Stateful Contracts for Affine Types*

Jesse A. Tov and Riccardo Pucella

Northeastern University, Boston, MA 02115, USA

{[tov](mailto:tov@ccs.neu.edu),[riccardo](mailto:riccardo@ccs.neu.edu)}@ccs.neu.edu

Abstract. Affine type systems manage resources by preventing some values from being used more than once. This offers expressiveness and performance benefits, but difficulty arises in interacting with components written in a conventional language whose type system provides no way to maintain the affine type system’s aliasing invariants. We propose and implement a technique that uses behavioral contracts to mediate between code written in an affine language and code in a conventional typed language. We formalize our approach via a typed calculus with both affine-typed and conventionally-typed modules. We show how to preserve the guarantees of both type systems despite both languages being able to call into each other and exchange higher-order values.

1 Introduction

Substructural type systems augment conventional type systems with the ability to control the number and order of uses of a data structure or operation [20]. Linear type systems [1, 3, 11, 19], for example, ensure that values with linear type cannot be duplicated or dropped, but must be eliminated exactly once. Other substructural type systems refine these constraints. Affine type systems, which we consider here, prevent values from being duplicated but allow them to be dropped: a value of affine type may be used once or not at all.

Affine types are useful to support language features that rely on avoidance of aliasing. One example is session types [6], which are a method to represent and statically check communication protocols. Suppose that the type declared by

$$\mathbf{type}_{\mathcal{A}} \text{ prot} = (\text{int } \text{send} \rightarrow \text{string } \text{recv} \rightarrow \text{unit}) \text{ chan} \quad (1)$$

represents a channel whose protocol allows us to send an integer, then receive a string, and finally end the session. Further, suppose that *send* and *recv* consume a channel whose type allows sending or receiving, as appropriate, and return a channel whose type is advanced to the next step in the protocol. Then we might write a function that takes two such channels and runs their protocols in parallel:

$$\begin{aligned} &\mathbf{let}_{\mathcal{A}} \text{ twice } (c1: \text{prot}, c2: \text{prot}, z: \text{int}): \text{string} \otimes \text{string} = \\ &\quad \mathbf{let} \text{ once } (c: \text{prot}) (_ : \text{unit}) = \\ &\quad \quad \mathbf{let} \ c \quad = \text{send } c \ z \ \mathbf{in} \\ &\quad \quad \mathbf{let} \ (s, _) = \text{recv } c \quad \mathbf{in} \ s \\ &\quad \mathbf{in} \ (\text{once } c1) \ ||| \ (\text{once } c2) \end{aligned} \quad (2)$$

* Our prototype implementation and the full details of our soundness theorem may be found at <http://www.ccs.neu.edu/~tov/pubs/affine-contracts/>.

The protocol is followed correctly provided that $c1$ and $c2$ are *different* channels. Calling $\text{twice}(c, c, 5)$, for instance, would violate the protocol. An affine type system can prevent this.

In addition to session types and other forms of typestate [15], substructural types have been used for memory management [8], for optimization of lazy languages [18], and to handle effects in pure languages [2]. Given this range of features, a programmer may wish to take advantage of substructural types in real-world programs. Writing real systems, however, often requires access to comprehensive libraries, which mainstream programming languages usually provide but experimental implementations often do not. The prospect of rewriting a large library to work in a substructural language strikes these authors as unappealing.

It is therefore compelling to allow conventional and substructural languages to interoperate. We envision complementary scenarios:

- A programmer wishes to import legacy code for use by affine-typed client code. Unfortunately, legacy code unaware of the substructural conditions may duplicate values received from the substructural language.
- A programmer wishes to export substructural library code for access from a conventional language. A client may duplicate values received from the library and resubmit them, causing aliasing that the library could not produce on its own and bypassing the substructural type system’s guarantees.

Our Contributions. We present a novel approach to regulating the interaction between an affine language and a conventionally-typed language and implement a multi-language system having several notable features:

- The non-affine language may gain access to affine values and may apply affine-language functions.
- The non-affine type system is utterly standard, making no concessions to the affine type system.
- And yet, the composite system preserves the affine language’s invariants.

We model the principal features of our implementation in a multi-language calculus that enjoys type soundness. In particular, the conventional language, although it has access to the affine language’s functions and values, cannot be used to subvert the affine type system.

Our solution is to wrap each exchanged value in a software contract [4], which uses *one bit* of state to track when an affine value has been used. While this idea is simple, the details can be subtle.

Design Rationale and Background. Our multi-language system combines two sublanguages with different type systems. The $\mathcal{C}$ (“conventional”) language is based on the call-by-value, polymorphic λ calculus [7, 12] with algebraic datatypes and SML-style *abstype* [10]. The $\mathcal{A}$ (“affine”) language adds affine types and the ability to declare new abstract affine types, allowing us to implement affine abstractions such as session types and static read-write locks.

A program in our language consists of top-level module, value, and type definitions, each of which may be written in either of the two sublanguages. (In the example above (2), the subscripts on **type** _{$\mathcal{A}$} and **let** _{$\mathcal{A}$} indicate the $\mathcal{A}$ language.) Each language has access to modules written in the other language, although they view foreign types through a translation into the native type system. Affine modules are checked by an affine type system, and non-affine modules are checked by a conventional type system. Notably, non-functional affine types appear as abstract types to the conventional type system, which requires no special knowledge about affine types other than comparing them for equality.

In our introductory example, a protocol violation occurs only if the two arguments to *twice* are aliases for the same session-typed channel, which the $\mathcal{A}$ language type system prevents. Problems would arise if we could use the $\mathcal{C}$ language to subvert $\mathcal{A}$ language’s type system non-aliasing invariants. To preserve the safety properties guaranteed by each individual type system and allow the two sublanguages to invoke one another and exchange values, we need to perform run-time checks in cases where the non-affine type system is too weak to express the affine type system’s invariants. Because the affine type system can enforce all of the conventional type system’s invariants, we may dispense with checks in the other direction.

For instance, the affine type system guarantees that an affine value created in an affine module will not be duplicated within the affine sublanguage. If, however, the value flows into a non-affine module, then static bets are off. In that case, we resort to a dynamic check that prevents the value from flowing back into an affine context more than once. Since our language is higher-order, we use a form of higher-order contract [4] to keep track of each module’s obligations toward maintaining the affine invariants.

Our approach to integrating affine and conventional types borrows heavily from recent literature on multi-language interoperability [5, 13]. Our approach borrows from that of Typed Scheme [16, 17] and of Matthews and Findler [9], both of which use contracts to mediate between an untyped, Scheme-like language and a typed language.

2 Example: Taming the Berkeley Sockets API

The key feature of our system is the ability to write programs that safely mix code written in an affine-typed language and a conventionally-typed language. As an example, we develop a small networking library and application, using both of our sublanguages where appropriate.

The Berkeley sockets API is the standard C language interface to network communication [14]. Transmission Control Protocol (TCP), which provides reliable byte streams, is the standard transport layer protocol used by most internet applications (*e.g.*, SMTP, HTTP, and SSH). Setting up a TCP session using Berkeley sockets is a multi-step process (Fig. 1). A client must first create a communication end-point, called a *socket*, via the **socket** system call. It may optionally select a port to use with **bind**, and then it establishes a connection

Fig. 1. States and transitions for TCP (simplified)

```

module  $\mathcal{C}$  Socket : sig
  type socket
  val socket : unit  $\rightarrow$  socket           (*  $\emptyset \Rightarrow \text{initial} *$  *)
  val bind  : socket  $\rightarrow$  port  $\rightarrow$  unit   (*  $\text{initial} \Rightarrow \text{bound} *$  *)
  val listen : socket  $\rightarrow$  unit           (*  $\text{bound} \Rightarrow \text{listening} *$  *)
  val accept : socket  $\rightarrow$  socket         (*  $\text{listening} \Rightarrow \text{connected} \otimes \text{listening} *$  *)
  val send  : socket  $\rightarrow$  string  $\rightarrow$  bool (*  $\text{connected} \Rightarrow \text{connected} \oplus \text{closed} *$  *)
  ... end
  
```

Fig. 2. Selected $\mathcal{C}$ language socket operations, annotated with state transitions

with **connect**. Once a connection is established, the client may **send** and **recv** until either the client or the other side closes the connection.

For a server, the process is more involved: it begins with **socket** and **bind** as the client does, and then it calls **listen** to allow connection requests to begin queuing. The server calls **accept** to accept a connection request. When **accept** succeeds, it returns a *new* socket that is connected to a client, and the old, listening socket is available for further **accept** calls. (For simplicity, we omit error transitions, except for failure of **send** and **recv**.)

