{true} \ \&\& \ \text{false}) \ \text{false}$ is not a well-formed type. This restriction reduces expressiveness by ruling out type-level computations, but greatly simplifies the compilation to CIL. The second premise of ($K\text{-Dep}$) uses the typing judgment—we describe it shortly. ($K\text{-Refine}$) only allows non-affine types τ to be refined by non-affine formulas ϕ .

Expression typing: $S; \Gamma; X \vdash_p e : \tau$

Where, $\mu = \text{fix} \mid \cdot$, $X ::= \cdot \mid x, X$; $Q(X, \tau) = !\tau$, $Q(\cdot, \tau) = \tau$; and $?\tau$ denotes τ or $!\tau$

$$\begin{array}{c}
\frac{S(D) = (p, \tau)}{S; \Gamma; \cdot \vdash_p D : \tau} \text{(T-Data)} \quad \frac{\Gamma(x) = (p, \mu, \tau) \quad S; \Gamma \vdash \tau :: \star}{S; \Gamma; \cdot \vdash_p x : \tau} \text{(T-Var)} \quad \frac{\Gamma(x) = (p, \mu, \tau)}{S; \Gamma; x \vdash_p x : \tau} \text{(T-VarA)} \\
\\
\frac{q \sqsubseteq p \in S \quad S; \Gamma; X \vdash_q e : \tau}{S; \Gamma; X, X' \vdash_p e : \tau} \text{(T-Drop)} \quad \frac{S; \Gamma \vdash \tau :: \star \quad S; \Gamma, f:(p, \text{fix}, \tau); \cdot \vdash_p v_p : \tau}{S; \Gamma; \cdot \vdash_p \text{fix } f:\tau.v_p : \tau} \text{(T-Fix)} \\
\\
\frac{S; \Gamma \vdash \tau_1 :: \kappa \quad \kappa \in \{\star, A\} \quad S; \Gamma, x:(p, \tau_1); X, x \vdash_p e : \tau_2}{S; \Gamma; X \vdash_p \lambda x:\tau_1. e : Q(X, x:\tau_1 \rightarrow \tau_2)} \text{(T-Abs)} \quad \frac{\kappa \in \{\star, A\} \quad S; \Gamma, \alpha::\kappa; X \vdash_p e : \tau'}{S; \Gamma; X \vdash_p \Lambda \alpha::\kappa. e : Q(X, \forall \alpha::\kappa. \tau')} \text{(T-TAbs)} \\
\\
\frac{S; \Gamma; X \vdash_p e_1 : \tau_1 \quad S; \Gamma \vdash \tau_2 :: \kappa \quad S; \Gamma, x:(p, \tau_1); X', x \vdash_p e_2 : \tau_2}{S; \Gamma; X, X' \vdash_p \text{let } x = e_1 \text{ in } e_2 : \tau_2} \text{(T-Let)} \quad \frac{S; \Gamma; X \vdash_p v_p : ?x:\tau_1 \rightarrow \tau_2 \quad S; \Gamma; X' \vdash_p v'_p : \tau_1 \quad S; \Gamma \vdash \tau_2[v'_p/x] :: \kappa}{S; \Gamma; X, X' \vdash_p v_p v'_p : \tau_2[v'_p/x]} \text{(T-App)} \\
\\
\frac{S; \Gamma; X \vdash_p v_p : ?\forall \alpha::\kappa. \tau \quad S; \Gamma \vdash \tau' :: \kappa}{S; \Gamma; X \vdash_p v_p \tau' : \tau[\tau'/\alpha]} \text{(T-TApp)} \quad \frac{S; \Gamma; X \vdash_q e : \tau}{S; \Gamma; X \vdash_p \langle e \rangle_q : \tau} \text{(T-Bracket)} \\
\\
\frac{S; \Gamma; X \vdash_p v_p : \tau' \quad S; \Gamma, \bar{x}:(\bar{p}, \bar{\tau}_x); \bar{x} \vdash_p D \bar{\tau} \bar{x} : \tau'' \quad S; \Gamma \vdash \text{unify}(\tau', \tau'') : \bar{x} \doteq \bar{v} \quad S; \Gamma, \bar{x}:(\bar{p}, \bar{\tau}_x), \bar{x} \doteq \bar{v}, v_p \doteq D \bar{\tau} \bar{x}; X', \bar{x} \vdash_p e_1 : \tau \quad S; \Gamma; X' \vdash_p e_2 : \tau}{S; \Gamma; X, X' \vdash_p \text{match } v_p \text{ with } D \bar{\tau} \bar{x} \rightarrow e_1 \text{ else } e_2 : \tau} \text{(T-Match)} \\
\\
\frac{S; \Gamma; X \vdash_p v_p : \tau \quad S; \Gamma \vdash \{x:\tau \mid x = v_p\} :: \star}{S; \Gamma; X \vdash_p v_p : \{x:\tau \mid x = v_p\}} \text{(T-Val)} \quad \frac{S; \Gamma; X \vdash_q e : \tau' \quad S; \Gamma \vdash \tau' <: \tau}{S; \Gamma; X, X' \vdash_p e : \tau} \text{(T-Sub)}
\end{array}$$

The typing judgment $S; \Gamma; X \vdash_p e : \tau$ above states that an expression e , when typed with the privilege of principal p in an environment Γ and signature S , has type τ . The set X records a subset of the variables in Γ , and each element of X represents a capability to use an assumption in Γ . The rule ($T\text{-Data}$) requires data constructors to be used only in code granted the appropriate privilege. In the second premise of ($T\text{-Match}$), we type check a pattern $D \bar{\tau} \bar{x}$ to ensure that data constructors are also destructured in a context with the appropriate privilege.

In ($T\text{-Var}$) we type a non-affine variable x by looking up its type in the environment and checking that the privilege of the context matches that of the variable. ($T\text{-VarA}$) is similar, but additionally allows an affine variable to be used only when a capability for its use appears in X . Unlike linear typing, affine assumptions need not always be used. ($T\text{-Drop}$) allows an arbitrary number of assumptions X' to be forgotten, and for e to be checked with a privilege q that is not greater than privilege p that it has been granted. An expression is granted privilege by enclosing it in angle brackets, as shown in ($T\text{-Bracket}$).

Returning to the second premise of ($K\text{-Dep}$), we check a type-level term v_p with the privilege of \top . The intuition is that in well-typed programs, type-level terms have no operational significance and, as such, cannot violate information-hiding. We also check v_p in ($K\text{-Dep}$) with an empty set of capabilities X . According to the well-formedness rule of kind $\tau \rightarrow \kappa$, no well-formed type constructors can be applied to an affine value, so a type-level term like v_p never uses an affine assumption.

In (T-Fix), we require fixed variables f to be given a non-affine type, and for the recursive expression to not capture any affine assumptions. Note that the variable f is bound in the environment with the tag $\mu = \text{fix}$. In the last premise of (K-Dep), we rule out the use of fix-bound variables in type-level values. In (T-Abs), we check that the type of the formal parameter is well-formed, and type check the body in an extended context. We record the privilege p of the program point at which the variable x was introduced to ensure that x is not destructed in unprivileged code in the function-body e . In the conclusion of (T-Abs), we use the auxiliary relation $Q(X, \tau)$, which attaches an affine qualifier to τ if the function captures any affine assumptions from its environment. (T-TAbs) is similar. Typing let-expressions is standard, with the addition that the second premise of (T-Let) ensures that the let-bound variable x does not escape its scope in the type τ_2 . When typing an application $v_q v_r$ in (T-App), we split the affine assumptions among the sub-terms. We allow v_q to be a possibly affine function type—the shorthand $?\tau$ captures this, and we use the same notation in (T-TApp).

We illustrate pattern-matching using an example from FileRM. Consider matching a value v_q of type tracked string (F file) against a pattern L string $x y$. When checking the true-branch, we record several term equalities that capture the runtime behavior of pattern matching. These assumptions will be used by our theorem prover in discharging proofs of refinement formulas (via the type conversion relation, discussed shortly). In our example, one such equality assumption is, clearly, $v_q \doteq \text{L string } x y$. However, with FINE’s value-indexed types, we can also infer equalities for some of the pattern-bound variables. In particular, by unifying the type of the scrutinee, tracked string (F file), with the type of the pattern, tracked string y , we can infer $y \doteq \text{F file}$.

In (T-Match), we split the affine assumptions between v_q and the branches. In the second premise, we type the pattern and in the third premise, unify the type of the scrutinee with the type of the pattern to compute equalities among the term indices—the definition of the unification judgment is standard and we omit it from our presentation. The fourth premise checks e_1 with the computed equality assumptions. The last premise checks e_2 with no additional assumptions. A variation in which e_2 is checked with a disequality **forall** $\bar{x}. v_q \neq D \bar{\tau} \bar{x}$ is also feasible. However, in practice, we use n -way exhaustive pattern matching (**match** x **with** $P_1 \rightarrow e_1 \dots P_n \rightarrow e_n$) and derive disequalities by relying on axioms that discriminate data constructors, e.g., **forall** $D_1, D_2, \bar{x}_1, \bar{x}_2, \bar{\tau}_1, \bar{\tau}_2. D_1 \neq D_2 \Rightarrow D_1 \bar{\tau}_1 \bar{x}_1 \neq D_2 \bar{\tau}_2 \bar{x}_2$.

We use (T-Val) to give values a precise singleton type using an equality refinement. This is useful in bootstrapping the type conversion relation, used in the second premise of (T-Sub), and defined below. Type conversion $S; \Gamma \vdash \tau <: \tau'$ is a reflexive, transitive relation without any structural rules, e.g., contra- and co-variant subtyping in function types. The type system of CIL uses nominal subtyping, and structural rules of this form are not easily translated. The rule (S-Refine) is our interface to the solver—we discuss this with an example shortly. The rule (S-Sub) treats a refined type as a subtype of the underlying type. Type conversion includes an equivalence relation on types $S; \Gamma \vdash \tau \cong \tau'$. In this judgment, (E-Sub) allows a type-level term v_p to be equated with v'_p when an assumption $v_p \doteq v'_p$ appears in the context.

Type conversion: $S; \Gamma \vdash \tau <: \tau'$, $S; \Gamma \vdash \tau \cong \tau'$ and $S; \Gamma \vdash e \cong e'$

Where $S; \Gamma \models \phi$ is the first-order logic entailment relation

$$\begin{array}{c}
\frac{S; \Gamma \vdash \tau_1 \cong \tau_2}{S; \Gamma \vdash \tau_1 <: \tau_2} \text{(S-Eq)} \quad \frac{S; \Gamma \vdash \tau_1 <: \tau_2 \quad S; \Gamma \vdash \tau_2 <: \tau_3}{S; \Gamma \vdash \tau_1 <: \tau_3} \text{(S-Trans)} \quad \frac{S; \Gamma \vdash \tau <: \tau' \quad S; \Gamma, x:(p, \tau) \models \phi}{S; \Gamma \vdash \tau <: \{x:\tau' \mid \phi\}} \text{(S-Refine)} \\
\frac{}{S; \Gamma \vdash \{x:\tau \mid \phi\} <: \tau} \text{(S-Sub)} \quad \frac{}{S; \Gamma \vdash \tau \cong \tau} \text{(E-IdT)} \quad \frac{}{S; \Gamma \vdash v_p \cong v_p} \text{(E-IdV)} \quad \frac{S; \Gamma \vdash \tau_1 \cong \tau'_1 \quad S; \Gamma \vdash \tau_2 \cong \tau'_2}{S; \Gamma \vdash \tau_1 \tau_2 \cong \tau'_1 \tau'_2} \text{(E-App)} \\
\frac{S; \Gamma \vdash \tau_1 \cong \tau'_1 \quad S; \Gamma \vdash v_p \cong v'_p}{S; \Gamma \vdash \tau_1 v_p \cong \tau'_1 v'_p} \text{(E-Dep)} \quad \frac{v_p \doteq v'_p \in \Gamma \vee v'_p \doteq v_p \in \Gamma}{S; \Gamma \vdash v_p \cong v'_p} \text{(E-Subst)} \quad \frac{\forall i, j \quad S; \Gamma \vdash \tau_i \cong \tau'_i \quad S; \Gamma \vdash v_j \cong v'_j}{S; \Gamma \vdash D \bar{\tau} \bar{v} \cong D \bar{\tau}' \bar{v}'} \text{(E-D)}
\end{array}$$

