

Practical Affine Types

Jesse A. Tov Riccardo Pucella

Northeastern University, Boston, Massachusetts, USA

{tov,riccardo}@ccs.neu.edu

Abstract

Alms is a general-purpose programming language that supports practical affine types. To offer the expressiveness of Girard’s linear logic while keeping the type system light and convenient, Alms uses expressive kinds that minimize notation while maximizing polymorphism between affine and unlimited types.

A key feature of Alms is the ability to introduce abstract affine types via ML-style signature ascription. In Alms, an interface can impose stiffer resource usage restrictions than the principal usage restrictions of its implementation. This form of sealing allows the type system to naturally and directly express a variety of resource management protocols from special-purpose type systems.

We present two pieces of evidence to demonstrate the validity of our design goals. First, we introduce a prototype implementation of Alms and discuss our experience programming in the language. Second, we establish the soundness of the core language. We also use the core model to prove a principal kinding theorem.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features

General Terms Languages

Keywords Affine types, linear logic, type systems, polymorphism, modules

1. A Practical Affine Type System

Alms is a practical, general-purpose programming language with affine types. *Affine types* enforce the discipline that some values are not used more than once, which in Alms makes it easy to define new, resource-aware abstractions. *General-purpose* means that Alms offers a full complement of modern language features suitable for writing a wide range of programs. *Practical* means that Alms is neither vaporware nor a minimal calculus—it is possible to download Alms today and try it out.

Rationale. Resource-aware type systems divide into two camps: foundational calculi hewing closely to linear logic, and implementations of special-purpose type systems designed to solve special problems. We argue that a general, practical type system based on Girard’s linear logic (1987) can naturally and directly express many of the special cases, such as region-based memory management, aliasing control, session types, and typestate. To this end, the language must satisfy several desiderata:

```
module type RW_LOCK = sig
  type ( $\alpha, \beta$ ) array — array of  $\alpha$  identified by  $\beta$ 
  type excl           — exclusive access
  type shared        — shared access
  type  $\beta@_\gamma : \mathbf{A}$  — grants  $\gamma$ -level access to array  $\beta$ 
  val new            : int  $\rightarrow \alpha \rightarrow \exists \beta. (\alpha, \beta)$  array
  val acquireW      : ( $\alpha, \beta$ ) array  $\rightarrow \beta@_{\text{excl}}$ 
  val acquireR      : ( $\alpha, \beta$ ) array  $\rightarrow \beta@_{\text{shared}}$ 
  val release       : ( $\alpha, \beta$ ) array  $\rightarrow \beta@_\gamma \rightarrow \text{unit}$ 
  val set           : ( $\alpha, \beta$ ) array  $\rightarrow \text{int} \rightarrow \alpha \rightarrow \beta@_{\text{excl}} \rightarrow \beta@_{\text{excl}}$ 
  val get           : ( $\alpha, \beta$ ) array  $\rightarrow \text{int} \rightarrow \beta@_\gamma \rightarrow \alpha \times \beta@_\gamma$ 
end
```

Figure 1. An interface for reader-writer locks (§2)

Convenience. Unlimited values are the common case, so working with them is as smooth as in a conventional language.

Expressiveness. A wide variety of resource-aware type systems appear naturally as idioms.

Familiarity. It is easy to use and understand.

Pragmatics. It provides the trappings of a modern, high-level programming language, such as algebraic datatypes, pattern matching, exceptions, concurrency, and modules.

Soundness. It has a clear theoretical foundation.

We show that Alms meets these criteria.

Alms employs a dependent kind system to determine whether a particular type is affine or unlimited and to support polymorphism over affine and unlimited types. This approach may sound complicated, but in practice it is no stranger or harder to understand than the type systems of other functional programming languages.

Affine types, a weakening of linear types, forbid duplication of some values; unlike with linear types, all values may be dropped. This flexibility is appropriate to a high-level, garbage-collected language, and it interacts better with other features such as exceptions.

Our Contributions. This paper introduces the programming language Alms, its implementation, and its basic theory:

- We describe the design of Alms, whose novel features include precise kinds for affine type constructors, and demonstrate how it expresses a variety of resource-aware idioms (§2).
- Our implementation is a usable, full-featured prototype in which we have written several thousand lines of code (§3).
- Alms rests on a firm theoretical basis. We provide a formal model (§4) and establish essential theoretical properties (§5).

Our implementation and full proofs are available at www.ccs.neu.edu/~tov/pubs/alms.

2. Learn Alms in 2.5 Pages

Alms is a typed, call-by-value, impure functional language with algebraic data types, pattern matching, reference cells, threads, exceptions, and modules with opaque signature ascription. Alms reads like Ocaml (Leroy et al. 2008) but is explicitly typed. In most cases local type inference renders explicit type instantiation unnecessary.

We introduce Alms through a series of examples. Consider a simple Alms function, *deposit*, that updates one element of an array by adding an integer:

```
let deposit (a: int Array.array) (acct: int) (amount: int) =
  Array.set a acct (Array.get a acct + amount)
```

This function has a race condition between the read and the write, so we may want to use a lock to enforce mutual exclusion:

```
let deposit (a: int Array.array) (acct: int)
  (amount: int) (lock: Lock.lock) =
  Lock.acquire lock;
  Array.set a acct (Array.get a acct + amount);
  Lock.release lock
```

Affine data. Locks can ensure mutual exclusion, but using them correctly is error-prone. A rather coarse alternative to ensure mutual exclusion is to forbid aliasing of the array altogether. If we have the only reference to an array then no other process can operate on it concurrently. In Alms, we do this by declaring an interface that includes a new, abstract array type:

```
module type AF_ARRAY = sig
  type  $\alpha$  array : A

  val new : int  $\rightarrow$   $\alpha \rightarrow \alpha$  array
  val set :  $\alpha$  array  $\rightarrow$  int  $\xrightarrow{A}$   $\alpha \xrightarrow{A}$   $\alpha$  array
  val get :  $\alpha$  array  $\rightarrow$  int  $\xrightarrow{A}$   $\alpha \times \alpha$  array
end

module AfArray : AF_ARRAY = struct ... end
```

The notation “: **A**” specifies that type α AfArray.array has kind **A**, as in *affine*, which means that any attempt to duplicate a reference to such an array is a type error. Two points about the types of AfArray.get and AfArray.set are worth noting:

- Each must return an array because the caller cannot reuse the reference to the array supplied as an argument.
- Type $\tau_1 \xrightarrow{A} \tau_2$ has kind **A**, which means that it may be used at most once.¹ This is necessary because reusing a function partially applied to an affine value would reuse that value.

We now rewrite *deposit* to use the AF_ARRAY interface:

```
let deposit (a: int AfArray.array) (acct: int) (amt: int) =
  let (balance, a) = AfArray.get a acct in
  AfArray.set a acct (balance + amt)
```

If we attempt to use an AfArray.array more than once without single-threading it, a type error results:

```
let deposit (a: int AfArray.array) (acct: int) (amt: int) =
  let (balance, _) = AfArray.get a acct in
  AfArray.set a acct (balance + amt)
```

Alms reports that the affine variable *a* is duplicated.

Implementing AfArray is just a matter of wrapping the primitive array type and operations, and sealing the module with an opaque signature ascription:

```
module AfArray : AF_ARRAY = struct
  type  $\alpha$  array =  $\alpha$  Array.array

  let new = Array.new
  let set (a:  $\alpha$  array) (ix: int) (v:  $\alpha$ ) = Array.set a ix v; a
  let get (a:  $\alpha$  array) (ix: int) = (Array.get a ix, a)
end
```

The original array type α Array.array has kind **U**, as in *unlimited*, because it places no limits on duplication. We can use it to represent an abstract type of kind **A**, however, because **U** is a subkind of **A**, and Alms’s kind subsumption rule allows assigning an abstract type a greater kind than that of its concrete representation. This is somewhat akin to Standard ML’s treatment of equality types (Milner et al. 1997) and Ocaml’s treatment of type constructor variance (Leroy et al. 2008). In SML, eqtype is subsumed by type, in that signature matching can abstract an equality type to a non-equality type but not vice versa.

We need not change *new* at all, and *get* and *set* are modified slightly to return the array as required by the interface.

Affine capabilities. The affine array interface is quite restrictive. Because it requires single-threading an array through the program, it cannot support operations that do not actually require exclusive access to the array. However, Alms supports creating a variety of abstractions to suit our needs. One way to increase our flexibility is to separate the reference to the array from the *capability* to read and write the array. Only the latter needs to be affine.

For example, we may prefer an interface that supports “dirty reads,” which do not require exclusive access but are not guaranteed to observe a consistent state:

```
module type CAP_ARRAY = sig
  type ( $\alpha, \beta$ ) array
  type  $\beta$  cap : A

  val new : int  $\rightarrow$   $\alpha \rightarrow \exists \beta. (\alpha, \beta)$  array  $\times$   $\beta$  cap
  val set : ( $\alpha, \beta$ ) array  $\rightarrow$  int  $\rightarrow$   $\alpha \rightarrow \beta$  cap  $\rightarrow$   $\beta$  cap
  val get : ( $\alpha, \beta$ ) array  $\rightarrow$  int  $\rightarrow$   $\beta$  cap  $\rightarrow$   $\alpha \times \beta$  cap
  val dirtyGet : ( $\alpha, \beta$ ) array  $\rightarrow$  int  $\rightarrow$   $\alpha$ 
end
```

In this signature, (α, β) array is now unlimited and β cap is affine. Type array’s second parameter, β , is a “stamp” used to tie it to its capability, which must have type β cap (where β matches). In particular, the type of *new* indicates that it returns an existential containing an array and a capability with matching stamps. The existential guarantees that the stamp on an array can only match the stamp on the capability created by the same call to *new*.

Operations *set* and *get* allow access to an array only when presented with the matching capability. This ensures that *set* and *get* have exclusive access with respect to other *sets* and *gets*. They no longer return the array, but they do need to return the capability. On the other hand, *dirtyGet* does not require a capability and should not return one.

For example, the CAP_ARRAY interface allows us to shuffle an array while simultaneously computing an approximate sum:

```
let shuffleAndDirtySum (a: ( $\alpha, \beta$ ) CapArray.array)
  (cap:  $\beta$  CapArray.cap) =
  let th1 = Thread.fork ( $\lambda \_ \rightarrow$  inPlaceShuffle a cap) in
  let th2 = Thread.fork ( $\lambda \_ \rightarrow$  dirtySumArray a) in
  (Thread.wait th1, Thread.wait th2)
```

To implement CAP_ARRAY, we need suitable representations for its two abstract types. We represent CAP_ARRAY’s arrays by the primitive array type, and capabilities by type unit, which is adequate because these capabilities have no run-time significance.

¹ It is tempting to call this an affine function, but standard terminology says that an affine function uses its *argument* at most once, whereas here we have the type of a function *itself* usable at most once. Whether an Alms function is affine is determined by the kind of the type of its formal parameter.

```

module A = Array
module CapArray : CAP_ARRAY = struct
  type (α, β) array = α A.array
  type β cap      = unit
  let new (size: int) (init: α) = (A.new size init, ())
  let set (a: α A.array) (ix: int) (v: α) _ = A.set a ix v
  let get (a: α A.array) (ix: int) _ = (A.get a ix, ())
  let dirtyGet = A.get
end

```

Type unit has kind **U**, but as in the previous example, we can abstract it to **A** to match the kind of β CapArray.cap. The implementation of the operations is in terms of the underlying array operations, with some shuffling to ignore capability arguments (in *set* and *get*) and to construct tuples containing () to represent the capability in the result (in *new* and *get*).²

Capabilities are values. Capabilities such as β CapArray.cap often represent the state of a resource, but in Alms they are also ordinary values. They may be stored in immutable or mutable data structures, packed into exceptions and thrown, or sent over communication channels like any other value. For example, suppose we would like a list of array capabilities. Lists are defined thus in the standard library:

```

type  $\hat{\alpha}$  list = Nil | Cons of  $\hat{\alpha} \times \hat{\alpha}$  list

```

The type variables we have seen until now could only be instantiated with unlimited types, but the diacritic on type variable $\hat{\alpha}$ indicates that $\hat{\alpha}$ may be instantiated to any type, whether affine or unlimited.