Our $\mathcal{C}$ sublanguage provides the interface to sockets shown in Fig. 2. The socket operations are annotated with their pre- and post-conditions, but the implementation detects and signals state errors dynamically. For example, calling **listen** on a socket in state **initial** or calling **connect** on a socket that is already **connected** will raise an exception. If the other side hangs up, **send** and **recv** raise exceptions, but nothing in this interface prevents further communication attempts that are bound to fail.¹

By reimplementing the sockets API in language $\mathcal{A}$, we can use language $\mathcal{A}$'s type system to move the state transition information from comments into the type system itself. For example, we give **listen** in sublanguage $\mathcal{A}$ the type

$$\forall \alpha. \alpha \text{ socket} \rightarrow \alpha \text{ bound} \rightarrow \alpha \text{ listening}, \quad (3)$$

which means that given a socket and evidence that the socket is **bound**, **listen** changes the state to **listening** and returns evidence to that effect. These evidence

¹ This simplifies the Berkeley sockets API by omitting address families, protocols, half-closed sockets, non-blocking IO, etc., but the stateful essence remains.

```

val socket   : unit →  $\exists \alpha. \alpha \text{ socket} \otimes \alpha \text{ initial}$ 
val bind    :  $\forall \alpha. \alpha \text{ socket} \rightarrow \text{port} \rightarrow \alpha \text{ initial} \rightarrow \alpha \text{ bound}$ 
val listen   :  $\forall \alpha. \alpha \text{ socket} \rightarrow \alpha \text{ bound} \rightarrow \alpha \text{ listening}$ 
val accept   :  $\forall \alpha. \alpha \text{ socket} \rightarrow \alpha \text{ listening} \rightarrow (\alpha \text{ listening} \otimes \exists \beta. \beta \text{ socket} \otimes \beta \text{ initial})$ 
val connect :  $\forall \alpha. \alpha \text{ socket} \rightarrow \text{host} \rightarrow \text{port} \rightarrow (\alpha \text{ initial} \oplus \alpha \text{ bound}) \rightarrow \alpha \text{ connected}$ 
val send    :  $\forall \alpha. \alpha \text{ socket} \rightarrow \text{string} \rightarrow \alpha \text{ connected} \rightarrow \alpha \text{ connected}$ 
val recv    :  $\forall \alpha. \alpha \text{ socket} \rightarrow \text{int} \rightarrow \alpha \text{ connected} \rightarrow \text{string} \otimes \alpha \text{ connected}$ 
val close   :  $\forall \alpha. \alpha \text{ socket} \rightarrow \alpha \text{ connected} \rightarrow \text{unit}$ 

```

Fig. 3. The $\mathcal{A}$ language sockets API

tokens are *capabilities*, and the type parameter on each capability ties it to the particular socket whose state it describes. These capabilities have affine type so that when *listen* consumes the *bound* capability, we cannot call *listen* again on the same socket.

We reimplement the sockets API in language $\mathcal{A}$ in terms of the language $\mathcal{C}$ operations. From the vantage of language $\mathcal{A}$, $\mathcal{C}$ function types are mapped to $\mathcal{A}$ function types, but the $\mathcal{C}$ type `Socket.socket` is mapped to an *opaque type* `{Socket.socket}`. Type constructor `{·}` delimits foreign types referenced from the other sublanguage.

We declare a new abstract type for sockets in language $\mathcal{A}$, along with a type to represent each of the states:

```

abstype  $\mathcal{A}$   $\alpha \text{ socket} = \text{Sock of } \{\text{Socket.socket}\}$ 
and  $\alpha \text{ initial}$    qualifier A = Initial
and  $\alpha \text{ bound}$     qualifier A = Bound
and  $\alpha \text{ listening}$  qualifier A = Listening
and  $\alpha \text{ connected}$  qualifier A = Connected
with ... (* operations detailed below *) ... end

```

(4)

Several aspects of this *abstype* declaration bear further explanation:

- Each type has a phantom parameter α , which is used to associate a socket with the type witnessing its state.
- The syntax **qualifier A** on each the state type declares that outside the abstraction boundary, values of those types will appear as affine. Code inside the *abstype* declaration sees that they are ordinary, non-affine data types.
- Because each of the capabilities has only one constructor with no values, they need not be represented at run time.

The $\mathcal{A}$ language sockets interface appears in Fig. 3. The $\mathcal{A}$ sockets implementation relies on delegating to $\mathcal{C}$ language functions. From within $\mathcal{A}$, $\mathcal{C}$ types are viewed through a simple translation: function types, quantified types, and a few base types such as `int` pass through transparently, whereas other types are wrapped opaquely as `Socket.socket` as above. Thus, the type of `Socket.socket` is

```

let $\mathcal{A}$  clientLoop[ $\alpha$ ] (sock:  $\alpha$  socket) (f: string  $\rightarrow$  string) (cap:  $\alpha$  connected) =
  let rec loop (cap:  $\alpha$  connected): unit =
    let (str, cap) = recv sock 1024 cap in
    let cap = send sock (f str) cap in
    loop cap
  in try loop cap with SocketError  $\_ \rightarrow ()$ 

let interface threadFork :> (unit  $\overset{a}{\rightarrow}$  unit)  $\rightarrow$  {thread} $\mathcal{C}$  = threadFork $\mathcal{C}$ 

let rec $\mathcal{A}$  acceptLoop[ $\alpha$ ] (sock:  $\alpha$  socket) (f: string  $\rightarrow$  string) (cap:  $\alpha$  listening): unit =
  let (cap, Pack( $\beta$ , (clientsock, clientcap))) = accept sock cap in
    threadFork (fun ()  $\rightarrow$  clientLoop clientsock f clientcap);
    acceptLoop sock f cap

let $\mathcal{A}$  echoServe (port: int) (f: string  $\rightarrow$  string) =
  let Pack( $\alpha$ , (sock, cap)) = socket () in
  let cap = bind sock port cap in
  let cap = listen sock cap in
    acceptLoop sock f cap

```

Fig. 4. An echo server in language $\mathcal{A}$

becomes $\text{unit} \rightarrow \{\text{Socket.socket}\}$ when viewed from $\mathcal{A}$. Each $\mathcal{A}$ function is a minimal wrapper around its $\mathcal{C}$ counterpart:

```

let $\mathcal{A}$  socket () =
  let sock = Socket.socket $\mathcal{C}$  () in
    in Pack(unit, (Sock[unit] sock, Initial[unit])) as  $\exists \beta. \beta$  socket  $\otimes \beta$  initial

let $\mathcal{A}$  listen[ $\alpha$ ] (Sock sock as s:  $\alpha$  socket) ( $\_$ :  $\alpha$  bound) =
  try Socket.listen $\mathcal{C}$  sock;
    Listening[ $\alpha$ ]
  with IOError msg  $\rightarrow$  raise (StillBound (freezeBound s cap, msg))      (5)

```

For *socket* _{$\mathcal{A}$} , we call *Socket.socket* _{$\mathcal{C}$} to create the new socket, which we wrap in the Sock constructor and pack into an existential with a new initial capability. (The type abstracted by the existential is immaterial; unit will do.) Function *listen* _{$\mathcal{A}$} calls its $\mathcal{C}$ counterpart on the socket and returns a listening capability tied by α to the socket. On failure, the socket is still in state **bound**, so it raises an exception containing the **bound** capability. The remaining functions are equally straightforward, but when we're done, provided we got this trusted kernel correct, we have an $\mathcal{A}$ library that enforces the correct ordering of socket operations.

Calling the various $\mathcal{C}$ socket operations from $\mathcal{A}$ is safe because none has a type that enables it to gain access to an $\mathcal{A}$ language value. Other situations are not as simple. Figure 4 shows an implementation of an echo server in language $\mathcal{A}$. (The working code is included with our prototype implementation on our web site.) The server sends back the data it receives from each client after passing it through an unspecified $\text{string} \rightarrow \text{string}$ function *f*. The main function *echoServe*

creates a socket, binds it to the requested port, and begins to listen. The type system ensures that *echoServe* performs these operations in the right order, and because the capabilities have affine types, it disallows referring to any one of them more than once. Function *echoServe* calls *acceptLoop*, which blocks in *accept* waiting for clients. For each client, it spawns a thread to handle that client and continues waiting for another client. Spawning the thread is where the multi-language interaction becomes tricky.

As in other substructural type systems, $\mathcal{A}$ requires that a function be given a type whose usage (unlimited or affine) is at least as restrictive as any variable that it closes over. Thus far, we have seen only unlimited function types ($\rightarrow$), also written $\multimap$. Language $\mathcal{A}$ also has affine function types, written $\multimap^a$.

The new client capability *clientcap*, returned by *accept*, has affine type β connected. Because the thunk for the new thread, $(\text{fun } () \rightarrow \text{clientLoop clientsock } f \text{ clientcap})$, closes over *clientcap*, it has affine type as well: $\text{unit} \multimap^a \text{unit}$. This causes a problem: To create a new thread, we must pass the thunk to the $\mathcal{C}$ function *threadFork_C*, whose type as viewed from $\mathcal{A}$ is $(\text{unit} \rightarrow \text{unit}) \rightarrow \{\text{thread}\}_{\mathcal{C}}$. Such a type makes *no guarantee* about how many times *threadFork_C* applies its argument. In order to pass the affine thunk to it, we assert that *threadFork_C* has the desired behavior:

$$\text{let interface threadFork} :> (\text{unit} \multimap^a \text{unit}) \rightarrow \{\text{thread}\}_{\mathcal{C}} = \text{threadFork}_{\mathcal{C}} \quad (6)$$

This constitutes a checked assertion that the $\mathcal{C}$ value actually behaves according to the given $\mathcal{A}$ type. This gets the program past $\mathcal{A}$'s type checker, and if *threadFork_C* attempts to apply its argument twice at run time, a dynamic check prevents it from doing so and signals an error.

