Kindly bent to free us

Gabriel Radanne, Hannes Saffrich, Peter Thiemann

Kindly Bent to Free Us

GABRIEL RADANNE, Inria, Paris
HANNES SAFFRICH, University of Freiburg, Germany
PETER THIEMANN, University of Freiburg, Germany

Systems programming often requires the manipulation of resources like file handles, network connections, or dynamically allocated memory. Programmers need to follow certain protocols to handle these resources correctly. Violating these protocols causes bugs ranging from type mismatches over data races to use-after-free errors and memory leaks. These bugs often lead to security vulnerabilities.

While statically typed programming languages guarantee type soundness and memory safety by design, most of them do not address issues arising from improper handling of resources. An important step towards handling resources is the adoption of linear and affine types that enforce single-threaded resource usage. However, the few languages supporting such types require heavy type annotations.

We present Affe, an extension of ML that manages linearity and affinity properties using kinds and constrained types. In addition Affe supports the exclusive and shared borrowing of affine resources, inspired by features of Rust. Moreover, Affe retains the defining features of the ML family: it is an impure, strict, functional expression language with complete principal type inference and type abstraction. Affe does not require any linearity annotations in expressions and supports common functional programming idioms.

1 INTRODUCTION

A large proportion of systems programming is focused on the proper handling of resources, like file handles, network connections, or dynamically allocated memory. Each of these resources comes with a protocol that prescribes the correct use of its API. For examples, a file handle appears as the result of opening a file. If it was opened for reading, then read operations will succeed, but write operations will fail. Once the handle is closed, it cannot be used for reading or writing, anymore. Dynamic allocation of memory is similar. An API call returns a pointer to a memory area, which can then be read and written to until the area is released by another API call.

In both cases, a resource is created in a certain state and a resource handle is returned to the program. Depending on this state, certain API calls can safely be applied to it. Finally, there is another API call to release the resource, which renders the handle invalid. Taken to the extreme, each API call changes the state so that a different set of API calls is enabled afterwards. Ignoring such life cycle protocols is a common source of errors.

Most type systems provide type soundness and memory safety, but neglect the protocol aspect. Systems that can support reasoning about protocols build on linear types [14] and/or uniqueness types [6]. A value of linear type is guaranteed to be consumed exactly once. That is, a file that has been opened must be closed and memory that has been allocated must be released. A value

Authors' addresses: Gabriel Radanne, Inria, Paris, radanne@informatik.uni-freiburg.de; Hannes Saffrich, University of Freiburg, Germany, saffrich@informatik.uni-freiburg.de; Peter Thiemann, University of Freiburg, Germany, thiemann@informatik.uni-freiburg.de.

XXXX-XXXX/2020/6-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

of unique type is guaranteed to have a single reference to it. Thus, memory can be reused on consumption of the value.

These systems work well if one is prepared to write programs functionally in resource-passing style. In this style, all operations in the resource's API take the resource as a parameter and return it in possibly modified state [1]. In typestate-oriented programming, they would also modify its type [2]. Functional session types represent a popular example [13, 20].

Explicit resource passing places a heavy burden on the programmer and complicates the program structure. For imperative APIs, resource-passing style is not an option at all. To this end, Boyland and Retert [8] proposed the notion of *borrowing* a resource. The idea is that a linear resource can be borrowed to a function call. The function can work with a borrow of the resource, but it cannot release the resource. Only the original owner of the resource has all rights to it and can release it.

The concepts of ownership and borrowing have grown popular over time and they form the foundation of the type system of the Rust language [21], which considers any memory-allocated data structure a resource. Rust supports two kinds of borrows, shared and exclusive ones. Exclusive borrows enable modification of the data structure whereas shared borrows only grant read access. At any given time, either a single exclusive borrow is active or any number of shared borrows can be active. Moreover, Rust makes sure that the lifetime of a borrow is properly contained in the lifetime of its lender.