The key rule in type conversion related to refinement typing is (S-Refine). This rule allows a type τ to be promoted to a refined type $\{x:\tau' \mid \phi\}$ when τ is a subtype of τ' , and when our solver can deduce the formula ϕ from the typing context. The entailment relation $S; \Gamma \models \phi$ is standard—we illustrate its behavior using an example from §2.2. When typing the main loop of ConfWeb, we are required to construct a derivation of the form $S; \Gamma \vdash s : \{x:\text{st} \mid \text{In } a \ x\}$, where (dropping principals for clarity) $\Gamma = s:\text{st}, a:\text{attr}, b:\{x:\text{bool} \mid x=\text{true} \Rightarrow \text{In } a \ s\}, b \doteq \text{true}$. We construct this derivation by using (T-Sub) with (T-Val) in the first premise to derive $S; \Gamma \vdash s : \{x:\text{st} \mid x=s\}$, and a derivation of $S; \Gamma \vdash \{x:\text{st} \mid x=s\} <: \{x:\text{st} \mid \text{In } a \ x\}$ in the second premise. This latter derivation proceeds by using (S-Sub), where we deduce $S; \Gamma, x:\{x:\text{st} \mid x=s\} \models \text{In } a \ x$ by using Z3 to show that the theory $(s:\text{st}, a:\text{attr}, b:\text{bool}, b=\text{true} \Rightarrow \text{In } a \ s, b=\text{true}, x:\text{st}, x=s, \text{not}(\text{In } a \ x))$ is unsatisfiable. Importantly, FINE’s type system ensures that the theories we generate never contain any affine assumptions, thus eliminating the need for

a linear logic prover.

3.3 Dynamic semantics

The operational semantics of FINE is instrumented to account for two program properties. First, our semantics places affinely typed values in a memory M . Reads from the memory are destructive—this allows us to prove that in well-typed programs, affine values are never used more than once. The semantics also tracks the privilege of expressions by propagating brackets through reductions, which is useful in showing an information-hiding property for our module system. The main judgment is written $(M, e) \xrightarrow{p} (M', e')$, and states that given an initial memory M an expression e steps to e' and updates the memory to M' . The p -superscript indicates that e steps while using the privilege of the principal p .

Dynamic semantics

Where a memory $M ::= (x, v_p), M \mid \cdot$

$p\text{-Destruct ctx. } E_p ::= \bullet \mid E_p e \mid E_p \tau \mid \text{match } E_p \text{ with } D \bar{\tau} \bar{x} \rightarrow e_1 \text{ else } e_2$			
$\frac{(M, e_1) \xrightarrow{p} (M', e'_1)}{(M, \text{let } x = e_1 \text{ in } e_2) \xrightarrow{p} (M', \text{let } x = e'_1 \text{ in } e_2)} \text{(R-Ctx)}$	$\frac{e \xrightarrow{p} e'}{(M, e) \xrightarrow{p} (M, e')} \text{(R-Pure)}$	$\text{let } x = v_p \text{ in } e \xrightarrow{p} e[v_p/x] \text{(R-Let)}$	
$\langle v_p \rangle_p \xrightarrow{p} v_p \text{(R-Strip)}$	$\langle \langle v_q \rangle_q \rangle_r \xrightarrow{p} \langle v_q \rangle_q \text{(R-Nest)}$	$\langle \lambda x:\tau. e \rangle_q \xrightarrow{p} \lambda y:\tau. \langle e[\langle y \rangle_p/x] \rangle_q \text{(R-Extl)}$	$\langle \Lambda \alpha::\kappa. e \rangle_q \xrightarrow{p} \Lambda \alpha::\kappa. \langle e \rangle_q \text{(R-ExtL)}$
$\frac{e \xrightarrow{q} e'}{\langle e \rangle_q \xrightarrow{p} \langle e' \rangle_q} \text{(R-Br)}$	$\lambda x:t. e v_p \xrightarrow{p} e[v_p/x] \text{(R-Beta)}$	$\Lambda \alpha::\kappa. e \tau \xrightarrow{p} e[\tau/\alpha] \text{(R-TBeta)}$	$\text{fix } f:t.v_p \xrightarrow{p} v_p[\text{fix } f:t.v_p/f] \text{(R-Fix)}$
$\frac{S; \cdot \vdash v_p : \tau \quad S; \cdot \vdash \tau :: \mathbf{A} \quad M' = M, (x, v_q) \quad x \text{ fresh}}{M, v_p \xrightarrow{p} M', x} \text{(R-Alloc)}$		$\frac{M = M_1, (x, v_q), M_2}{M, E_p[x] \xrightarrow{p} (M_1, M_2), E_p[v_q]} \text{(R-Destruct)}$	
$\frac{v_p \neq x \quad \text{unify}(v_p, D \bar{\tau} \bar{x}) : \theta}{\text{match } v_p \text{ with } D \bar{\tau} \bar{x} \rightarrow e_1 \text{ else } e_2 \xrightarrow{p} \theta(e_1)} \text{(R-Then)}$		$\frac{v_p \neq x \quad \forall \theta. \neg \text{unify}(v_p, D \bar{\tau} \bar{x}) : \theta}{\text{match } v_p \text{ with } D \bar{\tau} \bar{x} \rightarrow e_1 \text{ else } e_2 \xrightarrow{p} e_e} \text{(R-Else)}$	

Reduction rules that do not involve reading from or writing to memory are written $e \xrightarrow{p} e'$. All the interesting rules that manage privileges and brackets fall into this fragment. Redundant brackets around p -values can be removed using (R-Strip). However, not all nested brackets can be removed, as (R-Nest) shows. In (R-Extl), a λ -binder is extruded from a function with q -privilege so that it can be applied to a p -value. We have to be careful to enclose occurrences of the bound variable in e within p -brackets, to ensure that e treats its argument abstractly. Finally, (R-Br) allows evaluation to proceed under a bracket $\langle \cdot \rangle_q$ with q -privilege. The rules (R-Alloc) and (R-Destruct) model memory operations. The rule (R-Alloc) is applicable non-deterministically. It allocates a new location x for an affine value v_p into the memory M and replaces v_p with x . When a location x is in destruct position, (R-Destruct) reads a value v_p from M and deletes x .

Theorem 1 establishes the soundness of FINE through the standard progress and preservation lemmas. In the statement below, all free variables are implicitly universally quantified. Additionally, we say that a memory M is typeable with an environment $S; \Gamma$, if $S; \cdot \vdash_p M(x) : \Gamma(x)$, for each location $x \in \text{dom}(M)$. In addition to showing that well-typed programs do not go wrong, our soundness result guarantees that affine values are destructed at most once—a result that shows that state changes are modeled accurately. The proof appears in the appendix.

Theorem 1 (Soundness): *For all well-formed signatures S ; environments Γ ; non-values e ; and memories M typeable with $S; \Gamma$, the following statements are true:*

- 1) *If $S; \Gamma; \text{dom}(M) \vdash_p e : \tau$ then there exists M', e' such that $M, e \xrightarrow{p} M', e'$.*
- 2) *If $S; \Gamma; X \vdash_p e : \tau$ and $M, e \xrightarrow{p} M', e'$ for some p, M', e' , and $X \subseteq \text{dom}(M)$; then, there exists Γ', X' such that $S; \Gamma'; X' \vdash_p e' : \tau$ and M' is typeable with $S; \Gamma'$. Furthermore, for $\Delta_X = (\text{dom}(M) \cup \text{dom}(M')) \setminus (\text{dom}(M) \cap \text{dom}(M'))$ if $\text{dom}(M') \supseteq \text{dom}(M)$ then $X' = X \cup \Delta_X$; otherwise $X' = X \setminus \Delta_X$.*

3.4 Reasoning about the security of FINE programs

FINE allows programmers to specify conditions for correct policy enforcement and the type system checks that these conditions are satisfied. But, the onus is on the programmer to get these specifications right. For example, in the FileRM module of §2.1, wrongly assuming ($\text{forall } p:\text{prin. CanRead } p \text{ f} \Rightarrow \text{CanRead } p \text{ g} \Rightarrow \text{CanFlow } (F \text{ f}) (F \text{ g})$) (instead of the Atomicflow assumption) would destroy any meaningful confidentiality property intended for FileRM to enforce. Similarly, in the Authentication module, forgetting to declare the cred type `private` would allow adversaries to forge credentials. In neither case would FINE’s type checker complain. However, the metatheory of FINE provides a useful set of primitives using which an expert can prove high-level security properties. In prior work on the Fable calculus, we adopted a similar approach and showed how the metatheory of Fable could be used to prove high-level security properties (e.g., noninterference) for encodings of information flow, provenance tracking, and role-based access control. We anticipate a similar strategy being effective for FINE. Additionally, in §4, we discuss how tools like model checkers can complement FINE and be used to establish that FINE programs correctly enforce high-level security goals.

In addition to type soundness, the metatheory of FINE yields two general purpose security properties—proofs appear in the appendix. The first, corresponding to a secrecy property, is value abstraction. The theorem below states that a program e without p -privilege cannot distinguish p -values. As a corollary, we can also derive an integrity property, namely that a program without p -privilege cannot manufacture a p -value to influence the behavior of code with p -privilege. Note that this theorem appeals only to the pure fragment of our reduction rules—affine typing plays no special role in value abstraction. Additionally, observe that this result applies to selective information sharing/hiding between multiple principals, as FINE’s module system includes a lattice of principals ordered by the $p \sqsubseteq q$ relation. Finally, although this theorem applies to the abstraction of a single value from the p module exported at type τ , the program e can contain code from several principals.

Theorem 2 (Value abstraction): *For well-formed signatures S and non-values e , if e uses a p -value x but is well-typed without p -privilege, (i.e., $S; x:(p, \tau); x \vdash_q e : \tau'$ and $p \sqsubseteq q \notin S$) and, except for $\langle x \rangle_p$, e is free of r -brackets $\langle \cdot \rangle_r$, for any r where $p \sqsubseteq r \in S$; then, for any pair of τ -typed values v_p^1 and v_p^2 , (i.e., $S; \cdot \vdash_p v_p^i : \tau$, $i \in \{1, 2\}$) such that $e[v_p^1/x] \stackrel{q}{\rightsquigarrow} e_1$, there exists e' such that $e_1 = e'[v_p^1/x]$ and $e[v_p^2/x] \stackrel{q}{\rightsquigarrow} e'[v_p^2/x]$.*

4 Compiler implementation and application experience

We have implemented a prototype compiler, currently approximately 20,000 lines of F# code extending a front-end and IL generation libraries derived from the F# compiler [23]. The type-preserving translation of FINE to CIL accounts for a significant fraction of the complexity. Our compiler currently generates .NET assemblies that allow FINE programs to easily interface with modules defined in F#. Interoperability with the rest of .NET allows us to write only security critical parts of an application in FINE, leaving the rest to other, more commonly used languages.

The table below shows several small reference monitors in FINE, their size, the number of proof obligations generated during type checking, and parsing and type checking time (on 3.2 GHz Pentium Core Duo desktop running Windows Vista). Most benchmarks contain dense security critical code, where nearly every function call demands proving refinement formulas. Our results show that using an external solver to discharge these proofs (as opposed to constructing them by hand as in Fable or Aura) is critical for practical programming. We expect the checking time to improve significantly as we move from naïve representations of typing environments (currently association lists) to more efficient data structures.

Name	LOC	# pf. obl.	parsing/type checker time (s)
AuthAC	34	1	0.28
FileRM	120	36	1.64
FileAutomaton	121	3	0.45
IFlow	127	22	0.84
HealthWeb	318	19	6.41
DynDKAL	336	34	1.26
Lookout	519	23	2.73
ConfRM and ConfWeb	647	57	4.01
ProofLib	9943	0	19.28
Total	2222 (+ 9943)	195	17.62 (+ 19.28)

4.1 Modeling CONTINUE

Our most substantial example is the modeling of the security policy of CONTINUE. CONTINUE’s authors provided us with a specification of its policy, partly in natural language and partly as specification in the Alloy modeling language [12]. Starting from this specification, we implemented ConfRM to enforce a policy that contains 9 phases and 12 actions. Policy assumptions in ConfPolicy describe when each action is permissible, and a function exposed in the external interface of ConfRM (with a suitable refinement type on the state) mediates access to this action. In addition, each action corresponds to a particular web request handled by ConfWeb.