The two sublanguages can interact in other ways:

- We may call *echoServe_A* from the $\mathcal{C}$ language, passing it a $\mathcal{C}$ function for *f*. This is safe because function *f* has type $\text{string} \rightarrow \text{string}$, and thus can never gain access to an affine value.
- We may use the $\mathcal{A}$ language sockets library from a $\mathcal{C}$ program:

```
letC sneaky () =
  let Pack( $\alpha$ , (sock, cap1)) = socketA () in
  let cap2 = connectA sock "sneaky.example.org" 25 cap1 in
  let cap3 = connectA sock "sneaky2.example.org" 25 cap1 in
  ...
```

(7)

This program passes $\mathcal{C}$'s type checker but is caught when it attempts to reuse the initial capability *cap1* at run time. This misbehavior is detected because *sneaky*'s interaction with $\mathcal{A}$ is mediated by a behavioral contract.

3 Implementing Stateful Contracts

In Findler and Felleisen's formulation [4], a contract is an agreement between two software components, or *parties*, about some property of a value. The *positive*

party produces a value, which must satisfy the specified property. The *negative party* consumes the value and is held responsible for treating it appropriately. Contracts are concerned with catching violations of the property and blaming the guilty party, which may help locate the source of a bug. For first-order values the contract may be immediately checkable, but for functional values nontrivial properties are undecidable, so the check must wait until the negative party applies the function, at which point the negative party is responsible for providing a suitable argument and the positive party for producing a suitable result. Thus, for higher-order functions, checks are delayed until first-order values are reached.

In our language, the parties to contracts are modules, which must be in entirely one language or the other, and top-level functions, which we consider as singleton modules.

Contracts on first-order values check assertions about their arguments, and either return the argument or signal an error. Contracts on functions return functions that defer checking until first-order values are reached. The result of applying a contract should contextually approximate the argument. We represent a contract for a type α as a function taking two parties and a value of type α , and returning a value of the same type α :

$$\text{type } \alpha \text{ contract} = \text{party} \times \text{party} \rightarrow \alpha \rightarrow \alpha \quad (8)$$

A simple contract might assert something about a first-order value:

$$\begin{aligned} &\text{let evenContract (neg: party, pos: party) (x: int) =} \\ &\quad \text{if isEven } x \text{ then } x \text{ else blame pos} \end{aligned} \quad (9)$$

The contract is instantiated with the identities of the contracted parties, and then may be applied to a value. We may also construct contracts for functional values, given contracts for the domain and codomain:

$$\begin{aligned} &\text{let makeFunctionContract}[\alpha, \beta] \text{ (dom: } \alpha \text{ contract, codom: } \beta \text{ contract)} \\ &\quad \text{(neg: party, pos: party) (f: } \alpha \rightarrow \beta \text{) =} \\ &\quad \text{fun (x: } \alpha \text{) } \rightarrow \text{codom (neg, pos) (f (dom (pos, neg) x))} \end{aligned} \quad (10)$$

When this contract is applied to a function, it can perform no checks immediately. Instead, it wraps the function so that, when the resulting function is applied, the domain contract is applied to the actual parameter and the codomain contract to the actual result.

We follow this approach closely, but with one small change—contracts for affine functions are stateful:

$$\begin{aligned} &\text{let makeAffineFunContract}[\alpha, \beta] \text{ (dom: } \alpha \text{ contract, codom: } \beta \text{ contract)} \\ &\quad \text{(neg: party, pos: party) (f: } \alpha \rightarrow \beta \text{) =} \\ &\quad \text{let stillGood = ref true in} \\ &\quad \text{fun (x: } \alpha \text{) } \rightarrow \\ &\quad \quad \text{if !stillGood} \\ &\quad \quad \text{then stillGood } \leftarrow \text{false;} \\ &\quad \quad \text{codom (neg, pos) (f (dom (pos, neg) x))} \\ &\quad \text{else blame neg} \end{aligned} \quad (11)$$

This approach works for functions because we can wrap a function to modify its behavior. But what about for other affine values such as the socket capabilities in Sect. 2? We must consider how non-functional values move between the two sublanguages.

In order to understand the solution, we need to show in greater detail how types are mapped between the two sublanguages. (The rest of the type system appears in the next section.) We define mappings $(\cdot)^{\mathcal{A}}$ and $(\cdot)^{\mathcal{C}}$ from $\mathcal{C}$ types to $\mathcal{A}$ types and $\mathcal{A}$ types to $\mathcal{C}$ types, respectively. Base types such as `int` and `bool`, which may be duplicated without restriction in both languages, map to themselves:

$$(\mathcal{B})^{\mathcal{A}} = \mathcal{B} \qquad (\mathcal{B})^{\mathcal{C}} = \mathcal{B} \qquad (12)$$

Function types convert to function types. $\mathcal{C}$ function types go to unlimited functions in $\mathcal{A}$, and both unlimited and affine $\mathcal{A}$ functions collapse to ordinary ($\rightarrow$) functions in $\mathcal{C}$ (where $\mathbf{q}$ ranges over $\mathbf{a}$ and $\mathbf{u}$):

$$(\tau_1 \rightarrow \tau_2)^{\mathcal{A}} = (\tau_2)^{\mathcal{A}} \multimap (\tau_1)^{\mathcal{A}} \qquad (\sigma_1 \xrightarrow{\mathbf{q}} \sigma_2)^{\mathcal{C}} = (\sigma_1)^{\mathcal{C}} \rightarrow (\sigma_2)^{\mathcal{C}} \qquad (13)$$

Quantified types map to quantified types, but they require renaming because we distinguish type variables between the two languages. In particular, $\mathcal{A}$ language type variables carry usage qualifiers, which indicate whether they may be instantiated to any type or only to unlimited types. (All type variables in Sect. 2 were of the $\mathbf{u}$ kind.)

$$(\forall \alpha. \tau)^{\mathcal{A}} = \forall \beta^{\mathbf{u}}. (\tau_1[\{\beta^{\mathbf{u}}\}/\alpha])^{\mathcal{A}} \qquad (\forall \alpha^{\mathbf{q}}. \sigma)^{\mathcal{C}} = \forall \beta. (\sigma_1[\{\beta\}/\alpha^{\mathbf{q}}])^{\mathcal{C}} \qquad (14)$$

Finally, the remaining types are uninterpreted by the mapping, and merely enclosed in $\{\cdot\}$:

$$(\tau^o)^{\mathcal{A}} = \{\tau^o\}, \text{ otherwise} \qquad (\sigma^o)^{\mathcal{C}} = \{\sigma^o\}, \text{ otherwise} \qquad (15)$$

Values in this class of types are inert: they have no available operations other than passing them back to their native sublanguage, which removes the $\{\cdot\}$. (We take $\{\{\tau\}\}$ to be equivalent to τ .)

This mapping implies that all non-functional, affine types in $\mathcal{A}$ map to opaque types in $\mathcal{C}$.² Since all that the $\mathcal{C}$ language can do with values of opaque type is pass them back to $\mathcal{A}$, we are free to wrap such values when they flow into $\mathcal{C}$ and unwrap them when they return to $\mathcal{A}$. Specifically, when an affine value v passes into $\mathcal{C}$, we wrap it in a λ abstraction, **fun** ($_:$ `unit`) $\rightarrow v$, and wrap that thunk with an affine function contract. If the wrapped value flows back into

² Opaque types may seem limiting, but Matthews and Findler [9] have shown that it is possible, in what they call the “lump embedding,” for each sublanguage to marshal its opaque values for the other sublanguage as desired. In practice, this amounts to exporting a fold to the other sublanguage.

$$\begin{aligned}
\mathcal{CA}[\![\text{int}]\!](n, p) &= id \\
\mathcal{CA}[\![\sigma_1 \stackrel{u}{\multimap} \sigma_2]\!](n, p) &= \text{makeFunctionContract } (\mathcal{AC}[\![\sigma_1]\!], \mathcal{CA}[\![\sigma_2]\!]) (n, p) \\
\mathcal{CA}[\![\sigma_1 \stackrel{a}{\multimap} \sigma_2]\!](n, p) &= \text{makeAffineFunContract } (\mathcal{AC}[\![\sigma_1]\!], \mathcal{CA}[\![\sigma_2]\!]) (n, p) \\
\mathcal{CA}[\![\sigma^o]\!](n, p) &= \text{fun } (v: \sigma^o) \rightarrow \text{makeAffineFunContract} \quad (if \sigma^o \text{ is} \\
&\quad (id, id) (n, p) (\text{fun } () \rightarrow v) \quad \text{affine}) \\
\mathcal{AC}[\![\text{int}]\!](n, p) &= id \\
\mathcal{AC}[\![\sigma_1 \stackrel{a}{\multimap} \sigma_2]\!](n, p) &= \text{makeFunctionContract } (\mathcal{AC}[\![\sigma_1]\!], \mathcal{AC}[\![\sigma_2]\!]) (n, p) \\
\mathcal{AC}[\![\sigma^o]\!](n, p) &= \text{fun } (v: \text{unit} \rightarrow \sigma^o) \rightarrow v () \quad (if \sigma^o \text{ is affine})
\end{aligned}$$

Fig. 5. Type-directed generation of coercions

$\mathcal{A}$, we unwrap it by applying the thunk, which produces a contract error if we attempt unwrapping it more than once.