Whether a list should be treated as affine or unlimited depends on whether the contents of the list is affine or unlimited. Alms represents this fact by giving the list type constructor a dependent kind, where kind $\langle \hat{\alpha} \rangle$ denotes the kind of $\hat{\alpha}$:

```

 $\hat{\alpha}$  list :  $\langle \hat{\alpha} \rangle$  — list has kind  $\Pi \hat{\alpha}. \langle \hat{\alpha} \rangle$ 

```

That is, the kind of a list is the same as the kind of its element type: type int list has kind **U**, whereas β CapArray.cap list has kind **A**.

In general, the kind of a type is the least upper bound of the kinds of the types that occur directly in its representation. For example:

```

type ( $\hat{\alpha}, \hat{\beta}$ ) r =  $\hat{\alpha} \times \hat{\beta}$  —  $\langle \hat{\alpha} \rangle \sqcup \langle \hat{\beta} \rangle$ 
type ( $\hat{\alpha}, \hat{\beta}$ ) s = int  $\times \hat{\beta}$  —  $\langle \hat{\beta} \rangle$ 
type ( $\hat{\alpha}, \hat{\beta}$ ) t =  $T_1$  of  $\hat{\alpha}$  |  $T_2$  of ( $\hat{\beta}, \hat{\alpha}$ ) t —  $\langle \hat{\alpha} \rangle \sqcup \langle \hat{\beta} \rangle$ 
type ( $\hat{\alpha}, \hat{\beta}$ ) u =  $U_1$  |  $U_2$  of ( $\hat{\beta}, \hat{\alpha}$ ) u — U
type ( $\hat{\alpha}, \hat{\beta}$ ) v =  $\hat{\alpha} \times (\text{unit} \rightarrow \hat{\beta})$  —  $\langle \hat{\alpha} \rangle$ 
type ( $\hat{\alpha}, \hat{\beta}$ ) w =  $\hat{\alpha} \times (\text{unit} \xrightarrow{\langle \hat{\beta} \rangle} \text{unit})$  —  $\langle \hat{\alpha} \rangle \sqcup \langle \hat{\beta} \rangle$ 

```

Because both $\hat{\alpha}$ and $\hat{\beta}$ are part of the representation of ($\hat{\alpha}, \hat{\beta}$) r, it must be affine if either of its parameters is affine. On the other hand, the phantom parameter $\hat{\alpha}$ is not part of the representation of ($\hat{\alpha}, \hat{\beta}$) s, so that has kind $\langle \hat{\beta} \rangle$. The kinds of t and u are the least solutions to these inequalities:

```

 $\kappa((\hat{\alpha}, \hat{\beta}) t) \sqsupseteq \kappa(\hat{\alpha})$  —  $\kappa((\hat{\alpha}, \hat{\beta}) u) \sqsupseteq \mathbf{U}$ 
 $\kappa((\hat{\alpha}, \hat{\beta}) t) \sqsupseteq \kappa((\hat{\beta}, \hat{\alpha}) t)$  —  $\kappa((\hat{\alpha}, \hat{\beta}) u) \sqsupseteq \kappa((\hat{\beta}, \hat{\alpha}) u)$ 

```

The kind of each type must be at least as restrictive as the kinds of all of its alternatives.

For ($\hat{\alpha}, \hat{\beta}$) v, the kind of $\hat{\beta}$ does not appear because the domain and codomain of a function type are not part of the function's representation. Instead, function types have their kind in a superscript as in the definition of type w. We saw this in the *AfArray* example,

²If you think this unit shuffling is unnecessary, we agree (§7).

where the superscripts were all **A**. (When the kind is **U**, we often omit it.) We discuss the kinds of function types and their subtyping in more depth in §4.3.

More possibilities. The rules of Alms are flexible enough to express a wide variety of designs. For example, the ability to store capabilities in data structures allows us to create a more dynamic interface than the static capabilities of *CapArray*:

```

module type CAP_LOCK_ARRAY = sig
  include CAP_ARRAY
  val new : int  $\rightarrow \alpha \rightarrow \exists \beta. (\alpha, \beta)$  array
  val acquire : ( $\alpha, \beta$ ) array  $\rightarrow \beta$  cap
  val release : ( $\alpha, \beta$ ) array  $\rightarrow \beta$  cap  $\rightarrow$  unit
end

```

This signature changes the type of *new* to return an array (with unique tag β) but no capability. To operate on the array, one needs to request a capability using *acquire*. Subsequent attempts to acquire a capability for the same array block until the capability is released.

We implement CAP_LOCK_ARRAY in terms of *CapArray* without any privileged knowledge about the representation of its capabilities. The implementation relies on mvars, synchronized variables based Id's *M-structures* (Barth et al. 1991). An $\hat{\alpha}$ mvar may hold a value of type $\hat{\alpha}$ or it may be *empty*. While an mvar may contain an affine value, the mvar itself is always unlimited. This is safe because calling *take* on a non-empty mvar removes the value and returns it, while *take* on an empty mvar blocks until another thread *puts* a value in it.

To implement CAP_LOCK_ARRAY we now represent an array as a pair of the underlying (α, β) CapArray.array and an mvar to store its capability:

```

module CapLockArray : CAP_LOCK_ARRAY = struct
  module A = CapArray
  type ( $\alpha, \beta$ ) array = ( $\alpha, \beta$ ) A.array  $\times \beta$  cap MVar.mvar
  let new (size: int) (init: α) =
    let ( $\beta, a, cap$ ) = A.new size init in
      ( $a, MVar.new cap$ )
  let acquire ((_, mvar): ( $\alpha, \beta$ ) array) = MVar.take mvar
  let release ((_, mvar): ( $\alpha, \beta$ ) array) (cap: β cap) =
    MVar.put mvar cap
  let set ((a, _) : ( $\alpha, \beta$ ) array) = A.set a
  let get ((a, _) : ( $\alpha, \beta$ ) array) = A.get a
  ...
end

```

The *new* operation creates a new array-capability pair and stores the capability in an mvar. Operations *acquire* and *release* use the mvar component of the representation, while the old operations such as *set* must be lifted to project out the underlying CapArray.array.

There are many more possibilities. Figures 1 and 2 show two interfaces for reader-writer locks, which at any one time allow either exclusive read-write access or shared read-only access. Signature RW_LOCK (figure 1 on the first page) describes dynamic reader-writer locks. The signature declares nullary types *excl* and *shared* and an affine, binary type constructor ($\cdot @ \cdot$). Capabilities now have type $\beta @ \gamma$, where β ties the capability to a particular array and γ records whether the lock is exclusive or shared. Operation *set* requires an exclusive lock ($\beta @ \text{excl}$), but *get* allows γ to be shared or *excl*.

Signature FRACTIONAL (figure 2) describes static reader-writer locks based on fractional capabilities (Boyland 2003). As in the previous example, the capability type (β, γ) cap has a second parameter, which in this case represents a fraction of the whole

```

module type FRACTIONAL = sig
  type ( $\alpha, \beta$ ) array
  type 1
  type 2
  type  $\gamma/\delta$ 
  type ( $\beta, \gamma$ ) cap : A
  val new : int  $\rightarrow$   $\alpha \rightarrow \exists \beta. (\alpha, \beta)$  array  $\times$  ( $\beta, 1$ ) cap
  val split : ( $\beta, \gamma$ ) cap  $\rightarrow$  ( $\beta, \gamma/2$ ) cap  $\times$  ( $\beta, \gamma/2$ ) cap
  val join : ( $\beta, \gamma/2$ ) cap  $\times$  ( $\beta, \gamma/2$ ) cap  $\rightarrow$  ( $\beta, \gamma$ ) cap
  val set : ( $\alpha, \beta$ ) array  $\rightarrow$  int  $\rightarrow$   $\alpha \rightarrow$ 
    ( $\beta, 1$ ) cap  $\rightarrow$  ( $\beta, 1$ ) cap
  val get : ( $\alpha, \beta$ ) array  $\rightarrow$  int  $\rightarrow$ 
    ( $\beta, \gamma$ ) cap  $\rightarrow$   $\alpha \times$  ( $\beta, \gamma$ ) cap
end

```

Figure 2. Another interface for reader-writer locks

capability. The fraction is represented using type constructors 1, 2, and (\cdot / \cdot). A capability of type ($\beta, 1$) cap grants exclusive access to the array with tag β , while a fraction less than 1 such as $1/2$ or $1/2/2$ indicates shared access. There are operations *split*, which splits a capability whose fraction is γ into two capabilities of fraction $\gamma/2$, and *join*, which combines two $\gamma/2$ capabilities back into one γ capability. Again, *set* requires exclusive access but *get* does not.

Syntax matters. Given *CapLockArray* as defined above, we can rewrite *deposit* to take advantage of it:

```

open CapLockArray
let deposit (a: (int,  $\beta$ ) array) (acct: int) (amt: int) =
  let cap = acquire a in
  let (balance, cap) = get a acct cap in
  let cap = set a acct (balance + amt) cap in
  release a cap

```

While this gets the job done, the explicit threading of the capability can be inconvenient and hard to read. To address this, Alms provides preliminary support for an alternative syntax inspired by a proposal by Mazurak et al. (2010):

```

let deposit (a: (int,  $\beta$ ) array) (acct: int) (amt: int) =
  let !cap = acquire a in
  set a acct (get a acct cap + amt) cap;
  release a cap

```

The pattern `!cap` bound by `let` marks `cap` as an “imperative variable,” which means that within its scope, functions applied to `cap` are expected to return a pair of their real result and the new version of `cap`. Alms transforms this code into the explicitly-threaded version above. Currently this transformation happens before type checking, which means that it cannot compromise soundness but also cannot exploit type information.

3. Implementation and Experience

Our prototype, implemented in 16k lines of Haskell, is available at www.ccs.neu.edu/~tov/pubs/alms. Besides a usable interpreter, it includes all the example code from this paper and other Alms examples illustrating a variety of concepts:

- A capability-based interface to Berkeley Sockets ensures that the protocol to set up a socket is followed correctly. This library also shows how exceptions may be used for error recovery in the presence of affine capabilities.
- An echo server is built on top of the socket library.

- Two session types (Gay et al. 2003) libraries demonstrate different approaches to alternation: one uses anonymous sums and the other uses algebraic datatypes for named branches.
- Our version of Sutherland and Hodgman (1974) re-entrant polygon clipping uses session types to connect stream transducers.
- The Alms standard library implements higher-order coercions for checked downcasts which can, for example, turn a function of type $(\text{unit} \xrightarrow{\Delta} \text{unit}) \xrightarrow{\Delta} \text{thread}$ into a function of type $(\text{unit} \xrightarrow{\Delta} \text{unit}) \xrightarrow{\Delta} \text{thread}$ by adding a dynamic check.

These examples are not the last word on what can be done in a language like Alms. Haskell’s type classes (Wadler and Blott 1989), a general mechanism invented to solve specific problems, have since found myriad unanticipated uses. Similarly, a practical, general form of substructural types as offered by Alms likely has many applications waiting to be uncovered.

We have now written several thousand lines of code in Alms, and this experience has led to improvements in both its design and our skill at using what it has to offer. For example, an earlier version of Alms had only unlimited ($\xrightarrow{\Delta}$) and affine ($\xrightarrow{\Delta}$) arrows, but Alms’s behavioral contract library motivated the introduction of arrows whose kinds involve type variables (e.g., $\xrightarrow{\Delta}$). In particular, we found ourselves writing the same function multiple times with different qualifiers in the argument types, and the addition of usage qualifier expressions eliminates this redundancy.