A significant fragment of the Alloy specification for CONTINUE is devoted to specifying validity conditions on the authorization state. For example, in any given state, validity requires the assignment of papers to reviewers to respect the conflict of interest constraints. We found it relatively straightforward to express several of these validity constraints, although our implementation has yet to cover all the features of CONTINUE’s specification. One simplifying assumption we make is that there is a unique phase for the entire conference. In contrast, the Alloy specification associates a phase with each paper, and different papers can be in different phases at any given time. Extending our `attr` type to account for this complexity is possible, though we have yet to implement this.

Our experience with CONTINUE illustrates an important aspect of FINE. Tools like Alloy are useful for reasoning abstractly about policies and establishing that these correctly specify high-level security goals. However, the abstract analysis of policies in Alloy is disconnected from system implementations that are expected to enforce these policies. FINE, in contrast, does not attempt to validate policies, but provides assurance that system implementations properly enforce their policy specifications. We view these two approaches as complementary and expect their combination to be a potent tool for security analysis of system implementations. For example, the Alloy specification includes assertions to check that no sequence of actions allows a principal to read or write a review when there is a conflict of interest. We plan to investigate using the metatheory of FINE and the types of ConfRM, in conjunction with a tool like Alloy, to prove such facts of our implementation.

4.2 Other benchmarks

The benchmark FileRM extends the example from §2.1 to account for confidentiality and integrity concerns when tracking information flow. Recall that in FileRM the lattice of security labels was derived from a specification of access control permissions using the AtomicFlow assumption. To type check FileRM using Z3, we needed to rewrite the AtomicFlow assumption to the form shown below. To reason about formulas that use nested quantifiers, Z3 relies on a pattern-based instantiation mechanism that requires all bound variables (p in AtomicFlow) to be guarded by non-equality predicates. Note that this is not a fundamental limitation of FINE. We are currently investigating the use of first-order solvers to reason directly about quantified formulas without this restriction. For example, a customized version of Coq’s `firstorder` tactic can discharge proofs of the CanFlow proposition using the assumption AtomicFlow as shown in §2.1.

```
assume CW:IsPrin Admin && IsPrin (U "Alice") && IsPrin (U "Bob") && ...
assume AtomicFlow: forall f:file, g:file.
  (forall p:prin. (IsPrin p && CanRead p g) => CanRead p f) => CanFlow (F f) (F g)
```

Of the other benchmarks, AuthAC is a small purely permission-based access control monitor for files combined with password-based authentication. FileAutomaton is a reference monitor that implements an automaton-like policy on files, where, through the use of dependent and affine types, a file handle is indexed with a value indicating its current state, e.g., Open, Closed etc. A similar idiom could be used in ConfRM to associate phases with papers, instead of a global phase for the entire conference. IFlow is an implementation of a traditional information flow policy using a three-point lattice of labels which does not require the nested quantifiers of FileRM. HealthWeb is a reference monitor for an application that manages a database of electronic medical records. It enforces a stateful authorization policy. DynDKAL is an interpreter for an authorization logic; it uses refinement types to ensure that instantiations of quantified assumptions in policies is performed correctly. Lookout is the core reference monitor of a plugin-based email client we have started to build. This program mixes stateful authorization in the style of ConfRM with information flow tracking in the style of FileRM.

Finally, ProofLib is an automatically generated program, our largest test case by far. This program makes no use of refinement types and is used as a utility by our type-preserving compiler to represent proof terms. We include it here to give the reader a sense of the cost of dependent type checking for larger programs.

5 Related work and conclusions

Several programming languages and proof assistants use dependent types, including Agda [17], Coq [2], and Epigram [16]. All of these systems can be used to verify full functional correctness of programs. However, to ensure logical consistency of the type system, these languages exclude arbitrary recursion, making them less applicable for general-purpose programming. Projects like YNot [4] and Guru [19] aim to mix effects like non-termination with dependently typed functional programming; YNot also supports programming with state in an imperative style. Restrictions in both languages ensure that proofs are pure, ensuring that logical consistency is preserved. All of these systems include automation and tactic languages, but programmers still need to construct interactive proofs for their code. In contrast, FINE targets weaker, security properties; forgoes logical consistency in favor of practical programming by including recursion; and automatically synthesizes proof terms using an SMT solver. FINE also provides affine types to allow the enforcement of state-modifying policies, which could be expressed in YNot, but not easily in the other languages.

Dependent types have also been used for security verification. Jif [5] uses a limited form of dependent typing to express dynamic information flow policies. Aura [13] is specialized for the enforcement of policies specified in a policy language based on an intuitionistic modal logic. This makes Aura less applicable to policies specified in other logics, e.g., the Datalog-based policy language of Dougherty et al. [7], and Aura cannot model stateful policies. Aura provides logical consistency by separating types from propositions and excluding arbitrary recursion in proof terms. However, proof terms in Aura are always programmer-provided. As such, Aura is positioned as an intermediate language, rather than a source-level language. Fable [21], is another intermediate language for security verification that uses dependent types. Fable uses a two-principal module system. FINE’s module system generalizes Fable’s, with support for a lattice of multiple principals. FINE is also related to λ AIR [22], a calculus that targets the enforcement of declassification policies. λ AIR’s combination of affine and dependent types does not lend itself to integration with a solver and it was never implemented.

Refinement types in FINE are related to a similar construct in RCF [1]. Refinement formulas in RCF are drawn from an unsorted logic, rather than using dependent-type constructors, as we do. The lack of dependent type constructors in RCF makes it difficult to derive typeable proof terms, crucial to our goal of a type-preserving compiler for FINE. Additionally, without dependent type constructors, it appears impossible to enforce information flow policies in RCF, although RCF’s implementation, F7, does include dependent type constructors. RCF also lacks support for stateful authorization policies, although recent work shows how stateful policies can be modeled in F7 using a refined state monad [3]. However, the soundness of this encoding relies on a trusted compilation of the program in a linear, store-passing style. FINE’s type system also allows the use of refined state monads, but, as discussed in §2.3, affine types in FINE also admit other stateful programming idioms.

FINE is also related to hybrid-typed languages that use refinement types, like Sage [9]. Sage uses a trusted external solver to discharge proofs; we extract typeable proof terms from Z3 rather than trusting it. Another difference is that Sage automatically insert runtime checks when the solver fails to discharge a proof obligation. Failed runtime checks can cause subtle leaks of information, so automatic insertion of runtime checks is not yet a feature of FINE, where security is the primary concern—we plan to investigate adding support for automatic policy checking in the future.

Conclusions. This paper has presented FINE, a language for enforcing rich, stateful authorization and information flow policies. Our experience constructing several reference monitors provides initial evidence that programming in FINE is practical, due in part to the use of an automated solver to ease the proof burden, and that FINE can be used to check the enforcement of security policies commonly applied to software.

Acknowledgments. Thanks to Shriram Krishnamurthi for providing us with a specification of CONTINUE’s policy; to Nikolaj Bjørner and Leonardo de Moura for help with Z3; to Karthik Bhargavan, Johannes Borgstroem, Cédric Fournet, and Andy Gordon for numerous discussions about this work.

A Additional judgments

In this section, we define three judgments omitted from the main body of the paper. The first judgment is of the form $S; \Gamma \vdash \text{unify}(\tau, \tau') : \bar{x} \doteq \bar{v}$ and stands for unification of the types τ and τ' using the substitution of value variables \bar{x}

in τ' . This judgment is used among the premises of the typing rule for pattern matching to compute induced equalities among pattern bound variables and the value indices in the type of a scrutinee.

The second judgment is of the form $v_p \doteq \theta(D \bar{\tau} \bar{x})$ and states that a value v_p matches a pattern $D \bar{\tau} \bar{x}$ producing the substitution θ of pattern-bound variables. This judgment is used in the dynamic semantics to determine which branch of a pattern match should be taken.

The final judgment models the first-order entailment relation $S; \Gamma \models \phi$ as a typing derivation of a proof term of type proof ϕ that witnesses the validity of ϕ , where $\text{proof} :: \star \rightarrow \star$ is a type constructor. This judgment makes use of an auxiliary function $\text{refinements}(\Gamma)$ which collects assumptions in Γ given refinement types as additional assumptions of proof terms.

Unification and formula entailment

Where $\theta = x \doteq v, \theta \mid \cdot$.

$$\frac{\frac{\text{unify}(\bar{v}_p, \bar{x}) : \theta}{S; \Gamma \vdash \text{unify}(T \bar{\tau} \bar{v}_p, T \bar{\tau} \bar{x}) : \theta} \quad \frac{\text{unify}(v, v') : \theta_1 \quad \text{unify}(\bar{v}, \theta_1(\bar{x})) : \theta_2}{\text{unify}((v, \bar{v}), (v', \bar{x})) : \theta_1, \theta_2} \quad \frac{}{\text{unify}(v, x) : x \doteq v} \quad \frac{\text{unify}(\bar{v}_p, \bar{x}) : \theta}{\text{unify}(D \bar{\tau} \bar{v}, D \bar{\tau} \bar{x}) : \theta}}{\frac{\Gamma' = \text{refinements}(\Gamma) \quad S; \Gamma, \Gamma'; \cdot \vdash_p v_p : \text{proof } \phi}{S; \Gamma \models \phi} \text{ (T-Entail)}}$$

where

$$\begin{aligned} \text{refinements}(y:(p, \{x:\tau \mid \phi\})) &= z:(p, \text{proof}(\phi[y/x])) && z \text{ fresh} \\ \text{refinements}(y:(p, \tau)) &= \cdot && \text{when } \tau \neq \{x:\tau' \mid \phi\} \\ \text{refinements}(v_1 \doteq v_2) &= \cdot \\ \text{refinements}(\alpha::\kappa) &= \cdot \\ \text{refinements}(\Gamma, \Gamma') &= \text{refinements}(\Gamma), \text{refinements}(\Gamma') \\ \text{refinements}(\cdot) &= \cdot \end{aligned}$$

B Soundness of FINE

Definition 1 (Well-formed data constructors). *A data constructor $D: (p, t)$ is well-formed iff the following is true:*

The type t is of the form

$$\text{forall } a1::k1, \dots, a_n::k_n. x1:t1 \rightarrow \dots x_n:t_n \rightarrow t'$$

That is, all type arguments precede all term arguments.
We call t' the constructed type of D .

The constructed type t' of D is of the form $T a1..a_n v1'..vm'$,
for some type constructor T . I.e., D is injective in its type arguments,
and always constructs a constructed type T .

Additionally, a program e has well-formed data constructors if every data constructor D in the text of e is fully applied.