After type checking, our implementation translates $\mathcal{A}$ modules to $\mathcal{C}$ and wraps all interlanguage variable references with contracts that enforce the $\mathcal{A}$ language's view of the variable. In Fig. 5, we show several cases from a pair of metafunctions $\mathcal{AC}[\![\cdot]\!]$ and $\mathcal{CA}[\![\cdot]\!]$, which perform this wrapping. Metafunction $\mathcal{AC}[\![\cdot]\!]$ produces the coercion for references to $\mathcal{C}$ values from $\mathcal{A}$, and $\mathcal{CA}[\![\cdot]\!]$ is for references to $\mathcal{A}$ values from $\mathcal{C}$. Our formalization does not use this translation, but gives a semantics to the multi-language system directly.

4 Formalization

We model our language with a pair of calculi corresponding to the two sublanguages in the implementation. In this section, we first describe the two calculi independently, and then move on to explain how they interact.

To distinguish the two calculi, we typeset our affine calculus $\lambda^{\mathcal{A}}$ in a **sans-serif font** and our non-affine calculus $\lambda_{\mathcal{C}}$ in a **bold serif font**.

4.1 The Calculi $\lambda^{\mathcal{A}}$ and $\lambda_{\mathcal{C}}$

We model sublanguage $\mathcal{C}$ with calculus $\lambda_{\mathcal{C}}$, which is merely call-by-value System F [7] equipped with singleton modules, each of which for simplicity declares only one name bound to one value. The syntax of $\lambda_{\mathcal{C}}$ appears in Fig. 6, including module names, which are disjoint from variable names. We include integer literals, which serve as first-order values that should pass transparently into the affine subcalculus. A program comprises a mutually recursive collection of modules M and a main expression $\mathbf{e}$. We give only the semantics relevant to modules, as the rest is standard. The expression typing judgment has the form $\Delta; \Gamma \vdash_{\mathcal{C}}^M \mathbf{e} : \tau$, and it carries a module context M , which rule TC-MOD uses to type module expressions. To type a program, we must type each module with rule TM-C; note that the whole module context is available to each module, allowing for

<i>variables</i> $\mathbf{x}, \mathbf{y} \in \text{Var}_{\mathcal{C}}$	TC-MOD
<i>type variables</i> $\alpha, \beta \in \text{TVar}_{\mathcal{C}}$	$\text{module } \mathbf{f} : \tau = \mathbf{v} \in M \quad \cdot \vdash_{\mathcal{C}} \tau$
<i>module names</i> $\mathbf{f}, \mathbf{g} \in \text{MVar}_{\mathcal{C}}$	$\frac{}{\Delta; \Gamma \vdash_{\mathcal{C}}^M \mathbf{f} : \tau}$
<i>programs</i> $\mathbf{P} ::= M \mathbf{e}$	TM-C
<i>module contexts</i> $M ::= \mathbf{m}_1 \dots \mathbf{m}_k$	$\frac{; \cdot \vdash_{\mathcal{C}}^M \mathbf{v} : \tau}{\vdash^M \text{module } \mathbf{f} : \tau = \mathbf{v} \text{ okay}}$
<i>modules</i> $\mathbf{m} ::= \text{module } \mathbf{f} : \tau = \mathbf{v}$	C-MOD
<i>types</i> $\tau ::= \forall \alpha. \tau \mid \alpha$	$\frac{(\text{module } \mathbf{f} : \tau = \mathbf{v}) \in M}{\mathbf{f} \xrightarrow[M]{} \mathbf{v}}$
$\mid \tau \rightarrow \tau \mid \text{int}$	
<i>expressions</i> $\mathbf{e} ::= \mathbf{x} \mid \mathbf{f} \mid \mathbf{e}[\tau] \mid \mathbf{e} \mathbf{e}$	
$\mid \Lambda \alpha. \mathbf{v} \mid [z] \mid \dots$	

Fig. 6. Selected syntax and semantics of $\lambda_{\mathcal{C}}$

recursion. Finally, C-MOD shows that module names reduce to the value of the module.

We model sublanguage $\mathcal{A}$ with calculus $\lambda^{\mathcal{A}}$, which extends $\lambda_{\mathcal{C}}$ with affine types. While $\lambda^{\mathcal{A}}$ includes all of $\lambda_{\mathcal{C}}$, we choose not to embed $\lambda_{\mathcal{C}}$ in $\lambda^{\mathcal{A}}$ to emphasize the generality of our approach, anticipating conventional language features that we do not know how to type in an affine language. The syntax of $\lambda^{\mathcal{A}}$ may be found in Fig. 7. Expressions are mostly conventional: values, which include λ and Λ abstractions, constants, and pairs; variables; application and type application; if expressions; pair construction; and pair elimination. Less conventionally, expressions also include *module names* ($\mathbf{f}$), which reduce to the value of the named module. We define the free variables of an expression in the usual way, but note that this includes only regular variables (*e.g.*, y), not module names (*e.g.*, $\mathbf{g}$), which we assume are distinguished syntactically.

Types include integers, function types with qualifier $\mathbf{q}$, universals, and the syntactically distinguished opaque types, which include type variables, products, and reference cells. Figure 8 defines a lattice on qualifiers, of which there are only two: $\mathbf{u}$ is bottom and $\mathbf{a}$ is top. A qualifier is assigned to each type, with the notation $|\sigma| = \mathbf{q}$. Integers are always assigned the unlimited qualifier $\mathbf{u}$, whereas references always have the affine qualifier $\mathbf{a}$. Function types and type variables are annotated with their qualifiers, and products get the stronger qualifier of either of their components. We define the qualifier of a value context Γ as well, to be the maximum qualifier of any type bound in it; in other words, Γ is affine if *any* variable is affine, but if none is then it is unlimited.

The subtyping relation appears in Fig. 8. It is reflexive and transitive, covariant on both pair components and function codomains, and contravariant on function domains, as usual. Subtyping arises from the qualifier lattice in two ways: an unlimited function may be used where an affine function is expected (but not vice versa), and a universal type whose bound variable has qualifier $\mathbf{a}$ may be instantiated by a type with qualifier $\mathbf{u}$ (but not vice versa).

Figure 8 defines context splitting, which is used by expression typing to distribute affine assumptions to only one use in a term, but unlimited variables to an unlimited number of mentions. When a value context must be split to type

<i>variables</i>	x, y	$\in \text{Var}_{\mathcal{A}}$
<i>qualifiers</i>	q	$\in \{a, u\}$
<i>type variables</i>	α^q, β^q	$\in \text{TVar}_{\mathcal{A}}$
<i>module names</i>	f, g	$\in \text{MVar}_{\mathcal{A}}$
<i>modules</i>	$m ::= \text{module } f : \sigma = v$	
<i>types</i>	$\sigma ::= \text{int} \mid \sigma \stackrel{q}{\multimap} \sigma \mid \forall \alpha^q. \sigma \mid \sigma^\circ$	
<i>opaque types</i>	$\sigma^\circ ::= \alpha \mid \sigma \otimes \sigma \mid \sigma \text{ ref}$	
<i>expressions</i>	$e ::= v \mid x \mid f \mid e e \mid e[\sigma] \mid \text{if } 0 e e \mid \langle e, e \rangle \mid \text{let } \langle x, x \rangle = e \text{ in } e$	
<i>values</i>	$v ::= c \mid \lambda x : \sigma. e \mid \Lambda \alpha^q. v \mid \langle v, v \rangle$	
<i>constants</i>	$c ::= \text{new}[\sigma] \mid \text{swap}[\sigma][\sigma] \mid [z] \mid - \mid (z-) \mid \dots$	
<i>value contexts</i>	$\Gamma ::= \cdot \mid \Gamma, x : \sigma$	
<i>type contexts</i>	$\Delta ::= \cdot \mid \Delta, \alpha^q$	