4. The Calculus ${}^a\lambda_{ms}$

We model Alms with a calculus based on System F_{\leq}^{ω} , the higher-order polymorphic λ calculus with subtyping (Pierce 2002). Our calculus, ${}^a\lambda_{ms}$, makes several important changes to F_{\leq}^{ω} :

- Our type system’s structural rules are limited. In particular contraction, which duplicates variables in a typing context to make them accessible to multiple subterms, applies only to variables whose type is of the unlimited kind. Variables of affine or polymorphic kind cannot be contracted.
- Our kind system is enriched with dependent kinds.
- Our kind system is also enriched with variance on type operators (Steffen 1997), which allows abstract type constructors to specify how their results vary in relation to their parameters.
- Type operators are limited to first-order kinds—that is, type operators may not take type operators as parameters.
- Universally-quantified type variables are bounded only by a kind, not by a type.
- The subtyping relation is induced by the subtyping rule for functions, whereby an unlimited-use function may be supplied where a one-use function is expected.

The calculus ${}^a\lambda_{ms}$ also includes more types and terms than a minimal presentation of F_{\leq}^{ω} . Because we are interested in practical issues, we think it is important that our model include products,³ sums, mutable references, and non-termination.

We do not model modules directly in ${}^a\lambda_{ms}$, but its higher-kinded type abstraction makes ${}^a\lambda_{ms}$ a suitable target for the first-order fragment of Rossberg et al.’s (2010) “F-ing modules” translation.

³In linear logic terms, our calculus (like Alms) supplies multiplicative products (\otimes) and additive sums (\oplus) directly. Additive products ($\&$) are easily encoded by

$$\tau_1 \& \tau_2 \triangleq \forall \alpha:A. (\tau_1 \xrightarrow{A} \alpha) + (\tau_2 \xrightarrow{A} \alpha) \xrightarrow{\varphi_1 \sqcup \varphi_2} \alpha$$

$$[e_1, e_2] \triangleq \Lambda \alpha:A. \lambda k. \text{case } k \text{ of } \iota_1 \ x_1 \rightarrow x_1 \ e_1; \iota_2 \ x_2 \rightarrow x_2 \ e_2,$$

where φ_1 and φ_2 are the kinds of τ_1 and τ_2 .


$$\begin{aligned}
\varphi &::= \mathbf{q} \mid \langle \alpha \rangle \mid \varphi_1 \sqcup \varphi_2 \mid \varphi_1 \sqcap \varphi_2 & (\text{usage qualifiers}) \\
\kappa &::= \varphi \mid \Pi \alpha^{\mathbf{v}}. \kappa & (\text{kinds}) \\
\tau &::= \alpha \mid \lambda \alpha. \tau \mid \tau_1 \tau_2 \mid \forall \alpha: \kappa. \tau \mid \tau_1 \xrightarrow{\varphi} \tau_2 & (\text{types}) \\
&\quad \mid \mathbf{1} \mid (+) \mid (\times) \mid \text{aref} \\
e &::= x \mid \lambda x: \tau. e \mid e_1 e_2 \mid \Lambda \alpha: \kappa. v \mid e[\tau] & (\text{terms}) \\
&\quad \mid \langle \rangle \mid \langle e_1, e_2 \rangle \mid \text{case } e_1 \text{ of } \langle x_1, x_2 \rangle \rightarrow e_2 \\
&\quad \mid \iota_1 e \mid \iota_2 e \mid \text{case } e \text{ of } \iota_1 x \rightarrow e_1; \iota_2 y \rightarrow e_2 \\
&\quad \mid \text{fix } e \mid \text{new } e \mid \text{swap } e_1 e_2 \mid \text{delete } e \mid \text{ptr } \ell \\
v &::= \lambda x: \tau. e \mid \Lambda \alpha: \kappa. v \mid \langle \rangle \mid \langle v_1, v_2 \rangle & (\text{values}) \\
&\quad \mid \iota_1 v \mid \iota_2 v \mid \text{ptr } \ell \\
\Gamma, \Sigma &::= \cdot \mid \alpha: \kappa \mid x: \tau \mid \ell: \tau \mid \Gamma_1, \Gamma_2 & (\text{typing contexts})
\end{aligned}$$


Figure 3. Syntax of ${}^a\lambda_{ms}$

4.1 Syntax

We begin with the syntax of ${}^a\lambda_{ms}$ in figure 3. Terms (e) include the usual terms from System F (variables, abstractions, applications, type abstractions, and type applications), several forms for data construction and elimination (nil, pairs, pair elimination, sum injections, and sum elimination), recursion (fix e), and several operations on reference cells (allocation, linear swap, and deallocation). Location names (ptr ℓ) appear at run time but are not present in source terms. Values (v) are standard.

Types (τ) include type variables, type-level abstraction and application, universal quantification, function types, and type constructor constants for unit, sums, products, and references. As in Alms, the function arrow carries a *usage qualifier* (φ), which specifies whether the function is unlimited or one-use.

The two constant usage qualifiers (\mathbf{q}), \mathbf{U} for unlimited and \mathbf{A} for affine, are the bottom and top of the two-element lattice in figure 3. Now consider the K combinator $\Lambda \alpha: \langle \alpha \rangle. \Lambda \beta: \langle \beta \rangle. \lambda x: \alpha. \lambda y: \beta. x$ partially applied to a value: $K[\tau_1][\tau_2] v$. Whether it is safe to duplicate this term depends on whether it is safe to duplicate v , and this is reflected in the instantiation of α . To express this relationship, we introduce usage qualifier expressions (φ), which form a bounded, distributive lattice over type variables with \mathbf{U} and \mathbf{A} as bottom and top. We can thus give K type $\forall \alpha: \langle \alpha \rangle. \forall \beta: \langle \beta \rangle. \alpha \xrightarrow{\mathbf{U}} \beta \xrightarrow{\langle \alpha \rangle} \alpha$.

Qualifier expressions are the base kinds of ${}^a\lambda_{ms}$ —that is, the kinds that classify proper types that may in turn classify values. To classify type operators, kinds (κ) also include dependent product kinds, written $\Pi \alpha^{\mathbf{v}}. \kappa$. This is the kind of a type operator that, when applied to a type with kind φ , gives a type with kind $[\varphi/\alpha]\kappa$. For example, the kind of the Alms list type constructor in §2 is $\Pi \alpha^+. \langle \alpha \rangle$, which means that list τ has the same kind as τ .⁴

The superscript $+$ in kind $\Pi \alpha^+. \langle \alpha \rangle$ means that list is a covariant (or monotone) type constructor: if τ_1 is a subtype of τ_2 then list τ_1 is a subtype of list τ_2 . Variances (\mathbf{v}) form a four-point lattice (figure 3). A type operator may also be contravariant ($-$), where the result varies inversely with the argument; omnivariant (ϕ), where argument may vary freely without affecting the result; or invariant (\pm) where the argument may not vary at all without producing a

⁴ Whereas Alms uses ML’s conventional postfix notation for type-level application, ${}^a\lambda_{ms}$ uses prefix application.

$$\boxed{s; e \mapsto s'; e'} \quad (\text{small-step reduction})$$

$$\begin{aligned}
s; (\lambda x: \tau. e) v &\mapsto s; [v/x]e \\
s; (\Lambda \alpha: \kappa. v)[\tau] &\mapsto s; [\tau/\alpha]v \\
s; \text{fix } v &\mapsto s; v (\text{fix } v) \\
s; \text{new } v &\mapsto s \uplus \{\ell \mapsto v\}; \text{ptr } \ell \\
s \uplus \{\ell \mapsto v_1\}; \text{swap } (\text{ptr } \ell) v_2 &\mapsto s \uplus \{\ell \mapsto v_2\}; \langle \text{ptr } \ell, v_1 \rangle \\
s \uplus \{\ell \mapsto v\}; \text{delete } (\text{ptr } \ell) &\mapsto s; \langle \rangle \\
\frac{s; e \mapsto s'; e'}{s; E[e] \mapsto s'; E[e']} &
\end{aligned}$$

$$\begin{aligned}
E &::= [] \mid E e \mid v E \mid E[\tau] \mid \langle E, e \rangle \mid \langle v, E \rangle & (\text{evaluation contexts}) \\
&\quad \mid \text{case } E \text{ of } \langle x_1, x_2 \rangle \rightarrow e \mid \iota_1 E \mid \iota_2 E \\
&\quad \mid \text{case } E \text{ of } \iota_1 x \rightarrow e_1; \iota_2 y \rightarrow e_2 \\
&\quad \mid \text{fix } E \mid \text{new } E \mid \text{swap } E e \mid \text{swap } v E \mid \text{delete } E \\
s &::= \{ \} \mid \{ \ell \mapsto v \} \mid s_1 \uplus s_2 & (\text{stores})
\end{aligned}$$

Figure 4. Operational semantics (selected rules)

subtyping-unrelated result. We define a composition operation (\cdot) on variances, where $\mathbf{v}_1 \cdot \mathbf{v}_2$ is the variance of the composition of type operators having variances \mathbf{v}_1 and \mathbf{v}_2 .

The kinds of the type constructors for sums and references may aid understanding. The sum type constructor ($+$) has kind $\Pi \alpha^+. \Pi \beta^+. \langle \alpha \rangle \sqcup \langle \beta \rangle$. This means that the kind of a sum type is at least as restrictive as the kinds of its disjuncts. It is covariant in both arguments, which means that $\tau_1 + \tau_2$ is a subtype of $\tau'_1 + \tau'_2$ if τ_1 is a subtype of τ'_1 and τ_2 is a subtype of τ'_2 . The reference type constructor, on the other hand, has kind $\Pi \alpha^{\pm}. \mathbf{A}$. This means that reference cells are always affine and that their types do not support subtyping in either direction.

Typing contexts (Γ or Σ) associate type variables with their kinds, variables with their types, and locations with the types of their contents. By convention, we use Γ for typing contexts that include neither affine variables nor locations, and we use Σ for typing contexts that may include locations and affine (or indeterminate) variables. We use $\text{dom}(\Gamma)$ to refer to the type variables, variables, and locations mapped by a context.

We define the free type variables in a variety of syntactic forms ($\text{ftv}(e)$, $\text{ftv}(\tau)$, $\text{ftv}(\kappa)$, etc.) in the standard way. We use $\text{locs}(e)$ to denote the set of location names present in term e . There are no binders for locations.

4.2 Operational Semantics

The operational semantics of ${}^a\lambda_{ms}$ is mostly a standard call-by-value reduction semantics. We give a selection of rules in figure 4. The reduction relation (\mapsto) relates configurations ($s; e$) comprising a store and a term. A store maps locations (ℓ) to values (v). Stores are taken to be unordered and do not repeat location names.

The rules for reference operations are worth noting. In store s , new v chooses a fresh location ℓ , adding v to the store at location ℓ and reducing to the reference ptr ℓ . The operation swap (ptr ℓ) v_2 requires that the store have location ℓ holding some value v_1 . It swaps v_2 for v_1 in the store, returning a pair of a reference to ℓ and value v_1 . Finally, delete (ptr ℓ) also requires that the store contain ℓ , which it then removes from the store. This means that freeing a location can result in a dangling pointer, which would cause subsequent attempts to access that location to get stuck. Our type system prevents this.