Definition 2 (Well-formed signature). *Well-formedness of a signature S is defined inductively as*

1. $S = S', T::k \Rightarrow$
 S' is well-formed
and $S' \vdash k$
2. $S = S', D:(p, t) \Rightarrow$
 S' is well-formed
and $S'; - \vdash t :: k$
and k in $\{*, A\}$
and p in S'
and $D:(p, t)$ is well-formed
3. $S = \text{FOL}, p1 < p2, \dots, p < p', S'$ where FOL is the basic signature for first-order logic with equality specialized to a set of ground types $\text{TT} = \{T1 \dots Tn\}$ and all the principal names p_i are distinct, and the relation $p < p'$ forms a partial order.

```

FOL=
  T1::*, ..., Tn::*
  Eq_1::T1 -> T1 -> *
  ...
  Eq_n::Tn -> Tn -> *
  ReflT1: x:T1 -> Eq_1 x x
  ...
  ReflTn: x:Tn -> Eq_n x x
  And::* -> * -> *
  Or::* -> * -> *
  Not::* -> *
  True::*
  proof::* -> *
  tt: True

```

Definition 3 (Well-formed environment). An environment $Env=S;G;X$ is well-formed iff $S;G$ bind distinct names and all of the following are true

1. $Env=S;G;X, x \Rightarrow x \text{ in } \text{dom}(G)$ and
 $x \text{ not in } X$ and
 $S;G;X$ is well-formed
2. $Env=S;G, x:(p, t);- \Rightarrow \text{FreeVariables}(t) \leq \text{dom}(G)$ and
 $p \text{ in } S$ and
 $S;G;-$ is well-formed
3. $Env=S;G, e_1=e_2;- \Rightarrow \text{FreeVariables}(e_1) \leq \text{dom}(G)$ and
 $\text{FreeVariables}(e_2) \leq \text{dom}(G)$ and
 $S;G;-$ is well-formed
4. $Env=S;G, 'a::k;- \Rightarrow k \text{ in } \{*, A\}$ and
 $S;G;-$ is well-formed
5. $Env=S;-;- \Rightarrow S$ is a well-formed signature

Definition 4 (Well-formed memory). Given a signature S , and a memory M , the environment corresponding to M is written $S;G(M)$, where $G(M)$ is defined inductively as:

```

G(.) = .
G(M, (x, v_p)) = G(M), x:(p, t) where S;-;- |-_p v_p:t and S;- |- t :: A

```

A memory M is well-formed when $G(M)$ exists

Lemma 5 (Canonical forms). For all S, G, X, p, t, v_p, Φ ,

- (A1) $S;G;X$ well-formed
- (A2) $S;G;X \text{ |-}_p v_p : ?(x:t_1) \rightarrow t_2 \Rightarrow \text{exists } y. v_p=y \text{ or exists } e, v_p = \backslash x:t.e$
- (A3) $S;G;X \text{ |-}_p v_p : ?(\backslash 'a::k.e) \Rightarrow \text{exists } y. \text{exists } e, v_p = / \backslash 'a::k.e$
- (A4) $S;G;- \text{ |-}_p v_p : \{x:t \mid \Phi\}$
 $\Rightarrow \text{exists } v', S;G, \text{refinements}(G);- \text{ |-}_\text{bot } v' : \text{proof } \Phi[v/x]$

Proof. By induction on the structure of the typing derivation, appealing to fully-applied data constructors to exclude $(Dv1 : t_1 \rightarrow t_2)$ etc. and well-formedness of data constructors to ensure that data constructors do not introduce refined types. \square

Theorem 6 (Progress). For all $S M e t p$,

- (A1) $S;G(M); \text{dom}(M)$ well-formed and
- (A2) $S;G(M); \text{dom}(M) \text{ |-}_p e : t$
 $\Rightarrow \text{exists } v_p, e=v_p \text{ or exists } M' e', (M, e) \sim_p (M', e')$

Proof. Proof: By induction on the structure of (A2). (We rely on implicit applications of R-Pure in the proof below.)

Case (T-Var):

Impossible. From premises of (T-Var) we have $x \text{ in } \text{dom}(M)$ and $S;G(M) \text{ |- } t :: *$. From well-formedness of memory M , we have that for all $x \text{ in } \text{dom}(M)$, the type of x has kind A .

Case (T-VarA):

From the premises of (T-VarA) we have x in $\text{dom}(M)$, satisfying the premise of (R-Destruct) and allowing a step to be taken.

Case (T-Val, T-Data, T-Abs, T-Univ):

$v_p, \lambda x:t.e, \lambda a::k.e$ are all values for p .

Case (T-Fix):

$\text{fix } f:t.e \sim p \sim e$ [fix $f:t.e/f$] using [R-Fix]

Case (T-App):

($e = v_p v_p'$): From the first antecedent of (T-App), we have $S; - \vdash \vdash v_p : ?(x:t1) \rightarrow t2$.
From Lemma (Canonical forms), we can conclude that either

Sub-case (exists $e', v_p = \lambda x:t1.e'$.)

For the conclusion, we apply [R-Beta] producing $e' [v_p/x]$ as the witness for the right-side of the goal.

Sub-case (exists y in $\text{dom}(G(M)) = \text{dom}(M)$. $v_p = y$).

Step using (R-Destruct).

Case (T-TApp):

($e = v_p t$): From the first antecedent of (T-TApp) and Lemma (Canonical forms), we can conclude that either

Sub-case (exists $e', v_p = \lambda a::k.$)

For the conclusion we apply [R-TBeta] produces $e' [t/'a]$ as the witness on the right-side of the goal.

Sub-case (exists y in $\text{dom}(G(M)) = \text{dom}(M)$. $v_p = y$)

Step using (R-Destruct)

Case (T-Bracket):

a. ($e = \langle e' \rangle_q$): [R-Br] is applicable.

b. ($e = \langle v_q \rangle_q$), where $p \langle \rangle q$: We enumerate sub-cases on v_q .

i. ($v_q = u_q$):

$\langle u_q \rangle_q$ is a p -value, satisfying the left-side of the goal disjunction.

ii. ($v_q = \lambda x:t.e$):

$\langle \lambda x:t.e \rangle_q \sim p \sim \lambda y:t. \langle e[\langle y \rangle_p/x] \rangle_q$ using R-Ext1 satisfying the right side of the goal.

iii. ($v_q = \lambda a::k.e$):

$\langle \lambda a::k.e \rangle_q \sim p \sim \lambda a::k. \langle e \rangle_q$ using R-ExtL satisfying the right side of the goal.

iv. ($v_q = \langle u_r \rangle_r$):

$\langle \langle u_r \rangle_r \rangle_q \sim p \sim \langle u_r \rangle_r$ using R-Nest, satisfying the right-side of the goal.

c. $(e = \langle v_p \rangle_p)$:

$\langle v_p \rangle_p \sim^p v_p$ using [R-Strip], satisfying the right-side of the goal.

Case (T-Match):

$(e = \text{match } v_p \text{ with } \dots)$: The second premises of [R-Then] and [R-Else] form a tautology. So, as long as v_p is not a variable, we can satisfy the right-side of the goal using one of the two rules. If v_p is a variable, then we can step using [R-Destruct].

Cases (T-Sub, T-Drop):

The goal follows directly from the induction hypothesis.

Case (T-Let):

- a. $(e = \text{let } x = e_1 \text{ in } e_2)$: Step using (R-Ctx)
- b. $(e = \text{let } x = v_p \text{ in } e_2)$: Step using (R-Let)

(Note: the rule (R-Alloc) is unnecessary for progress)

□

Lemma 7 (Weakening). Lemma (Weakening for typing judgment):

For all $S \ G_1 \ G_2 \ G \ X_1 \ X_2 \ X \ e \ t \ p$,
(A1) $S; G_1, G_2; X_1, X_2$ well-formed
(A2) $S; G_1, G_2; X_1, X_2 \mid _p e : t$
(A3) $S; G_1, G, G_2; X_1, X, X_2$ well-formed
 $\Rightarrow S; G_1, G, G_2; X_1, X, X_2 \mid _p e : t$

Lemma (Weakening for kinding judgment):

For all $S \ G_1 \ G_2 \ G \ t \ k \ p$,
(B1) $S; G_1, G_2; -$ well-formed
(B2) $S; G_1, G_2 \mid _p t :: k$
(B3) $S; G_1, G, G_2; -$ well-formed
 $\Rightarrow S; G_1, G, G_2; - \mid _p t :: k$

Proof. By mutual induction on the structure of (A2) and (B2), generalizing on G_2 .

The interesting cases of (A2)

Case (T-Var), (T-VarA):

In both cases, we can establish the conclusion by using (T-Var/T-Var-A) in the premise of (T-Drop-A), to drop the additional affine assumptions in X , and note that G_1, G, G_2 binds x iff G_1, G_2 binds x .

Case (T-Abs):

For the first premise of (A2), we use the mutual induction hypothesis to show that the kinding derivation can be weakened to

(A2.1') $S; G_1, G, G_2 \mid _p t :: k$

The third premise of (A2) in this case is of the form

(A2.3) $S; G_1, G_2, x; t; X_1, X_2 \mid _p e : t'$

In this case, we can use the induction hypothesis since we have been careful to generalize on the extended environment G_2 , since the induction hypothesis specifically allows weakening by inserting assumptions in the "middle" of a context. Thus, we can then establish

(A2.3.broken) $S;G1,G,G2,x:t; X1,X,X2,X' \vdash_{\text{p}} e : t'$

However, this derivation is not always of the right shape, since when $X1,X2$ is empty and X is not empty, then when constructing

$S;G1,G,G2;X1,X,X2 \vdash_{\text{p}} \lambda x:t.e : ?(x:t) \rightarrow t'$

with (A2.3.broken) in the premise, the qualifier on the introduced type $?(x:t) \rightarrow t'$ may differ from the qualifier on tf , the type introduced in (A2).

To remedy this, we establish the conclusion by first showing that $S;G1,G,G2;X1,X2$ is well-formed. Next, we use the induction hypothesis to construct

(A2.3') $S;G1,G,G2,x:t; X1,X2,X' \vdash_{\text{p}} e : t'$

Finally, we use an application of (T-Drop-A) in the context $(S;G1,G,G2; X1,X,X2)$ to discard the affine assumptions X . In the premises of (T-Drop-A), we use (T-Abs) with (A2.1') and (A2.3'). The other premises are unchanged.

Case (T-Fix): Similar to (T-Abs), always using (T-Drop-A) in the conclusion to discard the affine assumptions in X .

Case (T-Tabs): Similar to (T-Abs).

The interesting cases of (B2)

Case (K-Arrow): Similar to (T-Abs), where generalization of $G2$ and allowing weakening in the middle of the context is key.

Case (K-Univ): Similar to (K-Arrow).

Case (K-App): The mutual induction hypothesis allows us to establish a weakening for the typing derivation in the second premise.

□

Lemma 8 (Well-kinded typings). Lemma (Well-kinded typings):

For all $S G X e t p$,
(A1) $S;G;X \vdash_{\text{p}} e : t$
 \Rightarrow exists $k, S;G \vdash t :: k$ and k in $\{*, A\}$

Lemma 6.2 (Well-formed kindings):

For all $S G t k$,
(B1) $S;G \vdash t :: k$
 $\Rightarrow S \vdash k$

Lemma 6.3 (Well-formed type conversions):

For all $S G t t' k$,
(C1) $S;G \vdash t :: k$
(C2) $S;G \vdash t <: t'$
 $\Rightarrow S;G \vdash t' :: k$

Proof. Straightforward from mutual induction on the structure of (A1) and (B1).

The interesting cases in (A1)

Case (T-Match):

From the last premise, we ensure that the result type t does not contain any pattern-bound variables.

Case (T-App):

The last premise ensures the result by construction. This ensures that non-values do not escape into types.

Case (T-Abs):

The first two premises ensure that the ascribed type is well-formed.

The cases of (B1) and (C1) are all straightforward. □

Lemma 9 (Substitution for type unification). Given $S;G,x:(p,t),G' \vdash \text{unify}(t', t'') : xi=vi$, and a substitution $s=[v/x]$; then, $S;G,G' \vdash \text{unify}(s(t'), s(t'')) : xi=s(vi)$

Lemma 10 (Substitution). Lemma (Substitution for typing judgment):

For all $S G1 G2 X X' X'' x tx e t v p q tx' \text{Phi}$,
 (A0) $X'=x$ or $X'=\{\}$
 (A1) $S;G1,x:(p,tx),G2;X,X',X''$ well-formed
 (A2) $S;G1,x:(p,tx),G2;X,X',X'' \vdash_{-q} e : t$
 (A3) $S;G1;- \vdash_{-p} v_p : tx$
 (A4) substitution $s=[v_p/x]$
 $\Rightarrow S;G1,s(G2);X,X'' \vdash_{-q} s(e) : s(t)$

Lemma (Substitution for kinding judgment):

For all $S G1 G2 x tx t k v p tx' \text{Phi}$,
 (B1) $S;G1,x:(p,tx),G2;-$ well-formed
 (B2) $S;G1,x:(p,tx),G2 \vdash t :: k$
 (B3) $S;G1;- \vdash_{-p} v_p : tx$
 (B4) substitution $s=[v_p/x]$
 $\Rightarrow S;G1,s(G2) \vdash s(t) :: k$

Lemma (Substitution for type conversion):

For all $S G1 G2 x tx t v p tx' \text{Phi}$,
 (C1) $S;G1,x:(p,tx),G2;-$ well-formed
 (C2) $S;G1,x:(p,tx),G2 \vdash t <: t'$
 (C3) $S;G1;- \vdash_{-p} v_p : tx$
 (C4) substitution $s=[v_p/x]$
 $\Rightarrow S;G1,s(G2) \vdash s(t) <: s(t')$

Lemma (Substitution for type equivalence):

For all $S G1 G2 x tx t v p tx' \text{Phi}$,
 (D1) $S;G1,x:(p,tx),G2;-$ well-formed
 (D2) $S;G1,x:(p,tx),G2 \vdash t \sim t'$
 (D3) $S;G1;- \vdash_{-p} v_p : tx$
 (D4) substitution $s=[v_p/x]$
 $\Rightarrow S;G1,s(G2) \vdash s(t) \sim s(t')$

Proof. By mutual induction on the structure of (A2), (B2), (C2), (D2), generalizing on q , the tail of the environment $G2$, and the set of affine assumptions X,X',X'' .