Fig. 7. Syntax of $\lambda^{\mathcal{A}}$

$$\boxed{q \sqsubseteq q}, \boxed{|\tau| = q}, \boxed{|\Gamma| = q}$$

$$u \sqsubseteq q \quad q \sqsubseteq a \quad |\text{int}| = u \quad |\sigma_1 \stackrel{q}{\multimap} \sigma_2| = q \quad |\forall \alpha^q. \sigma| = |\sigma| \quad |\alpha^q| = q$$

$$|\sigma_1 \otimes \sigma_2| = |\sigma_1| \sqcup |\sigma_2| \quad |\sigma \text{ ref}| = a \quad |\Gamma| = \bigsqcup_{x \in \text{dom}(\Gamma)} |\Gamma(x)|$$

$$\boxed{\sigma <: \sigma}$$

S-REFL

$$\frac{}{\sigma <: \sigma}$$

S-TRANS

$$\frac{\sigma_1 <: \sigma_2 \quad \sigma_2 <: \sigma_3}{\sigma_1 <: \sigma_3}$$

S-PROD

$$\frac{\sigma_1 <: \sigma'_1 \quad \sigma_2 <: \sigma'_2}{\sigma_1 \otimes \sigma_2 <: \sigma'_1 \otimes \sigma'_2}$$

S-ARROW

$$\frac{\sigma'_1 <: \sigma_1 \quad \sigma_2 <: \sigma'_2 \quad q \sqsubseteq q'}{\sigma_1 \stackrel{q}{\multimap} \sigma_2 <: \sigma'_1 \stackrel{q'}{\multimap} \sigma'_2}$$

S-FORALL

$$\frac{q_2 \sqsubseteq q_1 \quad \sigma_1[\beta^{q_2}/\alpha^{q_1}] <: \sigma_2}{\forall \alpha^{q_1}. \sigma_1 <: \forall \beta^{q_2}. \sigma_2}$$

$$\boxed{\Gamma \boxplus \Gamma = \Gamma}$$

$$\frac{}{\cdot \boxplus \cdot = \cdot}$$

$$\frac{\Gamma_1 \boxplus \Gamma_2 = \Gamma_3 \quad |\sigma| = a}{\Gamma_1 \boxplus \Gamma_2, x : \sigma = \Gamma_3, x : \sigma}$$

$$\frac{\Gamma_1 \boxplus \Gamma_2 = \Gamma_3 \quad |\sigma| = a}{\Gamma_1, x : \sigma \boxplus \Gamma_2 = \Gamma_3, x : \sigma}$$

$$\frac{\Gamma_1 \boxplus \Gamma_2 = \Gamma_3 \quad |\sigma| = u}{\Gamma_1, x : \sigma \boxplus \Gamma_2, x : \sigma = \Gamma_3, x : \sigma}$$

Fig. 8. Statics of $\lambda^{\mathcal{A}}$ (qualifiers, subtyping, contexts)

$\Delta; \Gamma \vdash_{\mathcal{A}}^M e : \sigma$	
$\frac{\text{TA-SUBSUME} \quad \Delta; \Gamma \vdash_{\mathcal{A}}^M e : \sigma \quad \sigma <: \sigma'}{\Delta; \Gamma \vdash_{\mathcal{A}}^M e : \sigma'}$	$\frac{\text{TA-LAM} \quad \Delta; \Gamma, x : \sigma \vdash_{\mathcal{A}}^M e : \sigma' \quad \Delta \vdash_{\mathcal{A}} \sigma \quad \Gamma _{\text{FV}(\lambda x : \sigma. e)} = q}{\Delta; \Gamma \vdash_{\mathcal{A}}^M \lambda x : \sigma. e : \sigma \stackrel{q}{\multimap} \sigma'}$
$\frac{\text{TA-TAPP} \quad \Delta; \Gamma \vdash_{\mathcal{A}}^M e : \forall \alpha^q. \sigma' \quad \Delta \vdash_{\mathcal{A}} \sigma \quad \sigma \sqsubseteq q}{\Delta; \Gamma \vdash_{\mathcal{A}}^M e[\sigma] : \sigma'[\sigma/\alpha^q]}$	
$\frac{\text{TA-APP} \quad \Delta; \Gamma_1 \vdash_{\mathcal{A}}^M e_1 : \sigma' \stackrel{q}{\multimap} \sigma \quad \Delta; \Gamma_2 \vdash_{\mathcal{A}}^M e_2 : \sigma'}{\Delta; \Gamma_1 \boxplus \Gamma_2 \vdash_{\mathcal{A}}^M e_1 e_2 : \sigma}$	$\frac{\text{TA-MOD} \quad \text{module } f : \sigma = v \in M \quad \cdot \vdash_{\mathcal{A}} \sigma}{\Delta; \Gamma \vdash_{\mathcal{A}}^M f : \sigma}$
$\frac{\text{TA-NEW}}{\Delta; \Gamma \vdash_{\mathcal{A}}^M \text{new}[\sigma] : \sigma \stackrel{u}{\multimap} \sigma \text{ ref}}$	$\frac{\text{TA-SWAP}}{\Delta; \Gamma \vdash_{\mathcal{A}}^M \text{swap}[\sigma_1][\sigma_2] : (\sigma_1 \text{ ref} \otimes \sigma_2) \stackrel{u}{\multimap} (\sigma_1 \otimes \sigma_2 \text{ ref})}$

Fig. 9. Statics of $\lambda^{\mathcal{A}}$ (selected expressions)

locations $\ell \in \text{Loc}$
values $v ::= \dots \mid \ell$
stores $s ::= \{\ell \mapsto v, \dots, \ell \mapsto v\}$
configurations $C ::= (s, e)$
evaluation contexts $E ::= []_{\mathcal{A}} \mid E[\sigma] \mid E e \mid v E \mid \langle E, e \rangle \mid \langle v, E \rangle \mid \dots$

$C \mapsto_M C$	
(A-NEW)	$(s, \text{new}[\sigma] v) \xrightarrow{M} (s \uplus \{\ell \mapsto v\}, \ell)$
(A-SWAP)	$(s \uplus \{\ell \mapsto v_1\}, \text{swap}[\sigma_1][\sigma_2] \langle \ell, v_2 \rangle) \xrightarrow{M} (s \uplus \{\ell \mapsto v_2\}, \langle v_1, \ell \rangle)$

Fig. 10. Dynamics of $\lambda^{\mathcal{A}}$ (selected rules)

two subexpressions, in an application expression, for example (Fig. 9), variables of affine type are made available to either the operator or operand, but not both.

Selected expression typing rules appear in Fig. 9. Rules TA-LAM and TA-APP are the usual substructural rules for typing λ expressions and applications: for λ expressions, the qualifier q given to the resulting $\stackrel{q}{\multimap}$ type is the qualifier of the context Γ limited to the free variables of the expression; thus, the function is at least as restricted as any values it closes over. The type application rule TA-TAPP requires that a type variable be at least as restrictive as any type with which it is instantiated.

Finally, rule TA-SWAP takes a pair of a σ_1 reference and a σ_2 , and returns a σ_1 and a σ_2 reference. From the operational semantics, a small selection of which appears in Fig. 10, it should be clear that **swap** swaps the σ_2 argument into the location and returns the value previously in the location. Since TA-SWAP does

<i>programs</i> $P ::= M \mathbf{e}$	$\lambda_{\mathcal{C}}$ <i>expressions</i> $\mathbf{e} ::= \dots \mid \mathbf{f}^{\mathbf{g}}$
<i>module contexts</i> $M ::= m_1 \dots m_k$	$\lambda_{\mathcal{C}}$ <i>types</i> $\tau ::= \dots \mid \{\sigma\}$
<i>modules</i> $m ::= \mathbf{m} \mid \mathbf{m}$	$\lambda^{\mathcal{A}}$ <i>expressions</i> $\mathbf{e} ::= \dots \mid \mathbf{f}^{\mathbf{g}}$
$\mid \mathbf{interface} \mathbf{f} :> \sigma = \mathbf{g}$	$\lambda^{\mathcal{A}}$ <i>types</i> $\sigma ::= \dots \mid \{\tau\}$

Fig. 11. New syntax for $\lambda_{\mathcal{C}}^{\mathcal{A}}$

$$\boxed{\vdash P : \tau}, \boxed{\vdash^M m \text{ okay}}$$

$$\frac{\text{PROG} \quad \forall m \in M, \vdash^M m \text{ okay} \quad \cdot; \vdash_{\mathcal{C}}^M \mathbf{e} : \tau}{\vdash M \mathbf{e} : \tau}$$

$$\frac{\text{TM-I} \quad (\mathbf{module} \mathbf{g} : (\sigma)^{\mathcal{C}} = \mathbf{v}) \in M \quad |\sigma| = u}{\vdash^M \mathbf{interface} \mathbf{f} :> \sigma = \mathbf{g} \text{ okay}}$$

$$\boxed{\Delta; \Gamma \vdash_{\mathcal{C}}^M \mathbf{e} : \tau}, \boxed{\Delta; \Gamma \vdash_{\mathcal{A}}^M \mathbf{e} : \sigma}$$

$$\frac{\text{TA-MODC} \quad (\mathbf{module} \mathbf{f} : \tau = \mathbf{v}) \in M \quad \cdot \vdash_{\mathcal{C}} \tau}{\Delta; \Gamma \vdash_{\mathcal{A}}^M \mathbf{f} : (\tau)^{\mathcal{A}}}$$

$$\frac{\text{TC-MODA} \quad (\mathbf{module} \mathbf{f} : \sigma = \mathbf{v}) \in M \quad \cdot \vdash_{\mathcal{A}} \sigma}{\Delta; \Gamma \vdash_{\mathcal{C}}^M \mathbf{f} : (\sigma)^{\mathcal{C}}}$$

$$\frac{\text{TA-MODI} \quad (\mathbf{interface} \mathbf{f} :> \sigma = \mathbf{g}) \in M \quad \cdot \vdash_{\mathcal{A}} \sigma}{\Delta; \Gamma \vdash_{\mathcal{A}}^M \mathbf{f} : \sigma}$$