$\Gamma \vdash \kappa$	kind κ is well formed	(fig. 6)
$\Gamma \vdash \kappa_1 <: \kappa_2$	kind κ_1 is subsumed by κ_2	(fig. 6)
$\Gamma \vdash \alpha \in \tau \uparrow \mathbf{v}$	type τ varies \mathbf{v} -ly when α increases	(fig. 7)
$\Gamma \vdash \tau : \kappa$	type τ has kind κ	(fig. 7)
$\Gamma \vdash \tau_1 <:^\mathbf{v} \tau_2$	type τ_1 is \mathbf{v} -related to type τ_2	(fig. 7)
$\Gamma \vdash \Sigma \preceq \varphi$	context Σ is bounded by qualifier φ	(fig. 8)
$\vdash (\Gamma_0; \Sigma_0), \Sigma' \rightsquigarrow \Gamma; \Sigma$	extending $\Gamma_0; \Sigma_0$ with Σ' gives $\Gamma; \Sigma$	(fig. 8)
$\Gamma; \Sigma \triangleright e : \tau$	term e has type τ	(fig. 9)
$\Sigma_1 \triangleright s : \Sigma_2$	store s has type Σ_2	(fig. 10)
$\triangleright s; e : \tau$	configuration $s; e$ has type τ	(fig. 10)

Figure 5. Type system judgments

$\Gamma \vdash \kappa$	(kind well-formedness)
$\Gamma \vdash \kappa_1 <: \kappa_2$	(subkinding)

$\Gamma \vdash \kappa$	(kind well-formedness)
$\Gamma \vdash \kappa_1 <: \kappa_2$	(subkinding)

$\Gamma \vdash \kappa$	(kind well-formedness)
$\Gamma \vdash \kappa_1 <: \kappa_2$	(subkinding)

$\Gamma \vdash \kappa$	(kind well-formedness)
$\Gamma \vdash \kappa_1 <: \kappa_2$	(subkinding)

$\Gamma \vdash \kappa$	(kind well-formedness)
$\Gamma \vdash \kappa_1 <: \kappa_2$	(subkinding)

Figure 6. Statics (i): kinds

4.3 Static Semantics

Our type system involves a large number of judgments, which we summarize in figure 5. In several cases, we omit premises concerning well-formedness of contexts, which clutter the presentation and do not add anything of interest. The omitted premises appear in the extended version of this paper.

Kind judgments. Judgments on kinds appear in figure 6. The first judgment, $\Gamma \vdash \kappa$, determines whether a kind κ is well formed in typing context Γ . A base kind (*i.e.*, a usage qualifier expression) is well formed provided that Γ specifies a kind for all of its free variables. A dependent product kind $\Pi\alpha^\mathbf{v}.\kappa$ is well formed if Γ maps all of its free variables, provided it satisfies a second condition: whenever the bound type variable α is free in κ —that is, when the kind is truly dependent—then variance \mathbf{v} must be $+$ or \pm . This rules out incoherent kinds such as $\Pi\alpha^-. \langle\alpha\rangle$ that classify no useful type operator but whose presence breaks the kinding relation’s monotonicity property (see lemma 4).

The second judgment is subkinding: $\Gamma \vdash \kappa_1 <: \kappa_2$. As we will see, if a type has kind κ_1 , then it may be used where κ_1 or any greater kind is expected. For dependent product kinds the subkinding order is merely the product order on the variance and the result kind, but for base kinds the relation relies on an interpretation of qualifier expressions.

We interpret qualifier expressions via a valuation \mathcal{V} , which is a map from type variables to qualifier constants. We extend \mathcal{V} ’s domain to qualifier expressions:

$$\begin{aligned} \mathcal{V}(q) &= q & \mathcal{V}(\varphi_1 \sqcup \varphi_2) &= \mathcal{V}(\varphi_1) \sqcup \mathcal{V}(\varphi_2) \\ \mathcal{V}(\langle\alpha\rangle) &= \mathcal{V}(\alpha) & \mathcal{V}(\varphi_1 \sqcap \varphi_2) &= \mathcal{V}(\varphi_1) \sqcap \mathcal{V}(\varphi_2) \end{aligned}$$

We need to interpret qualifier expressions under a typing context:

Definition 1 (Consistent valuations). A valuation \mathcal{V} is **consistent** with a typing context Γ if for all $\alpha:\varphi \in \Gamma$, $\mathcal{V}(\alpha) \sqsubseteq \mathcal{V}(\varphi)$.

Thus, a valuation is consistent with a context if it corresponds to a potential instantiation of the type variables, given that context.

Definition 2 (Qualifier subsumption). We say that φ_1 is **subsumed** by φ_2 in Γ , written $\Gamma \models \varphi_1 \sqsubseteq \varphi_2$, if for all valuations \mathcal{V} consistent with Γ , $\mathcal{V}(\varphi_1) \sqsubseteq \mathcal{V}(\varphi_2)$.

In other words, in all possible instantiations of the type variables in Γ , qualifier φ_1 being A implies that φ_2 is A.

Kinding and variance. The first two judgments in figure 7, for computing variances and giving kinds to types, are defined by mutual induction. It should be clear on inspection that the definitions are well-founded. Judgment $\Gamma \vdash \alpha \in \tau \uparrow \mathbf{v}$ means that type variable α appears in type τ at variance \mathbf{v} , or in other words, that type operator $\lambda\alpha.\tau$ has variance \mathbf{v} . Rules V-VAR and V-BOT say that type variables appear positively with respect to themselves and omnivariantly with respect to types in which they are not free. Rule V-ABS says that a type variable appears in a type operator $\lambda\beta.\tau$ at the same variance that it appears in the body τ . The remaining three rules are more involved:

- By rule V-APP, the variance of a type variable in a type application comes from both the operator and the operand. The variance of α in $\tau_1 \tau_2$ is at least the variance of α in τ_1 and at least the variance of α in the τ_2 composed with the variance of operator τ_1 . This makes sense: if τ is a contravariant type operator, then α appears negatively in $\tau \alpha$ but positively in $\tau(\tau \alpha)$.
- By rule V-ALL, the variance of α in $\forall\beta:\kappa.\tau$ is at least its variance in τ . However, if α appears in κ then it is invariant in $\forall\beta:\kappa.\tau$. This reflects the fact that universally-quantified types are related only if their bounds (κ) match exactly, so changing a type variable that appears in κ produces an unrelated type. (This means that ${}^a\lambda_{ms}$ is based on the *kernel* variant of F_{\leq}^∞ ; (Pierce 2002).)
- By rule V-ARR, the variance of α in a function type $\tau_1 \xrightarrow{\varphi} \tau_2$ is at least its variance in the codomain τ_2 and at least the opposite (composition with $-$) of its variance in the domain τ_1 . This reflects function argument contravariance. The variance of α is at least $+$ if it appears in the qualifier expression φ .

The second judgment, $\Gamma \vdash \tau : \kappa$, assigns kinds to well formed types. Rule K-VAR merely looks up the kind of a type variable in the context. Rules K-ABS and K-APP are the usual rules for dependent abstraction and application, with two small changes in rule K-ABS. First, it associates α with *itself* in the context, as $\alpha:\langle\alpha\rangle$, which ensures that occurrences of α in τ can be reflected in κ . Second, it appeals to the variance judgment to determine the variance of the type operator. Rule K-ALL assigns a universal type the same kind as its body, but with A replacing α . This is necessary because the resulting kind is outside the scope of α . Qualifier A is a safe bound for any instantiation of α , and no terms have types that lose precision by this choice. The kind of an arrow type, in rule K-ARR, is just the qualifier expression attached to the arrow. The remaining rules give kinds for type constructor constants.

Type equivalence and dereliction subtyping. The last judgment in figure 7 is subtyping. The subtyping relation is parametrized by a variance \mathbf{v} , which gives the direction of the subtyping: $\Gamma \vdash \tau_1 <:^\mathbf{v} \tau_2$ is the usual direction, judging τ_1 a subtype of τ_2 . In terms of subsumption, this means that values of type τ_1 may be used where values of type τ_2 are expected. The other variances are useful in defining the relation in the presence of \mathbf{v} -variant type operators: ($<:^-$) gives the inverse of the subtyping relation, ($<:^\pm$) relates only equivalent types, and ($<:^\emptyset$) relates all types. We can see how this works this in rule TSUB-APP. To determine whether $\tau_{11} \tau_{12}$ is a

$$\boxed{\Gamma \vdash \alpha \in \tau \Downarrow \mathbf{v}}$$

(variance of type variables with respect to types)

$$\begin{array}{c}
\text{V-VAR} \\
\frac{}{\Gamma \vdash \alpha \in \alpha \Downarrow +} \\
\\
\text{V-BOT} \\
\frac{\alpha \notin \text{ftv}(\tau)}{\Gamma \vdash \alpha \in \tau \Downarrow \phi} \\
\\
\text{V-ABS} \\
\frac{\Gamma, \beta: \langle \beta \rangle \vdash \alpha \in \tau \Downarrow \mathbf{v}}{\Gamma \vdash \alpha \in \lambda \beta. \tau \Downarrow \mathbf{v}} \\
\\
\text{V-APP} \\
\frac{\Gamma \vdash \alpha \in \tau_1 \Downarrow \mathbf{v}_1 \quad \Gamma \vdash \alpha \in \tau_2 \Downarrow \mathbf{v}_2 \quad \Gamma \vdash \tau_1 : \Pi \beta^{\mathbf{v}_3}. \kappa_3}{\Gamma \vdash \alpha \in \tau_1 \tau_2 \Downarrow \mathbf{v}_1 \sqcup (\mathbf{v}_2 \cdot \mathbf{v}_3)} \\
\\
\text{V-ALL} \\
\frac{\Gamma, \beta: \kappa \vdash \alpha \in \tau \Downarrow \mathbf{v}_1 \quad \mathbf{v}_2 = \text{if } \alpha \in \text{ftv}(\kappa) \text{ then } \pm \text{ else } \phi}{\Gamma \vdash \alpha \in \forall \beta: \kappa. \tau \Downarrow \mathbf{v}_1 \sqcup \mathbf{v}_2} \\
\\
\text{V-ARR} \\
\frac{\Gamma \vdash \alpha \in \tau_1 \Downarrow \mathbf{v}_1 \quad \Gamma \vdash \alpha \in \tau_2 \Downarrow \mathbf{v}_2 \quad \mathbf{v}_3 = \text{if } \alpha \in \text{ftv}(\varphi) \text{ then } + \text{ else } \phi}{\Gamma \vdash \alpha \in \tau_1 \xrightarrow{\varphi} \tau_2 \Downarrow -\mathbf{v}_1 \sqcup \mathbf{v}_2 \sqcup \mathbf{v}_3}
\end{array}$$

$$\boxed{\Gamma \vdash \tau : \kappa}$$

(kinding of types)

$$\begin{array}{c}
\text{K-VAR} \\
\frac{\alpha: \kappa \in \Gamma}{\Gamma \vdash \alpha: \kappa} \\
\\
\text{K-ABS} \\
\frac{\Gamma, \alpha: \langle \alpha \rangle \vdash \tau : \kappa}{\Gamma \vdash \lambda \alpha. \tau : \Pi \alpha^{\mathbf{v}}. \kappa} \\
\\
\text{K-APP} \\
\frac{\Gamma \vdash \tau_1 : \Pi \alpha^{\mathbf{v}}. \kappa \quad \Gamma \vdash \tau_2 : \varphi}{\Gamma \vdash \tau_1 \tau_2 : [\varphi/\alpha]\kappa} \\
\\
\text{K-ALL} \\
\frac{\Gamma, \alpha: \kappa \vdash \tau : \varphi}{\Gamma \vdash \forall \alpha: \kappa. \tau : [\mathbf{A}/\alpha]\varphi} \\
\\
\text{K-ARR} \\
\frac{\Gamma \vdash \tau_1 : \varphi_1 \quad \Gamma \vdash \tau_2 : \varphi_2}{\Gamma \vdash \tau_1 \xrightarrow{\varphi} \tau_2 : \varphi} \\
\\
\text{K-UNIT} \\
\frac{}{\Gamma \vdash 1 : \mathbf{U}} \\
\\
\text{K-SUM} \\
\frac{}{\Gamma \vdash (+) : \Pi \alpha^+. \Pi \beta^+. \langle \alpha \rangle \sqcup \langle \beta \rangle} \\
\\
\text{K-PROD} \\
\frac{}{\Gamma \vdash (\times) : \Pi \alpha^+. \Pi \beta^+. \langle \alpha \rangle \sqcup \langle \beta \rangle} \\
\\
\text{K-REF} \\
\frac{}{\Gamma \vdash \text{aref} : \Pi \alpha^{\pm}. \mathbf{A}}
\end{array}$$