Cases of (A2)

Case (T-Var):

$G=G1,x:(p,tx),G2$	(A2.1)
$G(y) = (q,ty)$	(A2.2)
$S; G \vdash ty :: *$	(A2.3)
$S; G; - \vdash_{-q} y : ty$	[T-Var]

We consider two sub-cases, depending on whether $x=y$.

Sub-case ($x<>y$):

In this case, we have $s(y)=y$. We have two further sub-cases, depending on whether y in $\text{dom}(G1)$ or y in $\text{dom}(G2)$.

Sub-sub-case (y in $\text{dom}(G1)$):

From the well-formedness of $S;G1,x:(p,tx);-$ and $G1(y)=(q,ty)$, we can conclude that x not in $\text{FV}(ty)$. Thus, $s(ty)=ty$, and we have, as required:

$$S;G1,s(G2);- \vdash_{-q} s(y) : s(ty)$$

Sub-sub-case (y in $\text{dom}(G2)$):

We have y in $\text{dom}(G2)$ and
 $G2(y)=(q,ty) \Rightarrow s(G2(y))=(q, s(ty))$.
The conclusion is immediate.

Sub-case ($x=y$):

In this case, $G(y) = G(x) = (p, tx) = (q, ty)$, and $s(y)=v_p$.

For the conclusion, we apply Lemma (weakening). In order to do this, we must first show that

$$(WF) \quad S;G1,s(G2);X \text{ is well-formed}$$

This is easily accomplished by induction on the length of $G2$, noting that x not in $\text{FV}(s(G2))$, and x not in X .

Now, using Lemma (weakening), WF and assumption (A3), we conclude

$$(Goal.0) \quad S;G1,s(G2);X \vdash_{-p} v_p : tx$$

Finally, for the goal, we note that $p=q$, we use Lemma (Well-kinded typings) on assumption (A3) to establish that $(S;G1;- \vdash_{-} tx :: k)$ and hence that x not in $\text{FV}(tx)$, and finally that $s(tx) = tx$ to get:

$$(Goal) \quad S;G1,s(G2);X \vdash_{-q} v_p : s(tx)$$

Case (T-Var-A):

Identical to (T-Var)

Case (T-Abs):

$$\begin{array}{ll} G=G1,x:(p,tx),G2 & (A2.1) \\ k \text{ in } \{*, A\} & (A2.2) \\ S; G \vdash_{-} t :: k & (A2.3) \\ S; G,y:(q,t); X,X',X'',y \vdash_{-q} e : t' & (A2.4) \\ X=. \Rightarrow tf = ?(y:t) \rightarrow t' & (A2.5) \\ X<.>. \Rightarrow tf = !(y:t) \rightarrow t' & (A2.6) \\ \hline S; G; X,X',X'' \vdash_{-q} \lambda y:t.e : tf & [T-Abs] \end{array}$$

We begin by applying the mutual induction hypothesis for substitution on the kinding judgment to (A2.3) to establish

$$(A2.3') \quad S;G1,s(G2) \vdash_{-} s(t) :: k$$

Next, we apply the induction hypothesis (having generalized on the tail of the environment $G2$, and X,X',X'') to obtain

$$(A2.4') \quad S; G1,s(G2),y:(q,s(t)); X,X'',y \vdash_{-q} s(e) : s(t')$$

We can now construct

(A2') $S; G1, s(G2); X, X'' \mid _q s(\backslash y:t.e) : ?(y:s(t)) \rightarrow s(t')$

using (T-Abs) with (A2.3') and (A2.4') in the premises.

We now consider four sub-cases

Sub-case ($X'=x$ and $X, X''=. .$):

In this case, we are required to construct a derivation with:
 $s(tf) = !(y:s(t)) \rightarrow s(t')$

The fifth premise of A2', gives us the freedom to chose
 $?(y:s(t)) \rightarrow s(t') = !(y:s(t)) \rightarrow s(t')$
as required.

Sub-case ($X'=x$ and $X, X'' <> .$):

Sub-case ($X'=. .$ and $X, X'' <> .$):

In these case, since $X, X', X'' <> \{\}$, A2 concludes with,
 $tf = !y:t \rightarrow t'$

And from the last premise of A2', we have
 $s(tf) = !y:s(t) \rightarrow s(t')$

as required.

Sub-cases ($X'=. .$ and $X, X''=. .$):

In this case, the fifth premise of A2 allows
 $tf = y:t \rightarrow t'$

or

$tf = !y:t \rightarrow t'$

In the the fifth premise of A2', we still have $X, X''=. .$ and we
can chose the qualifier on $s(tf)$ to match the qualifier on tf ,
as needed.

Case (T-Univ):

$G=G1, x:(p, tx), G2$ (A2.1)
 $k \text{ in } \{*, A\}$ (A2.2)
 $S; G, 'a::k; X \mid _q e : t$ (A2.3)
 $X=. => tf = ?(\backslash 'a::k.t)$ (A2.4)
 $X<>. => tf = !(\backslash 'a::k.t)$ (A2.5)
----- [T-Univ]
 $S; G; X \mid _q \backslash 'a::k.e : tf$

Similar to (T-Abs):

We begin by applying the induction hypothesis (having generalized on the tail of the environment
 $G2$, and X, X', X'') to obtain

(A2.3') $S; G1, s(G2), 'a::k; X, X'' \mid _q s(e) : s(t)$

We can now construct

(A2') $S; G1, s(G2); X, X'' \mid _q s(\backslash 'a::k.e) : ?(\backslash 'a::k.t)$

using (T-Univ) with (A2.3') in the premises.

We now consider four sub-cases, similar to (T-Abs), and establish
the affinity qualifier using the flexibility of (A2.4) as needed.

Case (T-Fix):

$$\begin{array}{ll}
G=G1, x:(p, tx), G2 & (A2.1) \\
S; G \vdash t :: * & (A2.2) \\
S; G, f:(q, t); - \vdash_{-q} e : t & (A2.3) \\
\hline
S; G; - \vdash_{-q} \text{fix } f:t.e : t & [T-Fix]
\end{array}$$

Similar to T-Abs.

Case (T-App):

$$\begin{array}{ll}
X1, X2 = X, X', X'' & \\
G=G1, x:(p, tx), G2 & (A2.1) \\
S; G; X1 \vdash_{-q} v_{-q} : ?(y:t1) \rightarrow t2 & (A2.2) \\
S; G; X2 \vdash_{-q} v_{-q}' : t1 & (A2.3) \\
S; G \vdash t2[v_{-q}'/y] :: k & (A2.4) \\
\hline
S; G; X, X', X'' \vdash_{-q} v_{-q} v_{-q}' : t2[v_{-q}'/y] & [T-App]
\end{array}$$

Induction hypothesis on (A2.2) and (A2.3) gives us

$$\begin{array}{ll}
(A2.2') & S; G1, s(G2); X1 \vdash_{-q} s(v_{-q}) : ?(y:s(t1)) \rightarrow s(t2) \\
(A2.3') & S; G1, s(G2); X2 \vdash_{-q} s(v_{-q}') : s(t1)
\end{array}$$

and by the mutual induction hypothesis we have

$$(A2.3') \quad S; G1, s(G2) \vdash s(t2[v_{-q}'/y]) :: s(k)$$

Note that on splitting X, X', X'' into $X1, X2$, the assumption x in X' (if present) goes either in $X1$ or in $X2$, or in neither (if it is absent). In each case, we can use the induction hypothesis with either the left side of premise (A0) or the right side of (A0).

For the conclusion, we require a derivation with the type $s(t2[v_{-q}'/y])$

Using (T-App) with (A2.2') and (A2.3') gives us the type

$$s(t2) [s(v_{-q}')/y]$$

We need to show

$$s(t2[v_{-q}'/y]) = s(t2) [s(v_{-q}')/y]$$

which following from the observation that $x < y$ (y is a bound variable and can be alpha renamed appropriately), and furthermore that y not in $FV(\text{range}(s))$.

Case (T-TApp):

$$\begin{array}{ll}
S; G; X, X', X'' \vdash_{-q} e : ?(\backslash/a::k.t') & \\
S; G \vdash t :: k & \\
\hline
S; G; X, X' \vdash_{-q} e t : t' [t'/a] & [T-TApp]
\end{array}$$

Similar to (T-App), except there is no need for any special management of the affine assumptions.

Case (T-Bracket):

$$\begin{array}{ll}
G=G1, x:(p, tx), G2 & (A2.1) \\
S; G; X, X', X'' \vdash_{-r} e : t & (A2.2)
\end{array}$$

----- [T-Bracket]
 S; G; X, X', X'' |-_q <e>_r : t

Having generalized on the index q on the turnstile of (A2), we can apply the induction hypothesis to (A2.2) to obtain

(A2.2') S; G1, s(G2); X, X'' |-_r s(e) : s(t)

The conclusion follows from an application of (T-Bracket) with (A2.2') in the premise.

Case (T-Match):

G=G1, x:(p, tx), G2 (A2.1)
 X1, X2 = X, X', X'' (A2.2)
 X3 < x1..xn (A2.3)
 S; G; X1 |-_q v_q : t' (A2.4)
 S; G, xi:(q, ti); X3 |-_q D t1..tn x1..xn : t'' (A2.5)
 S; G |- unify (t', t'') : xi=vi (A2.6)
 S; G, xi:(q, ti), xi=vi, v_q=D t1..tn x1..xn; X2, X3 |-_q e1 : t (A2.7)
 S; G; X2 |-_q e2 : t (A2.8)
 ----- [T-Match]
 S; G; X1, X2 |-_q match v_q with
 D t1..tn x1..xn -> e1 : t
 else e2

As in (T-App), the affine assumption in X', if present, either goes to X1 or X2. When using the induction hypothesis, we satisfy premise (A0) by using either side of the disjunct, depending on whether x in X1, X2 or neither.

We use the induction hypothesis to establish

(A2.4') S; G1, s(G2); X1 |-_q s(v_q) : s(t')
 (A2.5') S; G1, s(G2); X3 |-_q s(D t1..tn x1..xn) : s(t'')

Next, we use (A2.6) and appeal to substitution for type unification to show that:

(A2.6') S; G1, s(G2) |- unify(s(t'), s(t'')) : xi=s(vi)

The induction hypothesis on (A2.7) gives us:

(A2.6') S; G1, s(G2), s(xi:(q, ti)), s(xi=vi); s(v_q=D t1..tn x1..xn); X2, X3 |-_q s(e1) : s(t)
 (A2.7') S; G1, s(G2); X2 |-_q s(e2) : s(t)

and use each of these in the premises of (T-Match) for the goal.

Case (T-Sub):

G=G1, x:(p, tx), G2 (A2.1)
 S; G; X, X', X'' |-_q e : t' (A2.2)
 S; G |- t' <: t (A2.3)
 ----- [T-Sub]
 S; G; X, X', X'' |-_q e : t

From the induction hypothesis we get

(A2.2') S; G1, s(G2); X, X'' |-_q s(e) : s(t')

From the mutual induction hypothesis with Lemma (substitution for type conversion), we get

(A2.3') S; G1, s(G2) |- s(t') <: s(t)

The conclusion follows from an application of (T-Sub) with (A2.2') and (A2.3') in the premises.