Fig. 12. New statics for $\lambda_{\mathcal{C}}^{\mathcal{A}}$

not require these two types to be the same, **swap** performs a *strong update*—that is, it may change the type of the value residing in a reference cell. This is why the qualifier given to references must be **a**: if a reference is aliased, then it becomes possible to observe the type change in a way that destroys type safety. This feature of the calculus is a stand-in for the variety of invariants that an affine type system might enforce. In the mixed calculus, $\lambda_{\mathcal{C}}$ may gain access to $\lambda^{\mathcal{A}}$ references. It has no operations available to read or write them, but it must be prevented from passing an aliased reference cell back into $\lambda^{\mathcal{A}}$ where it can cause trouble.

4.2 Mixing It up with $\lambda_{\mathcal{C}}^{\mathcal{A}}$

The primary aim of this work is to construct (type-safe) programs by mixing modules written in an affine language and modules written in a non-affine language, and to have them interoperate as seamlessly as possible. We can then model an affine program calling into a library written in a legacy language, or a conventional program calling into code written in an affine language. In either case, we must ensure that the non-affine portions of the program do not break the affine portions' invariants. As noted in Sect. 3, we accomplish this via run-time checks in the style of higher-order contracts [4].

The additional syntax for mixed programs is in Fig. 11. The main expression in a mixed program is in subcalculus $\lambda_{\mathcal{C}}$. Modules now include $\lambda^{\mathcal{A}}$ modules, $\lambda_{\mathcal{C}}$ modules, and *interface* modules, which are used to assert a $\lambda^{\mathcal{A}}$ type about a $\lambda_{\mathcal{C}}$ module as we saw in Sect. 2.

We add to each subcalculus's expressions a production referring to modules from the other subcalculus. We decorate each such module name with the name of the module in which it appears (*e.g.*, $\mathbf{f}^{\mathbf{g}}$ for a reference to $\lambda_{\mathcal{C}}$ module $\mathbf{f}$ from $\lambda^{\mathcal{A}}$ module $\mathbf{g}$) and use this name as the negative party in contracts regulating the intercalculus boundary, in order to assign blame.

Static Semantics. The type system for the mixed calculus is the union of the type systems for $\lambda^{\mathcal{A}}$ and $\lambda_{\mathcal{C}}$ (Figs. 6, 8, and 9), along with additional typing rules (Fig. 12) for $\lambda^{\mathcal{A}}$ module invocations in $\lambda_{\mathcal{C}}$ expressions and $\lambda_{\mathcal{C}}$ module invocations in $\lambda^{\mathcal{A}}$ expressions.

Rule TC-MODA (Fig. 12) types occurrences of $\lambda^{\mathcal{A}}$ module names in $\lambda_{\mathcal{C}}$ expressions. The rule uses the type conversion function $(\cdot)^{\mathcal{C}}$, defined in Sect. 3 (p. 558) to give a $\lambda_{\mathcal{C}}$ type to the $\lambda^{\mathcal{A}}$ module invocation. Because $\lambda^{\mathcal{A}}$ types are richer than $\lambda_{\mathcal{C}}$ types— $\lambda^{\mathcal{A}}$ function types carry extra information in the qualifier—the conversion loses information, which may need to be recovered through dynamic checks. For example, given a $\lambda^{\mathcal{A}}$ module $\mathbf{g}$ with type $\text{int} \multimap \text{int} \multimap \text{int}$, the conversion rule assigns it the $\lambda_{\mathcal{C}}$ type $\text{int} \rightarrow \text{int} \rightarrow \text{int}$. Calculus $\lambda_{\mathcal{C}}$'s type system cannot enforce that the result of applying $\mathbf{g}$ be applied at most once, which will need to be checked at run time.

For a $\lambda_{\mathcal{C}}$ module with type τ invoked from a $\lambda^{\mathcal{A}}$ expression, we use the module at type $(\tau)^{\mathcal{A}}$. It would be reasonable for TA-MODC to give it any $\lambda^{\mathcal{A}}$ type in the pre-image of the $\lambda^{\mathcal{A}}$ -to- $\lambda_{\mathcal{C}}$ mapping, but $(\cdot)^{\mathcal{A}}$ makes the most permissive, statically safe choice, which is to map all $\lambda_{\mathcal{C}}$ arrows ($\rightarrow$) to the unlimited $\lambda^{\mathcal{A}}$ arrow ($\multimap$). Consider:

- If $\mathbf{f} : \text{int} \rightarrow \text{int}$ in $\lambda_{\mathcal{C}}$, then $\text{int} \multimap \text{int}$ is the right type in $\lambda^{\mathcal{A}}$. There is no reason to limit $\mathbf{f}$ to an affine function type, because $\lambda_{\mathcal{C}}$ does not impose that requirement, and subtyping allows us to use it at $\text{int} \multimap \text{int}$, if necessary.
- If $\mathbf{f} : (\text{int} \rightarrow \text{int}) \rightarrow \text{int}$ in $\lambda_{\mathcal{C}}$, then $(\text{int} \multimap \text{int}) \multimap \text{int}$ will allow the imported function to be passed unlimited functions but not affine functions. This is a safe choice, because $\lambda_{\mathcal{C}}$'s type system does not tell us whether $\mathbf{f}$ may call its argument more than once.

In the latter case, what if the programmer somehow knows that function $\mathbf{f}$ applies its argument at most once, as in the example of *threadFork* _{$\mathcal{C}$} (p. 556)? It should not violate $\lambda^{\mathcal{A}}$'s invariants to pass an affine function to *threadFork* _{$\mathcal{C}$} , but $\lambda^{\mathcal{A}}$ cannot know this. Therefore, rule TA-MODC gives $\lambda_{\mathcal{C}}$ modules a conservative $\lambda^{\mathcal{A}}$ type that requires no run-time checks. We can use an *interface* module to coerce a $\lambda_{\mathcal{C}}$ module's type τ to a more permissive $\lambda^{\mathcal{A}}$ type in the pre-image of τ , and this, too, requires a dynamic check.

Operational Semantics. We extend the syntax of our mixed language with several new forms (Fig. 13). Whereas our source syntax segregates the two subcalculi into separate modules, module invocation reduces to the body of the

module, which leads expressions of both subcalculi to nest at run time. Rather than allow $\lambda^{\mathcal{A}}$ terms to appear directly in $\lambda_{\mathcal{C}}$, and vice versa, we need a way to cordon off terms from one calculus embedded in the other and to ensure that the interaction is well-behaved. We call these new expression forms *boundaries*.

The new run-time syntax includes both boundary expressions ${}^{\sigma}_f \mathbf{AC}_{\mathbf{g}}(\mathbf{e})$ for embedding $\lambda_{\mathcal{C}}$ expressions in $\lambda^{\mathcal{A}}$ and boundary expressions ${}_f \mathbf{CA}_{\mathbf{g}}^{\sigma}(\mathbf{e})$ for embedding $\lambda^{\mathcal{A}}$ expressions in $\lambda_{\mathcal{C}}$. Each of these forms has a superscript σ , written on the $\lambda^{\mathcal{A}}$ side, which represents a contract between the two modules that gave rise to the nested expression. Some contracts, for example int , are fully enforced by both type systems. Other contracts, such as $\text{int} \stackrel{\text{a}}{\circ} \text{int}$, require dynamic checks. The type system guarantees that such a function receives and returns only integers, but this type also imposes an obligation on the negative party to apply the function at most once, which the $\lambda_{\mathcal{C}}$ type system alone does not enforce.

The right subscript of a boundary is a module name in the inner subcalculus, representing the positive party to the contract: It promises that if the enclosed subexpression reduces to a value, then the value will obey contract σ . The left subscript is the negative party, which promises to treat the resulting value properly. In particular, if the contract is affine, then the negative party promises to use the resulting value at most once.