$$\boxed{\Gamma \vdash \tau_1 <:^{\mathbf{v}} \tau_2}$$

(subtyping)

$$\begin{array}{c}
\text{TSUB-EQ} \\
\frac{\tau_1 \equiv \tau_2}{\Gamma \vdash \tau_1 <:^{\mathbf{v}} \tau_2} \\
\\
\text{TSUB-TRANS} \\
\frac{\Gamma \vdash \tau_1 <:^{\mathbf{v}} \tau_2 \quad \Gamma \vdash \tau_2 <:^{\mathbf{v}} \tau_3}{\Gamma \vdash \tau_1 <:^{\mathbf{v}} \tau_3} \\
\\
\text{TSUB-OMNI} \\
\frac{}{\Gamma \vdash \tau_1 <:^{\mathbf{0}} \tau_2} \\
\\
\text{TSUB-CONTRA} \\
\frac{\Gamma \vdash \tau_2 <:^{-\mathbf{v}} \tau_1}{\Gamma \vdash \tau_1 <:^{\mathbf{v}} \tau_2} \\
\\
\text{TSUB-ABS} \\
\frac{\Gamma, \alpha: \langle \alpha \rangle \vdash \tau_1 <:^{\mathbf{v}} \tau_2}{\Gamma \vdash \lambda \alpha. \tau_1 <:^{\mathbf{v}} \lambda \alpha. \tau_2} \\
\\
\text{TSUB-APP} \\
\frac{\Gamma \vdash \tau_{11} : \Pi \alpha^{\mathbf{v}_1}. \kappa_1 \quad \Gamma \vdash \tau_{21} : \Pi \alpha^{\mathbf{v}_2}. \kappa_2 \quad \Gamma \vdash \tau_{11} <:^{\mathbf{v}} \tau_{21} \quad \Gamma \vdash \tau_{12} <:^{\mathbf{v} \cdot (\mathbf{v}_1 \sqcup \mathbf{v}_2)} \tau_{22}}{\Gamma \vdash \tau_{11} \tau_{12} <:^{\mathbf{v}} \tau_{21} \tau_{22}} \\
\\
\text{TSUB-ALL} \\
\frac{\Gamma, \alpha: \kappa \vdash \tau_1 <:^{\mathbf{v}} \tau_2}{\Gamma \vdash \forall \alpha: \kappa. \tau_1 <:^{\mathbf{v}} \forall \alpha: \kappa. \tau_2} \\
\\
\text{TSUB-ARR} \\
\frac{\Gamma \vdash \tau_{11} <:^{-\mathbf{v}} \tau_{21} \quad \Gamma \vdash \tau_{12} <:^{\mathbf{v}} \tau_{22} \quad \Gamma \vdash \varphi_1 <:^{\mathbf{v}} \varphi_2}{\Gamma \vdash \tau_{11} \xrightarrow{\varphi_1} \tau_{12} <:^{\mathbf{v}} \tau_{21} \xrightarrow{\varphi_2} \tau_{22}}
\end{array}$$

Figure 7. Statics (ii): types

subtype of $\tau_{21} \tau_{22}$, we take \mathbf{v} to be $+$, yielding

$$\frac{\Gamma \vdash \tau_{11} : \Pi \alpha^{\mathbf{v}_1}. \kappa_1 \quad \Gamma \vdash \tau_{21} : \Pi \alpha^{\mathbf{v}_2}. \kappa_2 \quad \Gamma \vdash \tau_{11} <:^{+} \tau_{21} \quad \Gamma \vdash \tau_{12} <:^{\mathbf{v}_1 \sqcup \mathbf{v}_2} \tau_{22}}{\Gamma \vdash \tau_{11} \tau_{12} <:^{+} \tau_{21} \tau_{22}}$$

This means that for the subtyping relation to hold:

- The operators must be related in the same direction, so that τ_{11} is a subtype of τ_{21} .
- The operands must be related in the direction given by the variances of the operators. For example, if both operators are covariant, then the operands must vary in the same direction, so that τ_{12} is a subtype of τ_{22} . If both operators are contravariant, then the operands must vary in the opposite direction. If the operators are invariant then the operands cannot vary at all, but if they are omnivariant then $\tau_{11} \tau'_{12}$ is a subtype of $\tau_{21} \tau'_{22}$ for any τ'_{12} and τ'_{22} .

Rule TSUB-EQ says that subtyping includes type equivalence ($\tau_1 \equiv \tau_2$), which is merely β equivalence on types. Rule TSUB-OMNI allows any pair of types to be related by ϕ -variant subtyping, and rule TSUB-CONTRA says that the opposite variance sign gives the inverse relation. Rules TSUB-ABS and TSUB-ALL specify that type operators and universally-quantified types are related if their bodies are.

Rule TSUB-ARR is more than the usual arrow subtyping rule. Beyond the usual contravariance for arguments and covariance for results, it requires that qualifiers φ_1 and φ_2 relate in the same

direction. This rule is the source of non-trivial subtyping in ${}^a\lambda_{ms}$, without which subtyping would relate only equivalent types. The rule has two important implications.

First, an unlimited-use function can always be used where a one-use function is expected. This corresponds to linear logic's usual dereliction rule, which says that the ! ("of course!") modality may always be removed. ILL (Bierman 1993) has a rule:

$$\frac{\Delta \vdash e : !A}{\Delta \vdash \text{derelict } e : A} \text{DERELICTION.}$$

Dereliction is syntax-directed in this rule, but for practical programming we consider that as too inconvenient. Thus, our subtyping relation supports dereliction as needed.

For example, the function for creating a new thread in Alms, *Thread.fork*, has type $\forall \hat{\alpha}. (\text{unit} \xrightarrow{\mathbf{A}} \hat{\alpha}) \xrightarrow{\mathbf{U}} \hat{\alpha} \text{ thread}$, which means that *Thread.fork* will not call its argument more than once. However, this should not stop us from passing an unlimited-use function to *Thread.fork*, and indeed we can. Dereliction subtyping allows us to use a value of type $\text{unit} \xrightarrow{\mathbf{U}} \hat{\alpha}$ where a value of type $\text{unit} \xrightarrow{\mathbf{A}} \hat{\alpha}$ is expected. Alternatively, by domain contravariance, we can use *Thread.fork* where a value of type $\forall \hat{\alpha}. (\text{unit} \xrightarrow{\mathbf{U}} \hat{\alpha}) \xrightarrow{\mathbf{U}} \hat{\alpha}$ thread is expected. In this case subsumption allows us to forget *Thread.fork*'s promise not to reuse its argument.

The other important implication of dereliction subtyping will become clearer once we see how qualifier expressions are assigned to function types. Subsumption makes it reasonable to always assign functions the most permissive safe usage qualifier, because

$\boxed{\Gamma \vdash \Sigma \preceq \varphi}$	(bound of typing context)
$\frac{\text{B-NIL}}{\Gamma \vdash \cdot \preceq \mathbf{U}}$	$\frac{\text{B-CONSA} \quad \Gamma \vdash \Sigma \preceq \varphi_1 \quad \Gamma \vdash \kappa}{\Gamma \vdash \Sigma, \alpha : \kappa \preceq \varphi_1}$
$\frac{\text{B-CONSX} \quad \Gamma \vdash \Sigma \preceq \varphi_1 \quad \Gamma \vdash \tau : \varphi_2}{\Gamma \vdash \Sigma, x : \tau \preceq \varphi_1 \sqcup \varphi_2}$	$\frac{\text{B-CONSL} \quad \Gamma \vdash \Sigma \preceq \varphi_1 \quad \Gamma \vdash \tau : \varphi_2}{\Gamma \vdash \Sigma, \ell : \tau \preceq \mathbf{A}}$
$\boxed{\vdash (\Gamma_0; \Sigma_0), \Sigma' \rightsquigarrow \Gamma_1; \Sigma_1}$	(typing context extension)
$\frac{\text{X-NIL}}{\vdash (\Gamma; \Sigma), \cdot \rightsquigarrow \Gamma; \Sigma}$	
$\frac{\text{X-CONSU} \quad \Gamma_0 \vdash \tau : \mathbf{U}}{\vdash (\Gamma_0, x : \tau; \Sigma_0), \Sigma' \rightsquigarrow \Gamma; \Sigma}$	$\frac{\text{X-CONSA} \quad \Gamma_0 \vdash \tau : \varphi}{\vdash (\Gamma_0; \Sigma_0, x : \tau), \Sigma' \rightsquigarrow \Gamma; \Sigma}$
$\frac{}{\vdash (\Gamma_0; \Sigma_0), x : \tau, \Sigma' \rightsquigarrow \Gamma; \Sigma}$	$\frac{}{\vdash (\Gamma_0; \Sigma_0), x : \tau, \Sigma' \rightsquigarrow \Gamma; \Sigma}$

Figure 8. Statics (iii): typing contexts

subsumption then allows us to use them in a less permissive context. Dereliction subtyping applies only to function types because in both the ${}^a\lambda_{ms}$ calculus and Alms language only function types carry qualifiers. For instance, Alms has no separate types $\text{int}^{\mathbf{U}}$ for unlimited integers and $\text{int}^{\mathbf{A}}$ for affine integers. Integers are always unlimited. If a programmer wants an affine version of int , she can create it in Alms using the module system.

Context judgments. Figure 8 defines two judgments on contexts. Judgment $\Gamma \vdash \Sigma \preceq \varphi$, which will be important in typing functions, computes an upper bound φ on the qualifiers of all the types in context Σ . If a context contains any locations, it is bounded by \mathbf{A} ; otherwise, its bound is the least upper bound of the qualifiers of all the types of variables in the context.

The second judgment shows how environments are extended by variable bindings. The typing judgment for terms will use two typing contexts: Γ holds environment information that may be safely duplicated, such as type variables and variables of unlimited type, whereas Σ holds information, such as location types and affine variables, that disallows duplication. Given contexts Γ_0 and Σ_0 , judgment $\vdash (\Gamma_0; \Sigma_0), \Sigma' \rightsquigarrow \Gamma; \Sigma$ extends them by the variables and types in Σ' to get Γ and Σ . Any variables may be placed in Σ , but only variables whose types are known to be unlimited may be placed in Γ , since Γ may be duplicated.

Term judgment. The typing judgment for terms appears in figure 9. The judgment, $\Gamma; \Sigma \triangleright e : \tau$, uses two typing contexts in the style of DILL (Barber 1996): the unlimited environment Γ and the affine environment Σ . When typing multiplicative terms such as application, we distribute Γ to both subterms but partition Σ between the two:

$$\frac{\Gamma; \Sigma_1 \triangleright e_1 : \tau_1 \xrightarrow{\varphi} \tau_2 \quad \Gamma; \Sigma_2 \triangleright e_2 : \tau_1}{\Gamma; \Sigma_1, \Sigma_2 \triangleright e_1 e_2 : \tau_2} \text{T-APP}$$

Unlike DILL, not all types in Σ are necessarily affine. Since types whose usage qualifier involves type variables are not known to be unlimited, we place those in Σ , to ensure that we do not duplicate values that *might* turn out to be affine once universally-quantified types are instantiated.