Case (T-Let):

Straightforward application of induction hypothesis to each of the premises.

Cases (T-Drop, T-Data, T-Val):

Trivial, from the induction hypothesis.

Cases of (B2)

Case (K-Var):

Trivial, since $G1, G2$ binds 'a and kinds have no free variables.

Case (K-TC):

Trivial, since type constructors are bound in S and S is unchanged.

Case (K-Afn):

Follows from the induction hypothesis.

Case (K-Arrow):

$$\begin{array}{ll}
 G=G1, x: (p, tx), G2 & (B2.1) \\
 k, k' \text{ in } \{*, A\} & (B2.2) \\
 S; G \vdash t1 :: k & (B2.3) \\
 S; G, x:t1 \vdash t2 :: k' & (B2.4) \\
 \hline
 S; G \vdash (x:t1) \rightarrow t2 & [K-Arrow]
 \end{array}$$

From the induction hypothesis, we have

(B2.3') $S; G1, s(G2) \vdash s(t1) :: k$

From the induction hypothesis, having generalized on the tail of the environment $G2$, we have

(B2.4') $S; G1, s(G2), x:s(t1) \vdash s(t2) :: k'$

For the conclusion, we can use (K-Arrow) with (B2.3') and (B2.4') in the premises.

Case (K-Univ):

Similar to (K-Arrow), relies on generalization over $G2$.

Case (K-App):

Straightforward, from induction hypothesis applied to each premise.

Case (K-Dep):

$$\begin{array}{ll}
 G=G1, x: (p, tx), G2 & (B2.1) \\
 S; G \vdash t :: t' \rightarrow k & (B2.2) \\
 S; G; . \vdash_{-q} v_q : t' & (B2.3) \\
 \hline
 S; G \vdash t \ v_q : k & [K-Dep]
 \end{array}$$

From the induction hypothesis on (B2.2) we get

(B2.2') $S; G1, s(G2) \vdash s(t) :: t' \rightarrow k$

From Lemma (Well-formed kindings), we have that

$S \vdash t' \rightarrow k$

from which we can conclude that $\text{FreeVars}(t' \rightarrow k) = \{\}$, and hence $s(t') = t'$.

Next, we apply the mutual induction hypothesis on the typing judgment to (B2.3) to produce

(B2.3.1) $S'; G1, s(G2); - \vdash_{-q} s(v_q) : s(t')$

Or,

(B2.3') $S'; G1, s(G2); - \vdash_{-q} s(v_q) : t'$

Finally, we require showing that for all v_q and s of the form $(x \rightarrow v_p)$, if v_q is a q -value, then $s(v_q)$ is a r -value, for some r . This is straightforward.

For the conclusion, we apply (K-Dep) with (B2.2') and (B2.3') in the premises.

Case (K-Refine):

Similar to (K-Arrow), since Phi is simply a type.

Cases of (C2)

Case (S-Eq):

By the mutual induction hypothesis on substitution for the type-equivalence judgement.

Case (S-Trans):

By the induction hypothesis on each premise.

Case (S-Refine)

$G=G1, x:(p, tx), G2$	(C2.1)
$S; G \vdash t :: *$	(C2.2)
$S; G \vdash t <: t'$	(C2.3)
$G' = \text{refinements}(G, y:(p, t))$	(C2.4)
$S; G, y:(p, t), G'; - \vdash_{-q} v_q : \text{proof } \text{Phi}'$	(C2.5)
----- [S-Refine]	
$S; G \vdash t <: \{y:t' \mid \text{Phi}'\}$	

Where $G' = \text{refinements}(G1), \text{refinements}(x:(p, tx)), \text{refinements}(G2, y:t)$

Our goal is to show that

$S; G1, s(G2) \vdash s(t) <: \{y:s(t') \mid s(\text{Phi}')\}$

From the mutual induction hypothesis, we have

(C2.2') $S; G1, s(G2) \vdash s(t) :: *$

From the induction hypothesis we have

(C2.3') $S; G1, s(G2) \vdash s(t) <: s(t')$

From the definition of refinements, we have

$$\begin{aligned} \text{refinements}(G1, s(G2)) &= \text{refinements}(G1), \text{refinements}(s(G2)) \\ &= \text{refinements}(G1), s(\text{refinements}(G2)) \end{aligned}$$

We consider two subcases. Here, we write $\text{ref}(G)$ as a shorthand for $\text{refinements}(G)$

Subcase $(tx \langle \rangle \{z:tx' \mid \text{Phi}\})$:

In this case, $\text{ref}(x:(p, tx)) = .$ and we have C2.5

C2.5 $S;G1, x:(p,tx), G2, y:(p, t), \text{ref}(G1), \text{ref}(G2); - \mid - v_q : \text{proof Phi}'$

We use the mutual induction hypothesis to show that

C2.5.1
 $S;G1, s(G2), s(\text{ref}(G1)), s(\text{ref}(G2)); - \mid - s(v_q) : \text{proof s(Phi}'$

From the well-formedness of $G1$, we have that x not in $\text{FV}(G1)$ and $s(\text{ref}(G1)) = \text{ref}(G1)$, and so

C2.5
 $S;G1, s(G2), \text{ref}(G1), s(\text{ref}(G2)); - \mid - s(v_q) : \text{proof s(Phi}'$

We can now construct the goal using C2.2', C2.3' and C2.5'

Subcase $(tx = \{z:tx' \mid \text{Phi}\})$:

In this case, $\text{ref}(x:(p, tx)) = x':\text{proof Phi}$ and we have

C2.5 $S;G1,x:(p,tx),G2,y:(p,t),\text{ref}(G1), x':\text{proof Phi},\text{ref}(G2); - \mid - v_q : \text{proof Phi}'$

We first apply (T-Sub) with (S-Sub) to (C3) to produce

(C3') $S;G1;- \mid -_p v_p : tx'$

And apply the mutual induction hypothesis to construct

C2.5.1
 $S;G1,s(G2),s(\text{ref}(G1)), s(x':(p, \text{proof Phi})),s(\text{ref}(G2)); - \mid - s(v_q) : \text{proof s(Phi}'$

Next, we appeal to Lemma (canonical forms) for values of refined types and weakening to construct

PPhi $S;G1, s(G2), \text{ref}(G1);- \mid -_p v_p' : \text{proof Phi}$

And apply the mutual induction hypothesis to C2.5.1 and PPhi to construct

C2.5.2
 $S;G1,s(G2),s(\text{ref}(G1)), s'(s(\text{ref}(G2))); - \mid - s'(s(v_q)) : \text{proof s'(s(Phi}'$

where $s' = (x', v_p')$

However, since x' is fresh (from definition of ref), we have the composition of s' with s to be the identity. ($s'.s = s$)

So, we have:

C2.5' $S;G1,s(G2),s(\text{ref}(G1)), s(\text{ref}(G2)); - \mid - s(v_q) : \text{proof s(Phi}'$

And we can construct the goal using C2.2', C2.3' and C2.5'.

Case (S-Sub):

Immediate.

Cases of (D2)

Case (E-Id, E-App, E-Dep, E-D):

Straightforward from induction hypotheses.

Case (E-Subst):

We consider two sub-cases depending on whether the used match assumption appears in G1 or G2.

Sub-case (match assumption in G1):

----- [E-Subst]
S; G1', v=v', x:(p,tx), G2 |- v ~ v'

From the well-formedness of the environment, the x not in FreeVars(e, e'). Thus s(v) = v, s(v')=v' and the goal follows.

(Goal) S;G1',v=v', s(G2) |- s(v) ~ s(v)'

Sub-case (match assumption in G2):

----- [E-Subst]
S; G1,x:(p,tx),v=v' G2 G' |- v ~ v'

The goal follows immediately.

(Goal) S;G1',s(v)=s(v'), s(G2) |- s(v) ~ s(v)'

□

Lemma 11 (Type substitution). Lemma (Type substitution for typing judgment):

For all S G1 G2 X e t t1 k 'a p,
(A1) S;G1,'a::k,G2;X well-formed
(A2) S;G1,'a::k,G2;X |-_p e : t
(A3) S;G1 |- t' :: k
(A4) substitution s=[t1/'a]
=> S;G1,s(G2);X |-_p s(e) : s(t)

Lemma (Type substitution for kinding judgment):

For all S G1 G2 t1 k1 t2 k2 'a p,
(B1) S;G1,'a::k2,G2;- well-formed
(B2) S;G1,'a::k2,G2 |- t :: k1
(B3) S;G1 |- t2 :: k2
(B4) substitution s=[t2/'a]
=> S;G1,s(G2) |-_p s(t) :: k1

Lemma (Type substitution for type conversion):

For all S G1 G2 t t' t2 k2 'a p,
(B1) S;G1,'a::k2,G2;- well-formed
(B2) S;G1,'a::k2,G2 |- t <: t'
(B3) S;G1 |- t2 :: k2
(B4) substitution s=[t2/'a]
=> S;G1,s(G2) |- s(t) <: s(t')

Lemma (Type substitution for type equivalence):

For all S G1 G2 t t' t2 k2 'a p,
(B1) S;G1,'a::k2,G2;- well-formed
(B2) S;G1,'a::k2,G2 |- t ~ t'
(B3) S;G1 |- t2 :: k2
(B4) substitution s=[t2/'a]

$\Rightarrow S;G1, s(G2) \vdash s(t) \sim s(t')$

Proof. Straightforward mutual induction on the structure of (A2), (B2), (C2), (D2). □

Lemma 12 (Strengthening for inaccessible affine assumptions). For all $S G G' X x p t e q t'$,

- (A1) $S;G, x:(p, t), G'; X \vdash_{-q} e : t'$
 - (A2) $x \text{ not in } X$
 - (A3) $S;G; - \vdash t :: A$
- $\Rightarrow S;G, G'; X \vdash_{-q} e : t'$

Proof. Observation 1: $S;G, G'; X$ is well-formed, since
 $-- x \text{ not in } FV(G)$, by well-formedness of $S;G; -$
 $-- x \text{ not in } FV(G')$, since from (A2), x is affine
 And, from well-formedness of kinds,
 forall $S, G, t, k. S;G \vdash t :: k \Rightarrow x \notin FV(t)$

By induction on the structure of (A1), noting that (T-VarA) requires x in X . □

Corollary 13 (Strengthening for affine assumptions in kinding). For all $S G G' t t' k$,

- (A1) $S;G, x:(p, t), G' \vdash t' :: k$
 - (A2) $S;G \vdash t :: A$
- $\Rightarrow S;G, G'; \vdash t' :: k$

Lemma 14 (Destruction of affine assumption). For all $S M p e t M' x$,

- (A1) $S;G(M); X \vdash_{-p} e : t$
 - (A2) $(M, e) \sim_p \rightarrow (M', e')$
 - (A3) $x \text{ in } \text{dom}(M) \wedge x \text{ not in } \text{dom}(M')$
- $\Rightarrow x \text{ in } X$

Lemma 15 (Construction of affine assumption). For all $S M p e t M' x$,

- (A1) $S;G(M); X \vdash_{-p} e : t$
 - (A2) $(M, e) \sim_p \rightarrow (M', e')$
 - (A3) $x \text{ in } \text{dom}(M') \wedge x \text{ not in } \text{dom}(M)$
- $\Rightarrow FV(e') \subseteq X, x$

Proof. Simple induction on the structure of (A2), noting from well-formed memory that values in the store are always closed. □

Lemma 16 (Redundant match assumptions). For all $S M X p e t v$,

- (A1) $S;G, v=v, G'; X \vdash_{-p} e : t$
- $\Rightarrow S;G, G'; X \vdash_{-p} e : t$

Proof. Straightforward by noting that every application of (E-Subst) can be replaced by (E-Id). □

Lemma 17 (Unrolling fixed points). For all $S G G' f t t' p e e'$,

- (A1) $S;G, f:(p, \text{fix}, t), G'; X \vdash_{-p} e : t'$
 - (A2) $S;G; . \vdash_{-p} e' : t$
- $\Rightarrow S;G, f:(p, \text{fix}, t), G'; X \vdash_{-p} e[e'/f] : t'$

Proof. A simpler version of the substitution lemma, without need for mutual recursion at the type-level, since (K-Dep) rules out f from being used in types. □

Theorem 18 (Subject reduction). For all $S M M' X e t p$,

- (A0) $S;G(M); X$ well-formed
 - (A1) $S;G(M); X \vdash_{-p} e : t$
 - (A2) $(M, e) \sim_p \rightarrow (M', e')$
- \Rightarrow exists X' . $S;G(M'); X' \vdash_{-p} e' : t \wedge$
 $X' = X \cup (\text{dom } M' \setminus \text{dom } M)$ if $\text{dom } M' \supseteq \text{dom } M$
 $X' = X \setminus (\text{dom } M \setminus \text{dom } M')$ otherwise

Proof. By induction on the structure of the typing derivation (A1).