Boundaries first arise when a module in one calculus refers to a module in the other calculus. When the name of a $\lambda_{\mathcal{C}}$ module appears in a $\lambda^{\mathcal{A}}$ term, A-MODC wraps the module name with an $\mathbf{AC}$ boundary, using the $\lambda^{\mathcal{A}}$ -conversion of the module's type τ as the contract. For interface modules, the contract is as declared by the interface, and the name of the interface is the positive party (A-MODI). From the other direction, a $\lambda^{\mathcal{A}}$ module invoked from a $\lambda_{\mathcal{C}}$ expression is wrapped in a $\mathbf{CA}$ boundary by rule C-MODA.

We add evaluation contexts for reduction under boundaries, which means it is now possible to construct a $\lambda_{\mathcal{C}}$ evaluation context with a $\lambda^{\mathcal{A}}$ hole, and vice versa. If the expression under a boundary reduces to a value, it is time to apply the boundary's contract to the value. There are three possibilities:

- Some values, such as integers, always satisfy the contract, so the boundary is discarded.
- Functional values and opaque affine values must have their checks deferred: functions until application time, and opaque values until they pass back into their original subcalculus. For deferred checks, we leave the value in a “sealed” boundary, ${}_f \mathbf{CA}[\ell]_{\mathbf{g}}^{\sigma}(\mathbf{v})$ or ${}^{\sigma}_f \mathbf{AC}[\]_{\mathbf{g}}(\mathbf{v})$, which is itself a value form.
- When a previously sealed opaque value reaches a boundary back to its original subcalculus, both that boundary and the sealed boundary are discarded.

Rule C-WRAP implements contract application for $\lambda^{\mathcal{A}}$ values embedded in $\lambda_{\mathcal{C}}$ expressions, as indicated by metafunction $\text{coerce}_{\mathcal{C}}$. The first case of $\text{coerce}_{\mathcal{C}}$ handles immediate checks, and its second case unseals previously sealed $\lambda_{\mathcal{C}}$ values that have returned home. The second case of $\text{coerce}_{\mathcal{C}}$ seals and *blesses* a $\lambda^{\mathcal{A}}$ value, by allocating a location ℓ , to which it stores a distinguished value **blssd**; it adds this location to the boundary, which marks the sealed value as not yet

$$\begin{array}{ll}
\lambda_{\mathcal{C}} \text{ terms } \mathbf{e} ::= \dots \mid \mathbf{CA}_{\mathbf{f} \mathbf{f}}^{\sigma}(\mathbf{e}) & \lambda^{\mathcal{A}} \text{ terms } \mathbf{e} ::= \dots \mid \mathbf{AC}_{\mathbf{f} \mathbf{f}}^{\sigma}(\mathbf{e}) \\
\lambda_{\mathcal{C}} \text{ values } \mathbf{v} ::= \dots \mid \mathbf{CA}_{\mathbf{f} \mathbf{f}}[\ell]^{\sigma}(\mathbf{v}) & \lambda^{\mathcal{A}} \text{ values } \mathbf{v} ::= \dots \mid \mathbf{AC}_{\mathbf{f} \mathbf{f}}^{\sigma}[\ell](\mathbf{v}) \\
\lambda_{\mathcal{C}} \text{ eval. cxts. } \mathbf{E} ::= \dots \mid \mathbf{CA}_{\mathbf{f} \mathbf{f}}^{\sigma}(\mathbf{E}) & \lambda^{\mathcal{A}} \text{ eval. cxts. } \mathbf{E} ::= \dots \mid \mathbf{AC}_{\mathbf{f} \mathbf{f}}^{\sigma}(\mathbf{E}) \\
\\
\text{configurations } C ::= (s, \mathbf{e}) \mid \mathbf{blame f} & \\
\text{answers } A ::= (s, \mathbf{v}) \mid \mathbf{blame f} & \\
\text{stores } s ::= \{\} \mid s \uplus \{\ell \mapsto \mathbf{v}\} \mid s \uplus \{\ell \mapsto \mathbf{v}\} & \\
\\
(\text{C-ModA}) & (s, \mathbf{f}^{\mathbf{g}}) \xrightarrow{M} (s, \mathbf{CA}_{\mathbf{f} \mathbf{f}}^{\sigma}(\mathbf{f})) \quad (\text{module } \mathbf{f} : \sigma = \mathbf{v}) \in M \\
(\text{A-ModC}) & (s, \mathbf{f}^{\mathbf{g}}) \xrightarrow{M} (s, \mathbf{AC}_{\mathbf{f} \mathbf{f}}^{\tau}(\mathbf{f})) \quad (\text{module } \mathbf{f} : \tau = \mathbf{v}) \in M \\
(\text{A-ModI}) & (s, \mathbf{f}^{\mathbf{g}}) \xrightarrow{M} (s, \mathbf{AC}_{\mathbf{f} \mathbf{f}}^{\sigma}(\mathbf{f}')) \quad (\text{interface } \mathbf{f} : \sigma = \mathbf{f}') \in M \\
(\text{C-Wrap}) & (s, \mathbf{CA}_{\mathbf{f} \mathbf{g}}^{\sigma}(\mathbf{v})) \xrightarrow{M} \text{coerce}_{\mathcal{C}}(s, \sigma, \mathbf{v}, \mathbf{f}, \mathbf{g}) \\
(\text{A-Wrap}) & (s, \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\sigma}(\mathbf{v})) \xrightarrow{M} \text{coerce}_{\mathcal{A}}(s, \sigma, \mathbf{v}, \mathbf{f}, \mathbf{g}) \\
(\text{C-B-A}) & (s, \mathbf{CA}_{\mathbf{f} \mathbf{g}}[\ell]^{\forall \alpha^{\mathbf{q}}. \sigma}(\mathbf{v})[\tau]) \xrightarrow{M} \text{check}(s, \ell, |\sigma|, \mathbf{CA}_{\mathbf{f} \mathbf{g}}^{\sigma[(\tau)^{\mathcal{A}} / \alpha^{\mathbf{q}}]}(\mathbf{v}[(\tau)^{\mathcal{A}}]), \mathbf{f}) \\
(\text{C-}\beta\text{-A}) & (s, \mathbf{CA}_{\mathbf{f} \mathbf{g}}[\ell]^{\sigma_1 \xrightarrow{\mathbf{q}} \sigma_2}(\mathbf{v}_1) \mathbf{v}_2) \xrightarrow{M} \text{check}(s, \ell, \mathbf{q}, \mathbf{CA}_{\mathbf{f} \mathbf{g}}^{\sigma_2}(\mathbf{v}_1 \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\sigma_1}(\mathbf{v}_2)), \mathbf{f}) \\
(\text{A-B-C}) & (s, \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\forall \alpha^{\mathbf{q}}. \sigma}[\ell](\mathbf{v})[\sigma_{\mathbf{a}}]) \xrightarrow{M} (s, \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\sigma[\sigma_{\mathbf{a}} / \alpha^{\mathbf{q}}]}(\mathbf{v}[(\sigma_{\mathbf{a}})^{\mathcal{C}}])) \\
(\text{A-}\beta\text{-C}) & (s, \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\sigma_1 \xrightarrow{\mathbf{q}} \sigma_2}[\ell](\mathbf{v}_1) \mathbf{v}_2) \xrightarrow{M} (s, \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\sigma_2}(\mathbf{v}_1 \mathbf{CA}_{\mathbf{g} \mathbf{f}}^{\sigma_1}(\mathbf{v}_2))) \\
\\
\text{coerce}_{\mathcal{C}}(s, \sigma, \mathbf{v}, \mathbf{f}, \mathbf{g}) = \begin{cases} (s, \lceil z \rceil) & \text{if } \mathbf{v} = \lceil z \rceil \\ (s, \mathbf{v}') & \text{if } \mathbf{v} = \mathbf{g}'^{\tau^{\circ}} \mathbf{AC}_{\lceil \mathbf{f}' \rceil}(\mathbf{v}') \\ (s \uplus \{\ell \mapsto \mathbf{blssd}\}, \mathbf{CA}_{\mathbf{f} \mathbf{g}}[\ell]^{\sigma}(\mathbf{v})) & \text{otherwise} \end{cases} \\
\\
\text{coerce}_{\mathcal{A}}(s, \sigma, \mathbf{v}, \mathbf{f}, \mathbf{g}) = \begin{cases} (s, \lceil z \rceil) & \text{if } \mathbf{v} = \lceil z \rceil \\ \text{check}(s, \ell, |\sigma^{\circ}|, \mathbf{v}', \mathbf{g}') & \text{if } \mathbf{v} = \mathbf{g}'^{\tau^{\circ}} \mathbf{CA}_{\lceil \mathbf{f}' \rceil}[\ell]^{\sigma^{\circ}}(\mathbf{v}') \\ (s, \mathbf{AC}_{\mathbf{f} \mathbf{g}}^{\sigma}[\ell](\mathbf{v})) & \text{otherwise} \end{cases} \\
\\
\text{check}(s, \ell, \mathbf{q}, e, \mathbf{f}) = \begin{cases} (s, e) & \text{if } \mathbf{q} = \mathbf{u} \\ (s' \uplus \{\ell \mapsto \mathbf{dfunct}\}, e) & \text{if } s = s' \uplus \{\ell \mapsto \mathbf{blssd}\} \\ (s, \mathbf{blame f}) & \text{otherwise} \end{cases}
\end{array}$$

Fig. 13. Dynamics of $\lambda_{\mathcal{C}}^{\mathcal{A}}$ (run-time syntax and reduction rules)

used. This corresponds directly to the reference cell allocated by *makeAffineFunContract* in Sect. 3.