The other multiplicative rules are T-PAIR for product introduction, T-UNPAIR for product elimination, and T-SWAP for reference updates. Note that T-SWAP does not require that the type of the

reference in its first parameter match the type of the value in its second—in other words, swap performs a *strong update*. To type the term $\text{case } e \text{ of } \langle x_1, x_2 \rangle \rightarrow e_1$, rule T-UNPAIR first splits the affine environment into Σ_1 for typing subterm e and Σ_2 for subterm e_1 . It invokes the context extension relation (figure 8) to extend Γ and Σ_2 with bindings for x_1 and x_2 in order to type e_1 . The context extension relation requires that variables not known to be unlimited be added to Σ_2 .

The rule for sum elimination, T-CHOOSE, is both multiplicative and additive: the affine context is split between the term being destructed and the branches of the case expression. However, the portion of the context given to the branches is shared between them, because only one or the other will be evaluated. Rule T-CHOOSE also uses the context extension relation to bind the pattern variables for the branches.

Rules T-NEW and T-DELETE introduce and eliminate reference types in the usual way. Likewise, the sum introduction rules T-IN κ and type abstraction rule T-TABS are standard. Rules T-VAR, T-PTR, and T-UNIT are standard for an affine calculus but not a linear one, as they implicitly support weakening by allowing Σ to contain unused bindings. Rule T-FIX is also standard, modulo the reasonable constraint that its parameter function be unlimited, since the reduction rule for fix makes a copy of the parameter.

The type application rule T-TAPP supports subkinding, because it requires only that the kind of the actual type parameter be a sub-kind of that of the formal parameter. This is the rule that supports the sort of type abstraction that we used in our examples of §2 to construct affine capabilities. For example, the rule lets us instantiate affine type variable α with unlimited unit type $\mathbf{1}$:

$$\frac{\Gamma; \Sigma \triangleright (\Lambda \alpha : \mathbf{A}. \lambda x : \alpha. e) : \forall \alpha : \mathbf{A}. \alpha \xrightarrow{\varphi} \tau \quad \Gamma \vdash \mathbf{1} : \mathbf{U} \quad \Gamma \vdash \mathbf{U} < : \mathbf{A}}{\Gamma; \Sigma \triangleright (\Lambda \alpha : \mathbf{A}. \lambda x : \alpha. e)[\mathbf{1}] : \mathbf{1} \xrightarrow{\varphi} \tau} \text{T-TAPP}$$

Within its scope, α is considered *a priori* affine, regardless of how it may eventually be instantiated. This term types only if x appears in affine fashion in e .

This brings us finally to T-ABS, the rule for typing term-level λ abstractions. To type a term $\lambda x : \tau_1. e$, rule T-ABS uses the context extension relation to add $x : \tau_1$ to its contexts and types the body e in the extended contexts. It also must determine the qualifier φ that decorates the arrow. Because abstractions close over their free variables, duplicating a function also duplicates the values of its free variables. Therefore, the qualifier of a function type should be at least as restrictive as the qualifiers of the abstraction's free variables. To do this, rule T-ABS appeals to the context bounding judgment (figure 8) to find the least upper bound of the usage qualifiers of variables in the affine environment, and it requires that the function type's qualifier be equally restrictive.

This refines linear logic's usual promotion rule, which says that the $!$ modality may be added to propositions that in turn depend only on $!$ -ed resources. In ILL, we have

$$\frac{! \Delta \vdash e : A}{! \Delta \vdash \text{promote } e : ! A} \text{PROMOTION,}$$

where $! \Delta$ is a context in which all assumptions are of the form $x : ! B$. As with dereliction, in our system it only makes sense to apply promotion to function types.

Our treatment of promotion indicates why we need the explicit weakening rule T-WEAK, which allows discarding unused portions of the affine environment. In order to give a function type the best qualifier possible, we need to remove from Σ any unused variables or locations that might otherwise raise the bound on Σ , and the algorithmic version of the type system as implemented in

$\boxed{\Gamma; \Sigma \triangleright e : \tau}$

(typing of terms)

$\frac{\text{T-SUBSUME} \quad \Gamma; \Sigma \triangleright e : \tau' \quad \Gamma \vdash \tau' <: ^+ \tau}{\Gamma; \Sigma \triangleright e : \tau}$	$\frac{\text{T-WEAK} \quad \Gamma; \Sigma \triangleright e : \tau}{\Gamma; \Sigma, \Sigma' \triangleright e : \tau}$	$\frac{\text{T-VAR} \quad x : \tau \in \Gamma, \Sigma}{\Gamma; \Sigma \triangleright x : \tau}$	$\frac{\text{T-PTR} \quad \ell : \tau \in \Sigma}{\Gamma; \Sigma \triangleright \text{ptr } \ell : \text{aref } \tau}$
$\frac{\text{T-ABS} \quad \vdash (\Gamma; \Sigma), x : \tau_1 \rightsquigarrow \Gamma'; \Sigma' \quad \Gamma'; \Sigma' \triangleright e : \tau_2 \quad \Gamma \vdash \Sigma \preceq \varphi}{\Gamma; \Sigma \triangleright \lambda x : \tau_1. e : \tau_1 \xrightarrow{\varphi} \tau_2}$	$\frac{\text{T-APP} \quad \Gamma; \Sigma_1 \triangleright e_1 : \tau_1 \xrightarrow{\varphi} \tau_2 \quad \Gamma; \Sigma_2 \triangleright e_2 : \tau_1}{\Gamma; \Sigma_1, \Sigma_2 \triangleright e_1 e_2 : \tau_2}$	$\frac{\text{T-TABS} \quad \Gamma, \alpha : \kappa; \Sigma \triangleright v : \tau}{\Gamma; \Sigma \triangleright \Lambda \alpha : \kappa. v : \forall \alpha : \kappa. \tau}$	
$\frac{\text{T-TAPP} \quad \Gamma; \Sigma \triangleright e : \forall \alpha : \kappa. \tau \quad \Gamma \vdash \tau' : \kappa' \quad \Gamma \vdash \kappa' <: \kappa}{\Gamma; \Sigma \triangleright e[\tau'] : [\tau'/\alpha]\tau}$	$\frac{\text{T-UNIT}}{\Gamma; \Sigma \triangleright \langle \rangle : 1}$	$\frac{\text{T-PAIR} \quad \Gamma; \Sigma_1 \triangleright e_1 : \tau_1 \quad \Gamma; \Sigma_2 \triangleright e_2 : \tau_2}{\Gamma; \Sigma_1, \Sigma_2 \triangleright \langle e_1, e_2 \rangle : \tau_1 \times \tau_2}$	
$\frac{\text{T-UNPAIR} \quad \Gamma; \Sigma_1 \triangleright e : \tau_1 \times \tau_2 \quad \vdash (\Gamma; \Sigma_2), x_1 : \tau_1, x_2 : \tau_2 \rightsquigarrow \Gamma'; \Sigma' \quad \Gamma'; \Sigma' \triangleright e_1 : \tau}{\Gamma; \Sigma_1, \Sigma_2 \triangleright \text{case } e \text{ of } \langle x_1, x_2 \rangle \rightarrow e_1 : \tau}$		$\frac{\text{T-IN}k \quad (k \in \{1, 2\}) \quad \Gamma; \Sigma \triangleright e : \tau_k}{\Gamma; \Sigma \triangleright \iota_k e : \tau_1 + \tau_2}$	
$\frac{\text{T-CHOOSE} \quad \Gamma; \Sigma \triangleright e : \tau_1 + \tau_2 \quad \vdash (\Gamma; \Sigma'), x_1 : \tau_1 \rightsquigarrow \Gamma_1; \Sigma_1 \quad \Gamma_1; \Sigma_1 \triangleright e_1 : \tau \quad \vdash (\Gamma; \Sigma'), x_2 : \tau_2 \rightsquigarrow \Gamma_2; \Sigma_2 \quad \Gamma_2; \Sigma_2 \triangleright e_2 : \tau}{\Gamma; \Sigma, \Sigma' \triangleright \text{case } e \text{ of } \iota_1 x_1 \rightarrow e_1; \iota_2 x_2 \rightarrow e_2 : \tau}$			
$\frac{\text{T-FIX} \quad \Gamma; \Sigma \triangleright e : \tau \xrightarrow{u} \tau}{\Gamma; \Sigma \triangleright \text{fix } e : \tau}$	$\frac{\text{T-NEW} \quad \Gamma; \Sigma \triangleright e : \tau}{\Gamma; \Sigma \triangleright \text{new } e : \text{aref } \tau}$	$\frac{\text{T-SWAP} \quad \Gamma; \Sigma_1 \triangleright e_1 : \text{aref } \tau_1 \quad \Gamma; \Sigma_2 \triangleright e_2 : \tau_2}{\Gamma; \Sigma_1, \Sigma_2 \triangleright \text{swap } e_1 e_2 : \text{aref } \tau_2 \times \tau_1}$	$\frac{\text{T-DELETE} \quad \Gamma; \Sigma \triangleright e : \text{aref } \tau}{\Gamma; \Sigma \triangleright \text{delete } e : 1}$

Figure 9. Statics (iv): terms

 $\boxed{\Sigma_1 \triangleright s : \Sigma_2}$

(store typing)

$\frac{\text{S-NIL}}{\Sigma_1 \triangleright \{ \} : \cdot}$	$\frac{\text{S-CONS} \quad \Sigma_{11} \triangleright s : \Sigma_2 \quad \cdot; \Sigma_{12} \triangleright v : \tau}{\Sigma_{11}, \Sigma_{12} \triangleright s \uplus \{ \ell \mapsto v \} : \Sigma_2, \ell : \tau}$
--	--

 $\boxed{\triangleright s; e : \tau}$

(configuration typing)

$\frac{\text{CONF} \quad \Sigma_1 \triangleright s : \Sigma_1, \Sigma_2 \quad \cdot; \Sigma_2 \triangleright e : \tau}{\triangleright s; e : \tau}$

Figure 10. Statics (v): stores and configurations

Alms does just that. In §5 we show that our implicit promotion mechanism selects the best usage qualifier for function types.

Store and configuration judgments. In order to prove our type soundness theorem, we need to lift our typing judgments to stores and run-time configurations.

The type of a store is a typing context containing the names of the store's locations and the types of their contents. The store typing judgment $\Sigma_1 \triangleright s : \Sigma_2$ gives store s type Σ_2 in the context of Σ_1 , which is necessary because values in the store may refer to other values in the store. Rule S-CONS shows that the resources represented by context Σ_1 (i.e., Σ_{11}, Σ_{12}) are split between the values in s .

Our preservation lemma concerns typing judgments on configurations, $\triangleright s; e : \tau$, which means that e has type τ in the context of store s . To type the configuration by rule CONF, we type the store, splitting its type into Σ_1 , which contains locations referenced from the store, and Σ_2 , which contains locations referenced from e .

5. Theoretical Results

We now state our two main theorems—principal qualifiers and type soundness—and sketch their proofs. Full versions of our proofs may be found in an extended version of this paper available at www.ccs.neu.edu/~tov/pubs/alms.

Principal qualifiers. Alms and ${}^a\lambda_{ms}$ go to a lot of trouble to find the best usage qualifier expressions for function types. To make programming with affine types as convenient as possible, we want to maximize polymorphism between one-use and unlimited versions of functions. While writing the Alms standard library, we found that usage qualifier constants **A** and **U**, even with dereliction subtyping, were insufficient to give a principal type to some terms.