Cases (T-Sub, T-Drop):

Induction hypothesis.

Case (T-Var):

Impossible, from the definition of well-formed memory, x is of kind A .

Case (T-VarA):

x steps using (R-Destruct) to v_q and, from the well-formedness of memory M , we have the result.

Case (T-Val):

v has kind $*$ and is thus irreducible. I.e, we cannot step via (R-Alloc).

Case (T-Data, T-Abs, T-Univ):

If v has kind $*$ then it is irreducible.

Otherwise, a step via (R-Alloc) is possible producing a fresh location x and memory $M'=(M, (x,v))$. And x is well-typed in $G(M')$ using T-VarA.

Case (T-Fix):

```
unrefined t
S; G |- t :: *           (A1.1)
S; G, f:(p,t); - |-_p v_p : t   (A1.2)
----- [T-Fix]
S; G; - |-_p fix f:t.v_p : t
```

Inversion of (A2) gives an application of (R-Fix):

```
----- [R-Fix]
fix f:t.v_p ~p~> (v_p[fix f:t.v_p / f])
```

From an application of the unrolling lemma, we get

```
S;G;- |-_p v_p [fix f:t.v_p / f] : s(t)
```

Incidentally, from well-kinding of typing, we have that f not in $\text{dom}(t)$. Thus, $s(t) = t$

Case (T-App):

```
S; G; X1 |-_p v1 : ?(x:t1) -> t2   (A1.1)
S; G; X2 |-_p v2 : t1             (A1.2)
S; G |- t2[v2/x] :: k             (A1.3)
----- [T-App]
S; G; X1,X2 |-_p v1 v2 : t2[v2/x]
```

We first use the canonical forms lemma to establish either

Sub-case ($v1 = \backslash x:t1.e$): Then, inversion of (A2) gives us an application of (R-Beta).

```
----- [R-Beta]
\ x:t1.e v2 ~p~> e[v2/x]
```

From an inversion of (A1.1), we get an application of (T-Abs) with

$$(A1.1.1) \quad S;G,x:t1;X,x \vdash_{-p} e : t2$$

Now, using (A1.1.1) and (A1.2), we apply the substitution lemma, to derive the goal

$$S;G;X \vdash_{-p} e[v2/x] : t2[v2/x]$$

Sub-case ($v1=y$, y in $\text{dom}(M)$): Then, inversion of (A2) gives us an application of (R-Destruct)

$$\text{----- [R-Destruct]} \\ (M, (y, v'), M'), (y \ v2) \rightsquigarrow (M, M'), v' \ v2$$

We appeal to well-formed environments to construct:

$$(A1.1') \quad S;G(M, M'); X1 \setminus \{y\} \vdash_{-p} v' : !(x:t1) \rightarrow t2$$

From splitting of affine assumptions, we have y not in $X2$.

Applying Lemma (strengthening for inaccessible affine assumptions) to (A1.2) we construct

$$(A1.2') \quad S;G(M, M'); X2 \vdash_{-p} v2 : t1$$

Similarly, we use strengthening on affine assumptions in the kinding judgment to construct

$$(A1.3') \quad S;G(M, M') \vdash_{-} t2[v2/x] :: k$$

For the conclusion, we use (T-App) with (A1.1') and (A1.2') and (A1.3').

Case (T-TApp):

Similar to (T-App) using canonical forms to split into two sub-cases, appealing to the type-substitution lemma in the first case. In the second case, we appeal to well-formed memories for the first premise, and observe that affine assumptions are unused in the kinding judgment for the type in the second premise.

Case (T-Bracket):

By inversion on (A2), we have one of several cases.

---Subcase (A2 is R-Br): Straightforward from induction hypothesis.

---Subcase (A2 is R-Strip): Use the premise of A1 for the conclusion.

---Subcase (A2 is R-Nest): Use T-Bracket with the nested premise of (A1) for the conclusion.

---Subcase (A2 is R-Ext1):

From inversion of (A1), we get:

$$S;G,x:(q,t);X,x \vdash_{-q} e : t' \quad (A1.1)$$

By weakening (A1.1), we get

$$S;G,y:(p,t),x:(q,t);X,y,x \vdash_{-q} e : t' \quad (A1.2)$$

We have that

$$S;G,y:(p,t);X,y \vdash_{-q} \langle y \rangle_p : t \quad (A1.3)$$

And $\langle y \rangle_p$ is a q-value.

So, from the substitution lemma, we get

$S; G, y: (p, t); X, y \vdash_{-q} e \llbracket \langle y \rangle_p / x \rrbracket : t' \quad (G1.1)$

For the conclusion, we apply (T-Fun) with (G1.1) in the premise.

---Subcase (A2 is E-ExtL):

We use (T-Tabs), with (T-Bracket) with the nested premise of (A1) for the conclusion.

Case (T-Match):

Sub-case (scrutinee v_p is a variable): We step using (R-Destruct). The conclusion follows from well-formedness of memories, splitting of affine assumptions, and strengthening of unused affine assumptions, as in (T-App).

Sub-case (v_p is a non-variable): Inversion of (A2) gives us two further sub-cases.

Sub-sub-case (A2 is R-Else):

If a step is taken to the false branch via (E-Match), then, the conclusion follows using the last premise of (A1), with (T-Drop) to introduce unused affine assumptions, if any.

Sub-sub-case (A2 is R-Then):

We have unify ($v_p, D \ t1..tn \ x1..xn$) and inversion gives us

$v_p = D \ t1..tn \ v1..vn$

with the substitution $\Theta = ([v1/x1], \dots, [vn/xn])$

The fourth premise of A1 is:

(A1.4) $S; G, x1:t1'..xn:tn', xi=vi...xk=vk, (D \ t1..tn \ x1..xn = D \ t1..tn \ v1..vn); X \vdash e1 : t$

where the equalities ($xi=vi...xk=vk$) is computed by unifying the types of v_p and the pattern.

Then, from the repeated application of the substitution lemma, we get

$S; G, vi=vi, \dots, vk=vk, (D \ t1..tn \ v1..vn = D \ t1..tn \ v1..vn); X \vdash e1 : \Theta(t),$

From repeated use of the redundant match assumptions lemma, we arrive at

$S; G; X \vdash e1 : \Theta(t)$

Finally, from well-kinded typings lemma applied to the last premise, we get that $\Theta(t) = t$, since $FV(t)$ does not include any of the pattern variables.

Case (T-Let):

----- [T-Let]
 $S; G; X \vdash_p \text{let } x = e1 \text{ in } e2 : t$

Sub-case ($e1$ is a value):

Inversion of (A2) gives either (R-Alloc) or (R-Let).

Sub-sub-case (R-Alloc): The resulting memory is well-formed from the first premise of (R-Alloc). The conclusion follows immediately.

Sub-sub-case (R-Let): We apply the substitution lemma to e_2 .

Sub-case (e_1 is a non-value):

Inversion of (A2) gives (R-Ctx). Conclusion follows from the induction hypothesis.

Case (T-Val):

If this is a non-variable, then it is irreducible.

If it is a variable, then a reduction via (R-Destruct) is possible. The conclusion follows from well-formed memories.

□

C Value abstraction for FINE

Theorem 19 (Value abstraction). forall $S \ e \ t \ x \ p \ q \ tx \ v1_p \ v2_p$.

(A0) $S; x: (p, tx); x$ well-formed

(A1) $S; x: (p, tx); x \vdash_{-q} e : t$

(A2) $q \text{ not } \geq p$ and e is a non-value free of $\langle \cdot \rangle_r$ brackets, forall r where $r \geq p$, except $\langle x \rangle_p$

(A3) forall i . $S; -; - \vdash_{-p} vi_p : tx$

(A4) $(\cdot, e[v1_p/x]) \sim_q (\cdot, e_1)$

\Rightarrow exists e_2 . $e_2[v1_p/x] = e_1 \wedge (\cdot, e[v2_p/x]) \sim_q (\cdot, e_2[v2_p/x])$

Proof. Note the restriction to the pure fragment. We write $e \sim_q e'$ as a shorthand for $(\cdot, e) \sim_q (\cdot, e')$.

Proof by induction on the structure of (A4).

In each case, we appeal to well-typing (A1) to ensure that $e[v1_p/x] \langle \cdot \rangle v1_p$, since by (T-Var), x can only be used when enclosed in p -brackets. Likewise, for $v2_p$.

Case (R-Br):

$$\frac{e[v1_p/x] \sim_r e'}{\langle e \rangle_r [v1_p/x] \sim_q \langle e' \rangle_r} \quad [\text{R-Br}]$$

e is free of $\langle \cdot \rangle_r$ brackets, for all $r \geq p$. So, either

Sub-case 1: ($r \text{ not } \geq p$): From the premise, and from the IH, we get exists e_2

$$(\text{FromIH}) \ e_2[v1_p/x] = e' \quad \text{and} \quad e[v2_p/x] \sim_r e_2[v2_p/x]$$

For the conclusion, we construct $\langle e_2 \rangle_r$ as the witness and use R-Br with FromIH in the premise.

Sub-case 2: ($r=p$ and $e=x$):

Impossible, since $e[v1_p/x]=v1_p$, and $v1_p$ is a p -value and is irreducible.

Case (R-Ext1):

By the definition of substitution, after a suitable alpha-conversion, we have

$$\langle \lambda y:t.e \rangle_{q'} [v1_p/x] = \langle \lambda y:t. e[v1_p/x] \rangle_{q'}$$

$$\frac{\text{fresh } z}{\langle \lambda y:t. e[v1_p/x] \rangle_{q'} \sim_q \langle \lambda z:t. \langle e[v1_p/x] \rangle_{[z]_q/y} \rangle_{q'}} \quad [\text{R-Ext1}]$$

For the conclusion, choose $e_2 = \langle \lambda z:t. \langle e \rangle_{[z]_q/y} \rangle_{q'}$

We have $e_2[v1_p/x] = \langle \lambda z:t. \langle e \rangle_{[z]_q/y} [v1_p/x] \rangle_{q'}$.

From $z \langle x \rangle$, the substitutions commute, yielding the left conjunct of the goal.

For the right conjunct, we use R-Ext1

$$\langle \lambda y:t. e[v2_p/x] \rangle_{q'} \sim q \sim \lambda z:t. \langle e [v2_p/x] \rangle_{[z]_q/y} \rangle_{q'}$$

Again, observing that $z \langle x \rangle$, we have commuting substitutions and have that

$$\begin{aligned} e2[v2_p/x] &= (\lambda z:t. \langle e [z]_q/y \rangle_{q'}) [v2_p/x] \\ &= (\lambda z:t. \langle e [z]_q/y \rangle_{[v2_p/x]_q'}) \\ &= \lambda z:t. \langle e [v2_p/x] \rangle_{[z]_q/y} \rangle_{q'} \quad (\text{by commutation}) \end{aligned}$$

Case (R-ExtL): Similar, but simpler since the type variable 'a is not wrapped and no commutation is involved.

Case (R-Strip):

$$\begin{array}{l} \text{----- (E-Strip)} \\ \langle v_q \rangle_q [v1_p/x] \sim q \sim v_q[v1_p/x] \end{array}$$

Sub-case $q=p$.

From our assumption of p-bracket freedom, we have that v_q must be x . Thus, we have

$$\langle x \rangle_p [v1_p/x] \sim p \sim v1_p$$

So, we pick $e2=x$, and the conclusion is immediate.