Rule A-WRAP implements contracts for $\lambda_{\mathcal{C}}$ values in $\lambda^{\mathcal{A}}$ expressions. Metafunction *coerce _{$\mathcal{A}$}* ’s first case is the same as *coerce _{$\mathcal{C}$}* ’s, and the third case seals a value for deferred checking; it need not allocate a location to track the usage of a $\lambda_{\mathcal{C}}$ value. The third case unseals a previously sealed $\lambda^{\mathcal{A}}$ value on its way back to $\lambda^{\mathcal{A}}$, and this requires checking that an affine value has not been previously unsealed. This step is specified by metafunction *check*, which also has three cases. Unlimited values are unsealed with no check. If an affine value remains blessed, *check* updates the store to mark it “defunct” and returns the unsealed value. If, on the other hand, there is an attempt to unseal a defunct affine value, *check* blames the negative party. This is the key dynamic check that enforces the affine invariant for non-functional values.

Rules C-B-A, C- β -A, A-B-C, and A- β -C all handle sealed abstractions, which are unsealed when they are applied. For sealed $\lambda^{\mathcal{A}}$ abstractions, the seal location ℓ must be checked, to ensure that an affine function or type abstraction is not unsealed and applied more than once. This is the dynamic check that enforces the affine invariant for functions.

Type Soundness. The presence of strong updates means that aliasing a location can result in a program getting “stuck”: if an aliased location is updated at a different type, reading from the alias produces a value of unexpected type. Calculus $\lambda^{\mathcal{A}}$ ’s type system prevents this, but adding $\lambda_{\mathcal{C}}$ means that a $\lambda^{\mathcal{A}}$ value may be aliased outside $\lambda^{\mathcal{A}}$. Our soundness criterion is that no program that gets stuck is assigned a type. In particular, all aliasing of affine values is either prevented by $\lambda^{\mathcal{A}}$ ’s type system or detected by a contract at run time.

In order to prove a Wright-Felleisen-style type soundness theorem [21], we identify precisely what property is preserved by subject reduction. We use an internal type system to track which portions of the store are reachable from $\lambda^{\mathcal{A}}$ values that have flowed into $\lambda_{\mathcal{C}}$. Under this type system, configurations enjoy standard progress and preservation, which allows us to state and prove a syntactic type soundness theorem using the internal type system’s configuration typing judgment $\triangleright^M C : \tau$:

Theorem (Type Soundness). *If $\vdash M e : \tau$ and $(\{\}, e) \mapsto_{M^*} C$ such that configuration C cannot take another step, then C is an answer with $\triangleright^M C : \tau$.*

Our full formalization, including complete definitions of the calculi and proofs, is available at <http://www.ccs.neu.edu/~tov/pubs/affine-contracts/>.

5 Conclusion

Our work is part of an ongoing program to investigate practical aspects of substructural type systems, and this paper describes one step in that program. Here, we have focused on the problem of interaction between substructural and non-substructural code, each governed by its own type system, and explored the

use of higher-order contracts to prevent the conventional language from breaking the substructural language’s invariants. Our answer to the problem at hand naturally raises more questions.

Exceptions. In a production language with a contract system, contract violations should not always terminate the program. Real programs may catch an exception and either try to mitigate the condition that caused it, try something easier instead, or report an error and go on with some other task. To ensure soundness, it suffices to prevent the questionable actions from occurring.

On one hand, we believe that ML-style exceptions should not provide too much difficulty in an affine setting. In our prototype, *try-with* expressions are multiplicative, in the sense that the type environment needs to be split between an expression and its exception handler, not given in whole to both.

On the other hand, we do not know how exceptions or any sort of blame might work in a linear setting—this is one reason why we chose an affine calculus. Terminating the program is problematic because of the implicit discarding of linear values, but catching an exception once part of a continuation containing linear values has been discarded seems even worse. Exceptions in linear languages remain an open question.

Linearity. Our work emphasizes contract-based interaction with affine type systems rather than linear type systems because it remains unclear to us what linear contracts ought to mean. We may want a conventional language to interoperate with a language that (at least sometimes) prohibits discarding values. However, unlike affine guarantees, which are safety properties, relevance guarantees—that a value is used at some point in the future—are a form of liveness property.

One approximation is to consider a contract representing a relevance guarantee to be violated if at any point we can determine that the contract necessarily will be violated. Detecting the violation of such a liveness property is undecidable in general, but tracing garbage collection approximates a liveness property very close to the one we desire. In an idealized semantics, we might garbage collect the store after each reduction step and signal a violation if the seal location of a not-yet-used linear value has become unreachable. In a real implementation, finalizers on linear values could detect discarding. If we detect a violation, we probably could do nothing to prevent it, but at worst we could file a bug report.

Our work suggests that adding substructural libraries to a conventional programming language such as ML does not require a particularly complicated implementation, and our results yield a realistic contract-based design.

Acknowledgments. We wish to thank Daniel Brown, Ryan Culpepper, Jed Davis, Matthias Felleisen, Alec Heller, Sam Tobin-Hochstadt, Aaron Turon, 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

1. Ahmed, A., Fluett, M., Morrisett, G.: L^3 : A linear language with locations. Tech. Rep. TR-24-04, Harvard University (2004)
2. Barendsen, E., Smetsers, S.: Uniqueness typing for functional languages with graph rewriting semantics. *Mathematical Structures in Computer Science* 6(6) (1996)
3. Benton, P.N.: A mixed linear and non-linear logic: Proofs, terms and models. In: Pacholski, L., Tiuryn, J. (eds.) *CSL 1994*. LNCS, vol. 933, pp. 121–135. Springer, Heidelberg (1995)
4. Findler, R.B., Felleisen, M.: Contracts for higher-order functions. In: *ICFP 2002*, pp. 48–59. ACM, New York (2002)
5. Flanagan, C.: Hybrid type checking. In: *POPL 2006*, vol. 41, pp. 245–256. ACM, New York (2006)
6. Gay, S.J., Hole, M.J.: Types and subtypes for client-server interactions. In: Swierstra, S.D. (ed.) *ESOP 1999*. LNCS, vol. 1576, pp. 74–90. Springer, Heidelberg (1999)
7. Girard, J.Y.: *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Ph.D. thesis, Université Paris VI (1972)
8. Jim, T., Morrisett, G., Grossman, D., Hicks, M., Cheney, J., Wang, Y.: Cyclone: A safe dialect of C. In: *Proc. USENIX Annual Technical Conference* (2002)
9. Matthews, J., Findler, R.B.: Operational semantics for multi-language programs. In: *POPL 2007*, vol. 42, pp. 3–10. ACM, New York (2007)
10. Milner, R., Tofte, M., Harper, R., MacQueen, D.: *The Definition of Standard ML*, revised edn. MIT, Cambridge (1997)
11. Plotkin, G.: Type theory and recursion. In: *LICS 1993*, p. 374. IEEE Computer Society, Los Alamitos (1993)
12. Reynolds, J.C.: Towards a theory of type structure. In: Robinet, B. (ed.) *Programming Symposium*. LNCS, vol. 19, pp. 408–423. Springer, Heidelberg (1974)
13. Siek, J.G., Taha, W.: Gradual typing for functional languages. In: *Workshop on Scheme and Functional Programming*, pp. 81–92. ACM, New York (2006)
14. Stevens, W.R.: *UNIX Network programming*. Prentice-Hall, New Jersey (1990)
15. Strom, R., Yemini, S.: Typestate: A programming language concept for enhancing software reliability. *IEEE Transactions on Software Engineering* 12(1) (1986)
16. Tobin-Hochstadt, S., Felleisen, M.: Interlanguage migration: From scripts to programs. In: *OOPSLA 2006*, pp. 964–974. ACM, New York (2006)
17. Tobin-Hochstadt, S., Felleisen, M.: The design and implementation of Typed Scheme. In: *POPL 2007*, pp. 395–406. ACM, New York (2008)
18. Turner, D.N., Wadler, P., Mossin, C.: Once upon a type. In: *FPCA 1995*, pp. 1–11. ACM, New York (1995)
19. Wadler, P.: Linear types can change the world. In: *Programming Concepts and Methods*. In: *Programming Concepts and Methods*, pp. 347–359. North Holland, Amsterdam (1990)
20. Walker, D.: Substructural type systems. In: Pierce, B.C. (ed.) *Advanced Topics in Types and Programming Languages*, ch. 1, pp. 3–44. MIT, Cambridge (2005)
21. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. *Information and Computation* 115(1), 38–94 (1994)