For example, consider function *default*, an eliminator for option types, sans function argument types:

```

let default (def: ...) (opt: ...) =
  match opt with
  | Some  $x \rightarrow x$ 
  | None    $\rightarrow \text{def}$ 

```

Without usage qualifier expressions, *default* has at least two incomparable types:

$$\begin{aligned} \text{default}_1 &: \forall \hat{\alpha}. \hat{\alpha} \xrightarrow{u} \hat{\alpha} \text{ option} \xrightarrow{A} \hat{\alpha} \\ \text{default}_2 &: \forall \alpha. \alpha \xrightarrow{u} \alpha \text{ option} \xrightarrow{U} \alpha. \end{aligned}$$

In the first case, because $\hat{\alpha}$ *might* be affine, the partial application of *default*₁ must be a one-use function, but in the second case we know that α is unlimited so partially applying *default*₂ and reusing the result is safe. Formally, these types are incomparable because the universally-quantified type variable $\hat{\alpha}$ in the former has a different kind than α in the latter, and Alms uses the *kernel* variant of rule T-SUB-ALL. However, even were we to replace

rule T_{SUB}-ALL with a rule analogous to F_{\leq}° 's full variant,

$$\frac{\Gamma, \alpha : \kappa \vdash \tau_1 <:^{\mathfrak{v}} \tau_2 \quad \Gamma, \alpha : \kappa \vdash \kappa_1 <:^{-\mathfrak{v}} \kappa_2}{\Gamma \vdash \forall \alpha : \kappa_1. \tau_1 <:^{\mathfrak{v}} \forall \alpha : \kappa_2. \tau_2} \text{T}_{\text{SUB-ALLFULL}},$$

the types would not be related by the subtyping order. More importantly, neither type is preferable in an informal sense. The type of *default*₁ allows $\hat{\alpha}$ to be instantiated to an affine or unlimited type, but the result of partially applying it is a one-use function even if $\hat{\alpha}$ is known to be unlimited:

$$\begin{aligned} \text{default}_1 5 & : \text{int option } \overset{\mathbf{A}}{\rightarrow} \text{int} \\ \text{default}_1 (\text{aref } 5) & : \text{int aref option } \overset{\mathbf{A}}{\rightarrow} \text{int aref.} \end{aligned}$$

If we choose *default*₂, the result of partial application is unlimited, but attempting to instantiate α to an affine type is a type error:

$$\begin{aligned} \text{default}_2 5 & : \text{int option } \overset{\mathbf{U}}{\rightarrow} \text{int} \\ \text{default}_2 (\text{aref } 5) & : \text{Type error!} \end{aligned}$$

Alms avoids both problems and instead discovers that the best usage qualifier for the arrow is the kind of the type variable:

$$\begin{aligned} \text{default} & : \forall \hat{\alpha}. \hat{\alpha} \overset{\mathbf{U}}{\rightarrow} \hat{\alpha} \text{ option } \overset{\langle \hat{\alpha} \rangle}{\rightarrow} \hat{\alpha} \\ \text{default } 5 & : \text{int option } \overset{\mathbf{U}}{\rightarrow} \text{int} \\ \text{default } (\text{aref } 5) & : \text{int aref option } \overset{\mathbf{A}}{\rightarrow} \text{int aref.} \end{aligned}$$

Because this is an important property, we prove a theorem that every typeable ${}^a\lambda_{ms}$ function has a principal usage qualifier.

Theorem 3 (Principal qualifiers). *If $\Gamma; \Sigma \triangleright \lambda x : \tau. e : \tau_1 \xrightarrow{\varphi} \tau_2$, then it has a least qualifier expression φ_0 ; that is,*

- $\Gamma; \Sigma \triangleright \lambda x : \tau. e : \tau_1 \xrightarrow{\varphi_0} \tau_2$ and
- $\Gamma \vdash \varphi_0 <: \varphi'$ for all φ' such that $\Gamma; \Sigma \triangleright \lambda x : \tau. e : \tau_1 \xrightarrow{\varphi'} \tau_2$.

Proof sketch. We obtain the principal qualifier φ_0 as follows. Let Σ_0 be the restriction of Σ to exactly the free variables and locations of $\lambda x : \tau. e$. Let φ_0 be the unique bound of Σ_0 given by $\Gamma \vdash \Sigma_0 \preceq \varphi_0$. By strengthening, $\Gamma; \Sigma_0 \triangleright \lambda x : \tau. e : \tau_1 \xrightarrow{\varphi_0} \tau_2$, and by rule T-WEAK we can get the same type in Σ .

A derivation of $\Gamma; \Sigma \triangleright \lambda x : \tau. e : \tau_1 \xrightarrow{\varphi'} \tau_2$ always involves rule T-ABS using some portion of Σ , followed by some number of subsumptions and weakenings. Subsumption will never let φ' be less than φ_0 . However, weakening might allow us to type $\lambda x : \tau. e$ with a different portion of Σ than Σ_0 . We know that any superset of Σ_0 has bound no less than φ_0 , and while a non-superset of Σ_0 may have a smaller bound, we chose Σ_0 so that only Σ_0 and its supersets are suitable to type the term and then weaken to Σ . \square

Thus, for a function in any given context, there is a least usage qualifier, and our implementation can find the least qualifier by considering only the portion of Σ that pertains to the free identifiers of the λ term, as suggested by the algorithmic rule

$$\text{(T-ABS}_{\text{ALG}}) \quad \varphi = \frac{\begin{array}{l} \vdash (\Gamma; \Sigma), x : \tau_1 \rightsquigarrow \Gamma'; \Sigma' \quad \Gamma'; \Sigma' \triangleright e : \tau_2 \\ \mathbf{A} \quad \text{if } \text{locs}(e) \neq \emptyset \\ \left[\bigcup \{ \varphi \mid x \in \text{fv}(e), \Gamma \vdash \Sigma(x) : \varphi \} \right] \quad \text{otherwise} \end{array}}{\Gamma; \Sigma \triangleright \lambda x : \tau_1. e : \tau_1 \xrightarrow{\varphi_0} \tau_2}.$$

Type soundness. The key obstacle in our type soundness proof is establishing a substitution lemma, which in turn relies on showing that the kind of the type of any value accurately reflects the resources contained in that value, which itself comes as a corollary to the proposition that the kinds of subtypes are themselves subkinds:

Lemma 4 (Monotonicity of kinding). *If $\Gamma \vdash \tau_1 <:^+ \tau_2$ where $\Gamma \vdash \tau_1 : \varphi_1$ and $\Gamma \vdash \tau_2 : \varphi_2$, then $\Gamma \vdash \varphi_1 <: \varphi_2$.*

This lemma is the reason for the premise in rule OK-ARR that for a kind $\Pi \alpha^{\mathfrak{v}}. \kappa$, variance \mathfrak{v} must be at least + if $\alpha \in \text{ftv}(\kappa)$. Otherwise, we could construct a counterexample to lemma 4:

- $\beta : \Pi \alpha^-. \langle \alpha \rangle \vdash \beta (1 \overset{\mathbf{A}}{\rightarrow} 1) <:^+ (1 \overset{\mathbf{U}}{\rightarrow} 1)$,
- $\beta : \Pi \alpha^-. \langle \alpha \rangle \vdash \beta (1 \overset{\mathbf{A}}{\rightarrow} 1) : \mathbf{A}$, and
- $\beta : \Pi \alpha^-. \langle \alpha \rangle \vdash \beta (1 \overset{\mathbf{U}}{\rightarrow} 1) : \mathbf{U}$,
- but $\beta : \Pi \alpha^-. \langle \alpha \rangle \vdash \mathbf{A} <: \mathbf{U}$ is not the case.

The kind well-formedness judgment rules out kinds like $\Pi \alpha^-. \langle \alpha \rangle$.

Proof for lemma 4. We define an extension of the subkinding relation, $\Gamma \vdash \kappa_1 \lesssim \kappa_2$, which is insensitive to the variances decorating Π kinds. Observe that on qualifier expressions this new relation coincides with subkinding. We generalize the induction hypothesis—if $\Gamma \vdash \tau_1 <:^+ \tau_2$ where $\Gamma \vdash \tau_1 : \kappa_1$ and $\Gamma \vdash \tau_2 : \kappa_2$, then $\Gamma \vdash \kappa_1 \lesssim \kappa_2$ —and complete the proof by induction on the structure of the subtyping derivation. \square

Corollary 5 (Kinding finds locations). *Suppose that $\Gamma; \Sigma \triangleright v : \tau$ and $\Gamma \vdash \tau : \varphi$ where $\text{dom}(\Sigma)$ contains only locations (ℓ). If any locations appear in v then $\Gamma \vdash \mathbf{A} <: \varphi$.*

Proof sketch. By induction on the typing derivation. We use the previous lemma in the case for the subsumption rule T-SUBSUME:

$$\text{Case } \frac{\Gamma; \Sigma \triangleright v : \tau' \quad \Gamma \vdash \tau' <:^+ \tau \quad \Gamma \vdash \tau : \varphi}{\Gamma; \Sigma \triangleright v : \tau}.$$

By the induction hypothesis, $\Gamma \vdash \tau' : \mathbf{A}$, and by lemma 4, $\Gamma \vdash \mathbf{A} <: \varphi$. \square

Corollary 5 lets us prove our substitution lemma. Then progress, preservation, and type soundness are standard:

Theorem 6 (Type soundness). *If $\triangleright \{ \}; e : \tau$ then either e diverges or there exists some store s and value v such that $\{ \}; e \mapsto^* s; v$ and $\triangleright s; v : \tau$.*

6. Related Work

In prior work, we showed how an Alms-like affine language may safely interoperate with a conventional (non-affine) language (Tov and Pucella 2010). In particular, the languages may freely share values, including functions. Attempts by the conventional language to subvert the affine language's invariants are prevented by dynamic checks in the form of behavioral software contracts. That paper focused specifically on multi-language interaction, using a predecessor of Alms.

System F° . Mazurak et al. (2010) describe a calculus of “light-weight linear types.” Their primary motivation is similar to ours: to remove needless overhead and provide a “simple foundation for practical linear programming.”

System F° and the prior iteration of Alms independently introduced several new ideas:

- Both use kinds to distinguish linear (in Alms, affine) types from unlimited types, where F° 's kinds \circ and \star correspond to our \mathbf{A} and \mathbf{U} , and their subkinding relation $\star \leq \circ$ corresponds to our $\mathbf{U} \sqsubseteq \mathbf{A}$.
- F° uses existentials and subkinding to abstract unlimited types into linear types. Alms (the language) uses modules and ${}^a\lambda_{ms}$ (the calculus) uses higher-kinded type abstraction to define abstract affine types, including type constructors with parameters.

Mazurak et al. mention the possibility of extending F° with abstraction over higher kinds but do not show the details.

- They sketch out a convenient notation for writing linear computations. This inspired our different implicit threading syntax, which is implemented in Alms as mentioned at the end of §2.

There are also notable differences:

- F° has linear types, which disallow weakening, whereas Alms has affine types, which support it. This is a trade-off. Linear types make it possible to enforce liveness properties, which may be useful, for instance, to ensure that manual memory management does not leak. On the other hand, we anticipate that safely combining linearity with exceptions requires a type-and-effect system to track when raising an exception would implicitly discard linear values. Alms can support explicit deallocation so long as failure to do so is backed up by a garbage collector.
- Alms’s unlimited-use function type is a subtype of its one-use function type. F° does not provide subtyping, though they do show how η expansion can explicitly perform the coercion that our subtyping does implicitly. Experience with our implementation confirms that dereliction subtyping is valuable, though we admit it comes at the cost of complexity.
- F° requires annotating abstractions ($\lambda^k x:\tau.e$) to specify the kind of the resulting arrow type, which may only be \star or \circ . Alms refines this with qualifier expressions and selects the least kind automatically.
- Mazurak et al. give a resource-aware semantics and prove that they can encode regular protocols. We do neither but conjecture that our system enjoys similar properties, except that weakening makes it possible to bail out of a protocol at any point.
- Their sketch of rules for algebraic datatypes is similar to how ours work, though ours are strictly stronger. For example, an option type in F° would have two versions:

$$\text{optionLin} : \circ \Rightarrow \circ \quad \text{optionUn} : \star \Rightarrow \star.$$

Our dependent kinds in Alms let us define one type constructor whose kind subsumes both:

$$\text{option} : \Pi \hat{\alpha}^+ . \langle \hat{\alpha} \rangle.$$

Clean. At first glance, Clean’s uniqueness types *appear* to be dual to affine types. Uniqueness types are descriptive—they indicate that a particular reference is unique—while affine (and linear) types are prescriptive, since they restrict what may be done to some reference in the future but do not necessarily know where it’s been. Similarly, Clean’s subtyping relation, which allows forgetting that a value is unique, appears dual to Alms’s, which allows pretending that an unlimited value is affine. However, the duality breaks down in the higher-order case. When a partially applied function captures some unique free variable, Clean’s type system must *prohibit* aliasing of the function in order to maintain descriptive uniqueness when the function is fully applied (Plasmeijer and Eekelen 2002). In Clean’s terminology, function types with the unique attribute are “essentially unique,” but we might call them “affine.”