Sub-case $q \langle p \rangle$.

Pick $e2 = v_q$.

Case (R-Nest): As in the previous case.

Case (R-Ctx, R-Let): Straightforward from IH and definitions of substitution.

Case (R-Beta):

From the definition of substitution, and alpha-converting the left-subterm to ensure that y is distinct, we get:

$$\begin{array}{l} \text{----- [R-Beta]} \\ (\lambda y:t. e v_q) [v1_p/x] \sim q \sim (e[v1_p/x]) [v_q [v1_p/x]/y] \end{array}$$

For the conclusion, we pick $e2 = e [v_q/y]$ and

For the left conjunct:

$$\begin{aligned} e2[v1_p/x] &= e [v_q/y] [v1_p/x] \\ &= e [v1_p/x] [v_q[v1_p/x] / y] \quad (\text{by def of subst. and } y \text{ not in } FV(v1_p)) \end{aligned}$$

For the right conjunct, first, we have from R-Beta:

$$(\lambda y:t. e v_q) [v2_p/x] \sim q \sim e [v2_p/x] [v_q[v2_p/x] / y]$$

Then, we observe:

$$\begin{aligned} e2[v2_p/x] &= e [v_q/y] [v2_p/x] \\ &= e [v2_p/x] [v_q[v2_p/x] / y] \quad (\text{by def of subst. and } y \text{ not in } FV(v2_p)) \end{aligned}$$

Which suffices for the goal.

Case (R-TBeta, R-Fix):

Similar to the previous case.

Case (R-Then):

$$\frac{\begin{array}{l} v_q \langle x \rangle y \\ \text{unify } (v_q[v1_p/x], D \ t1..tn \ x1..xm \ [v1_p/x]) = \text{Theta} \end{array}}{\text{match } v_q[v1_p/x] \text{ with } D \ t1..tn \ x1..xm \ [v1_p/x] \rightarrow e_t \ [v1_p/x] \text{ else } e_f \ [v1_p/x] \ \sim q \sim} \quad \text{[R-Then]}$$

$$\text{Theta}(e_t[v1_p/x])$$

Inverting the second premise of (R-Then), we get

$$v_q[v1_p/x] = D \ t1'[v1_p/x] \ .. \ tn'[v1_p/x] \ v1[v1_p/x] \ \dots \ vm[v1_p/x]$$

$$\begin{array}{l} t1[v1_p/x] = t1'[v1_p/x] \ \dots \\ tn[v1_p/x] = tn'[v1_p/x] \end{array}$$

$$\text{Theta} = [v1[v1_p/x]/x1, \ \dots \ vm[v1_p/x]/xm]$$

and

$$\text{match } \dots \ \sim q \sim \ \text{Theta} \ (e_t[v1_p/x])$$

(Note that v_q cannot be $\langle x \rangle_p$, since bracketed terms do not match any pattern.)

Factoring out the substitution $[v1_p/x]$, we get

$$\text{Theta}' = [v1/x1 \ \dots \ vm/xm] \ \text{and} \ \text{match } \dots \ \sim q \sim \ \text{Theta}' \ (e_t) \ [v1_p/x]$$

For the conclusion, we pick $e2 = \text{Theta}' \ (e_t)$

For the left side of the conjunct in the goal, we need to show

$$e2[v1_p/x] = \text{Theta}' \ (e_t) \ [v1_p/x], \ \text{which follows by construction.}$$

For the right side of the conjunct, we need to show

$$\begin{array}{l} \text{match } v_q[v2_p/x] \text{ with } D \ t1..tn \ x1..xm \ [v2_p/x] \ \text{then } e_t[v2_p/x] \ \text{else } e_f[v2_p/x] \\ \sim q \sim \\ e2[v2_p/x] \end{array}$$

To apply R-Then, we need to derive:

$$(G) \quad \begin{array}{l} \text{unify } (D \ t1'[v2_p/x] \ \dots \ tn'[v2_p/x] \ v1[v2_p/x] \ \dots \ vm[v2_p/x], \\ \quad D \ t1[v2_p/x] \ \dots \ tn[v2_p/x] \ x1 \ \dots \ xm) : \text{Theta}'[v2_p/x] \end{array}$$

For this, we must first show that

$$(G1) \quad \text{forall } i. \ ti[v2_p/x] = ti'[v2_p/x]$$

To show (G1), we appeal first to the well-typing premise (A1), where, by inversion we have an application of (T-Match).

$$\begin{array}{ll} S; x: (p, tx); x \vdash v_q : t & (A1.1) \\ S; x: (p, tx), x1..xm; x, x1..xm \vdash D \ t1..tn \ x1..xm : t' & (A1.2) \\ \text{unify}(t, t') : \text{Theta}'' & (A1.3) \\ \dots & \end{array}$$

$$\text{----- [T-Match (A1)]}$$

$$S; x: (p, tx); x \vdash_q \text{match } v_q \text{ with } D \ t1..tn \ x1..xm \ \text{then } e_t \ \text{else } e_f$$

Next, we appeal to well-formedness of data constructors to note that each of the $t_1..t_n$ appear as indices on t' , i.e.,

(G1.1) t' is of the form $T t_1 .. t_n v_1'..v_l'$, for some T and v_1', \dots, v_l'

Now, inverting (A1.3) and using (G1.1) we find that

(G1.2.1) t is of the form $T t_1 .. t_n v_1''..v_l''$, for some v_1'', \dots, v_l''

But, we know that $v_q = D t_1' .. t_n' v_1 .. v_m$ is well-typed at type t and from the well-formedness of data constructors, we have that

(G1.2.2) t is of the form $T t_1' .. t_n' v_1'' .. v_l''$

Hence,

(G1.3) $t_1=t_1', \dots, t_n=t_n'$,

which is sufficient for establishing (G1).

(I.e., in well-typed programs, the branch taken in a pattern match does not depend on the type-arguments of the pattern)

Rewriting G using (G1.3), we now have as our goal:

(G) $\text{unify } (D t_1[v_2_p/x] .. t_n[v_2_p/x] v_1[v_2_p/x] .. v_m[v_2_p/x],$
 $D t_1[v_2_p/x] .. t_n[v_2_p/x] x_1 .. x_m) : \text{Theta}'[v_2_p/x]$

Refactoring substitutions, we have

(G) $\text{unify } (D t_1 .. t_n v_1 .. v_m)[v_2_p/x], D t_1..t_n x_1..x_m [v_2_p/x]) : \text{Theta}'[v_2_p/x]$

Which is immediate from the definition of unify .

Case (R-Else):

$$\frac{\begin{array}{l} v_q \langle y \\ \text{forall Theta. not } (\text{unify } (v_q[v_1_p/x], D t_1..t_n x_1..x_m [v_1_p/x]) : \text{Theta}) \end{array}}{\text{match } v_q[v_1_p/x] \text{ with } D t_1..t_n x_1..x_m [v_1_p/x] \rightarrow e_t [v_1_p/x] \text{ else } e_f [v_1_p/x]} \quad \begin{array}{l} \text{(A4.1)} \\ \text{[R-Else]} \end{array}$$

$\sim q \rangle$
 $e_f[v_1_p/x]$

For the conclusion, we pick $e_2 = e_f$, and proceed by inversion on v_q to show that
 $\text{forall Theta. not } (\text{unify } (v_q[v_2_p/x], D t_1..t_n x_1..x_m [v_2_p/x]) : \text{Theta})$

Sub-case ($v_q = x$): impossible, by (A1)

Sub-case ($v_q = \langle v' \rangle_p$): Then $v_q[v_2_p/x] = \langle v' [v_2_p/x] \rangle_p$ and bracketed terms match no patterns.

Sub-case ($v_q = \backslash x:t.e, / \backslash a::k.e$): Abstractions do not match any patterns.

Sub-case ($v_q = D' t_1'..t_n' v_1..v_n$):

We have from (A4.1) that $\text{not}(\text{unify } (D' t_1'..t_n' v_1 ..v_m[v_1_p/x], D t_1..t_n x_1..x_m[v_1_p/x])$.

Suppose, for contradiction, we have

(ForContra) $\text{unify } (D' t_1'..t_n' v_1 ..v_m[v_2_p/x], D t_1..t_n x_1..x_m[v_2_p/x])$

Then, we have $D=D'$, and following an argument identical to the R-Then case, relying on well-typed programs, we have $t_1'..t_n' = t_1..t_n$.

Rewriting (A4.1) using these equalities, we reach

```
not(unify (D t1..tn v1..vm[v1_p/x], D t1..tn x1..xm[v1_p/x]))
```

But also, from the definition of unify, we can construct:a

```
unify (D t1..tn v1..vm[v1_p/x], D t1..tn x1..xm[v1_p/x]) : [v1/x1..vm/xm][v1_p/x]
```

A contradiction, as required.

□

References

- [1] J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffei. Refinement types for secure implementations. In *CSF*, 2008.
- [2] Y. Bertot and P. Castéran. *Coq'Art: Interactive Theorem Proving and Program Development*. Springer Verlag, 2004.
- [3] J. Borgstrom, A. Gordon, and R. Pucella. Roles, stacks, histories: A triple for hoare. Technical Report MSR-TR-2009-97, Microsoft Research, 2009.
- [4] A. Chlipala, G. Malecha, G. Morrisett, A. Shinnar, and R. Wisnesky. Effective interactive proofs for higher-order imperative programs. In *ICFP*, 2009.
- [5] S. Chong, A. C. Myers, N. Nystrom, L. Zheng, and S. Zdancewic. Jif: Java + information flow, July 2006. Software release.
- [6] L. de Moura and N. Björner. Z3: An efficient SMT solver. In *TACAS*, 2008.
- [7] D. J. Dougherty, K. Fidler, and S. Krishnamurthi. Specifying and reasoning about dynamic access-control policies. In *LNCS*, 2006.
- [8] ECMA. Standard ECMA-335: Common language infrastructure, 2006.
- [9] C. Flanagan. Hybrid type checking. In *POPL*, 2006.
- [10] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *PLDI*. 1993.
- [11] D. Grossman, G. Morrisett, and S. Zdancewic. Syntactic type abstraction. *ACM TOPLAS*, 22(6), 2000.
- [12] D. Jackson. Alloy: a lightweight object modelling notation. *TOSEM*, 11(2), 2002.
- [13] L. Jia, J. Vaughan, K. Mazurak, J. Zhao, L. Zarko, J. Schorr, and S. Zdancewic. Aura: A programming language for authorization and audit. In *ICFP*, 2008.
- [14] S. Krishnamurthi, P. W. Hopkins, J. McCarthy, P. T. Graunke, G. Pettyjohn, and M. Felleisen. Implementation and use of the PLT Scheme web server. *HOSC*, 20(4), 2007.
- [15] H. M. Levy. *Capability-Based Computer Systems*. Butterworth-Heinemann, 1984.
- [16] C. McBride and J. McKinna. The view from the left. *JFP*, 14(1), 2004.
- [17] U. Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Chalmers Institute of Technology, 2007.
- [18] V. Simonet. FlowCaml in a nutshell. In G. Hutton, editor, *APPSEM-II*, pages 152–165, 2003.
- [19] A. Stump, M. Deters, A. Petcher, T. Schiller, and T. Simpson. Verified programming in Guru. In *PLPV*, 2008.
- [20] N. Swamy, J. Chen, and R. Chugh. Enforcing stateful authorization and information flow policies in Fine. Technical Report MSR-TR-2009-164, Microsoft Research, 2009.
- [21] N. Swamy, B. J. Corcoran, and M. Hicks. Fable: A language for enforcing user-defined security policies. In *S&P*, 2008.
- [22] N. Swamy and M. Hicks. Verified enforcement of stateful information release policies. In *PLAS*, 2008.
- [23] D. Syme, A. Granicz, and A. Cisternino. *Expert F#*. Apress, 2007.
- [24] P. Wadler. Linear types can change the world. In *Prog. Concepts and Methods*, 1990.
- [25] R. Wahbe, S. Lucco, T. E. Anderson, and S. L. Graham. Efficient software-based fault isolation. In *SOSP*, 1993.