There is a strong similarity between our kinding judgment and Clean’s uniqueness propagation rules that relate the uniqueness of data structures to that of their constituent parts. While Clean supports subtyping, it does not have a subkinding relation analogous to Alms or F° ’s. In particular, Clean requires that the uniqueness attributes declared for an abstract type in a module’s interface exactly match the uniqueness attributes in the module’s implementation.

Use types and qualifiers. Wadler (1991) discusses several variants of linear type systems. He proposes something akin to dereliction subtyping (*i.e.*, $!A \leq A$) and points out that in such a system,

terms such as $\lambda f.\lambda x.f x$ have several unrelated types. (We made a similar observation in §5.) In order to recover principal types, he introduces use types, which decorate the exponential modality with a *use variable* $i: !^i$. The use variable ranges over 0 and 1, where $!^0 A = A$ and $!^1 A = !A$. This provides principal types, but at the cost of adding use-variable inequality constraints to type schemes.

A use-variable inequality of the form $i \leq j$ is essentially an implication $i \supset j$, where 1 is truth and 0 is falsity. De Vries et al. (2008) show, in the setting of uniqueness types, how such inequalities may be represented instead using Boolean logic. For example, if we have a type

$$\dots !^i \dots !^j \dots, [i \leq j],$$

we can discard the inequality constraint and represent it instead as

$$\dots !^i \dots !^{i \vee k} \dots,$$

because $i \leq (i \vee k)$. In general, any collection of use-variable inequalities (or uniqueness-attribute constraints) may be eliminated by replacing some of the use variables with propositional formulae over use variables. This insight is the source of Alms’s usage qualifier expressions.

If we follow use types to their logical conclusion, we reach λ^{URAL} (Ahmed et al. 2005), wherein each type is composed of a pretype that describes its representation and a qualifier that gives its usage. Alms does not follow this approach because we insist that qualified types are too verbose for a user-visible type system. Their system’s qualifier lattice includes two more than ours, R for relevant types that allow duplication but not discarding, and L for linear types. This results in a rich and elegant system, but we do not believe R and L would be useful in a language like Alms.

However, there is an interesting correspondence between our kinding rules and their type rules. For example, our product type constructor (\times) has kind $\Pi \alpha^+ . \Pi \beta^+ . \langle \alpha \rangle \sqcup \langle \beta \rangle$, which means that the kind of a product type is the least upper bound of the kinds of its components. The product typing rule in λ^{URAL} enforces a similar constraint, that the qualifier of a product type, ξ , must upper bound the qualifiers of its components τ_1 and τ_2 .

$$\frac{\begin{array}{l} \Delta \vdash \Gamma \rightsquigarrow \Gamma_1 \boxplus \Gamma_2 \quad \Delta \vdash \xi : \text{QUAL} \\ \Delta; \Gamma_1 \vdash v_1 : \tau_1 \quad \Delta \vdash \tau_1 \leq \xi \\ \Delta; \Gamma_2 \vdash v_2 : \tau_2 \quad \Delta \vdash \tau_2 \leq \xi \end{array}}{\Delta; \Gamma \vdash \langle v_1, v_2 \rangle : \xi(\tau_1 \otimes \tau_2)} \text{ (MPAIR).}$$

Vault. DeLine and Fähndrich’s Vault (2001) is a safe, low-level language with support for typestate. It tracks *keys*, which associate static capabilities with the identity of run-time objects, in the same manner that Alms uses existentially-quantified type variables to tie values to capabilities. This allows static enforcement of a variety of protocols. As an example, DeLine and Fähndrich give a tracked version of the Berkeley Sockets API. In previous work on Alms we show how Alms expresses the same interface.

Vault’s treatment of capabilities may be more convenient to use than Alms’s, because while Alms requires explicit threading of capability values, Vault’s key sets are tracked automatically within function bodies. On the other hand, because capabilities in Alms appear as ordinary values, we may combine them using the native intuitionistic logic of Alms’s type system. Instead, Vault must provide a simple predicate calculus for expressing pre- and post-conditions. For more complicated logic, Vault allows embedding capabilities in values, but since the values are untracked, extracting a capability from a value requires a dynamic check. Alms’s type system eliminates the need for such checks for affine values stored in algebraic datatypes, though it also allows dynamic management of affine values by storing them in reference cells.

Notably, Alms can also express Fähndrich and DeLine’s *adoption and focus* (2002).

Sing#. Microsoft’s experimental Singularity operating system is written in Sing#, a high-level systems programming language that extends Spec# (Fähndrich et al. 2006). Sing# has built-in support for *channel contracts*, which are a form of session type providing static checking of communication protocols between device drivers and other services. Unlike more idealistic linear systems, the design acknowledges the need to allow for failure: every protocol implicitly includes branches to close the channel at any point.

Sing# processes do not share memory but can allocate tracked objects on a common *exchange heap*. Only one process has access to an exchange heap object at a given time, but a process may give up access and transmit the object over a channel to another process, which then claims ownership of it.

Alms’s library includes two different implementations of session types supporting different interfaces, and the exchange heap concept is easily expressible as well.

7. Future Work and Conclusion

We already enjoy programming in Alms, but we are not done yet.

Unit subsumption. In §2, we found that adding capabilities to an existing interface often involves wrapping the old version of a function to ignore a new argument of type unit or construct a tuple containing unit for its result. This is unnecessary. While the client outside the abstraction barrier needs to see types that involve the affine capabilities, the implementation has no use for them.

To eliminate much of this noise, we can extend our subtyping relation to take advantage of the fact that unit is, well, a unit:

$$\frac{\Gamma \vdash \tau_2 \text{ singleton}}{\Gamma \vdash \tau_1 <:^\flat \tau_1 \times \tau_2} \quad \frac{\Gamma \vdash \tau_2 \text{ singleton}}{\Gamma \vdash \tau_1 <:^\flat \tau_2 \times \tau_1}$$

$$\frac{\Gamma \vdash \tau_2 \text{ singleton}}{\Gamma \vdash \tau_1 \xrightarrow{\flat} \tau_2 \xrightarrow{\flat} \tau_3 <:^\flat \tau_1 \xrightarrow{\flat} \tau_3}$$

This is implementable via a type erasure technique such as intensional polymorphism (Crary et al. 2002). Not representing compile-time capabilities at run time has performance benefits as well.

Type inference. Alms’s local type inference eliminates most explicit type applications, but needing to annotate all function arguments is irksome. To fix this, we are exploring possibilities for type inference. While we suspect that our limited subtyping should not impede full Damas-Milner-style inference (Damas and Milner 1982), Alms has several idioms that rely on existential types. We are exploring whether an extension for first-class polymorphism, such as HML (Leijen 2009), would be suitable for Alms.

Alms is not finished, but our prototype is at this point usable for experimentation. It is based on a calculus, λ_{ms} , whose type system we have proved sound. While some parts of the type system are complex, we have seen in practice that Alms types are tractable and Alms programs do not look very different from the functional programs to which we are accustomed. It currently implements algebraic datatypes, exceptions, pattern matching, concurrency, and opaque signature ascription. The language is rich enough to express Vault-style typestate, a variety of static and dynamic locking protocols, checked downcasts of one-use functions to unlimited-use functions, session types, and more.

Acknowledgments

We wish to thank J. Daniel Brown, Bryan Chadwick, Matthias Felleisen, Sam Tobin-Hochstadt, Vincent St-Amour, and the anonymous referees for their helpful comments, discussion, and corrections. This research was supported in part by AFOSR grant FA9550-09-1-0110.

References

- A. Ahmed, M. Fluet, and G. Morrisett. A step-indexed model of substructural state. In *ICFP’05*, pages 78–91. ACM, 2005.
- A. Barber. Dual intuitionistic linear logic. Technical Report ECS-LFCS-960347, U. of Edinburgh, 1996.
- P. S. Barth, R. S. Nikhil, and Arvind. M-structures: Extending a parallel, non-strict, functional language with state. In *Proc. FPCA’91*, volume 523 of *LNCS*. Springer, 1991.
- G. M. Bierman. *On Intuitionistic Linear Logic*. PhD thesis, U. of Cambridge, 1993.
- J. Boyland. Checking interference with fractional permissions. In *SAS’03*, volume 2694 of *LNCS*. Springer, 2003.
- K. Crary, S. Weirich, and G. Morrisett. Intensional polymorphism in type-erasure semantics. *JFP*, 12(06):567–600, 2002.
- L. Damas and R. Milner. Principal type-schemes for functional programs. In *POPL’82*, pages 207–212. ACM, 1982.
- R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI’01*. ACM, 2001.
- M. Fähndrich and R. DeLine. Adoption and focus: Practical linear types for imperative programming. In *PLDI’02*. ACM, 2002.
- M. Fähndrich, M. Aiken, C. Hawblitzel, O. Hodson, G. Hunt, J. R. Larus, and S. Levi. Language support for fast and reliable message-based communication in Singularity OS. In *EuroSys’06*, pages 177–190. ACM, 2006.
- S. J. Gay, V. T. Vasconcelos, and A. Ravara. Session types for inter-process communication. Technical Report TR-2003-133, U. of Glasgow, 2003.
- J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- D. Leijen. Flexible types: Robust type inference for first-class polymorphism. In *POPL’09*, pages 66–77. ACM, 2009.
- X. Leroy, D. Doligez, J. Garrigue, D. Rémy, and J. Vouillon. *The Objective Caml system*. INRIA, 3.11 edition, 2008.
- K. Mazurak, J. Zhao, and S. Zdancewic. Lightweight linear types in System F^o. In *TLDI’10*, pages 77–88. ACM, 2010.
- R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML*. MIT, revised edition, 1997.
- B. C. Pierce. *Types and Programming Languages*. MIT, 2002.
- R. Plasmeijer and M. van Eekelen. *Clean Language Report, Version 2.1*. Dept. of Software Technology, U. of Nijmegen, 2002.
- A. Rossberg, C. V. Russo, and D. Dreyer. F-ing modules. In *TLDI’10*, pages 89–102. ACM, 2010.
- M. Steffen. *Polarized Higher Order Subtyping*. PhD thesis, Universität Erlangen-Nürnberg, 1997.
- I. E. Sutherland and G. W. Hodgman. Reentrant polygon clipping. *CACM*, 17(1):32–42, 1974.
- J. A. Tov and R. Pucella. Stateful contracts for affine types. In *ESOP’10*, volume 6012 of *LNCS*, pages 550–569. Springer, 2010.
- E. de Vries, R. Plasmeijer, and D. M. Abrahamson. Uniqueness typing simplified. In *IFL’07*, pages 201–218. Springer, 2008.
- P. Wadler. Is there a use for linear logic? In *PEPM’91*, pages 255–273. ACM, 1991.
- P. Wadler and S. Blott. How to make ad-hoc polymorphism less ad hoc. In *POPL’89*, pages 60–76. ACM, 1989.