The design of Rust is geared towards programmers with a low-level imperative programming background, like C or C++. Its management of lifetimes supports the manual way of memory management customary in these languages very well and makes it safe. However, programmers with a background in managed languages feel alienated from the lack of garbage collected data. They would prefer a setting where automatic memory management with garbage collection was the default, but where they could seemlessly switch to safe, manual resource management if that was required. As a concrete example, consider a functional programmer who wants to safely interact with a C library. Invoking a C function is easy via the existing foreign function interface, but managing the underlying resources like malloc'd storage is not: it cannot be left to the garbage collector, but proper release of the storage via calls to `free()` must be ensured by programming conventions.

Our work provides a safe solution to programmers in this situation. We propose an extended type system for ML-like languages that comes with linear and affine types, exclusive and shared borrows, but all that integrated with full principal type inference, garbage collected data, and automatic placement of borrowing regions. In our system, it is a type error to omit the call to release the storage given a suitably typed API for storage allocation.

The most closely related contenders in this design space are Linear Haskell [7], henceforth LH, Quill [24], and ALMS [40]. Compared to LH and Quill, the goals and means are similar as these systems also permit abstraction over the number of uses of values and retain type inference, but the details are different.

(1) Multiplicities in LH and Quill are either linear or unrestricted whereas we also distinguish affine values.
(2) In Affe and in Quill multiplicities are directly attached to the type of a value. For example, in Affe the function type $\alpha \xrightarrow{\text{lin}} \beta$ denotes the type of a *single-use function* that can be called just once, whereas the multiplicities in LH choose between $\alpha \to \beta$ and $\alpha \multimap \beta$ where the latter is a function that promises to *use its argument exactly once.*
(3) Affe makes use of multiplicity contraints (like Quill) and kind subsumption (unlike Quill). Kind subsumption results in significantly simpler, more readable inferred types.

```
1 module File : sig
2   type t: lin
3   val fopen: path → t
4   val write: &!t → string -aff-> unit
5   val close: t → unit
6 end
```

(a) File API

```
1 let main () =
2   let h = File.fopen "foo" in
3   File.write &!h "Hello␣";
4   File.write &!h "world!";
5   File.close h
```

(b) File example

```
1 let main () =
2   let h = File.fopen "foo" in
3   {| File.write &!h "Hello␣" |};
4   {| File.write &!h "world!" |};
5   File.close h
```

(c) File example with regions

Fig. 1. Writing files

(4) Neither LH nor Quill have borrowing whereas Affe supports two flavors: affine (exclusive, mutable) and unrestricted (shared, immutable) borrows.

See Section 7 for further in-depth discussion of these and other related works.

1.1 First examples

As a first, well-known example we consider a simplified API for writing files shown in Fig. 1a. It introduces a linear abstract type `File.t`. A call like `File.fopen "foo"` returns a linear handle to a newly created file, which *must* be released later on with `File.close` as shown in Fig. 1b. Failing to do so is a static type error. To write to the file, we must take an exclusive borrow `&!h` of the handle and pass it to the `File.write` function. Exclusive borrows are affine: they must not be duplicated, but there is no requirement to use them. This affinity shows up in the annotation $\xrightarrow{\texttt{aff}}$ of the second arrow in the type of `File.write`: a partial application like `File.write &!h` captures the affine borrow and hence the resulting function is also affine. It would be an error to use the affine closure twice as in

```
let w = File.write &!h in w "Hello␣"; w "world!" (*type error*)
```

The remaining arrows in the API are unrestricted and we write $\rightarrow$ instead of the explicitly annotated $\xrightarrow{\texttt{un}}$.

Every borrow is restricted to a *region*, i.e., a lexically scoped program fragment from which the borrow must not escape. In Fig. 1b, there are two regions visualized in Fig. 1c, one consisting of Line 3 and another consisting of Line 4. Both are fully contained in the scope of the linear handle `h`, hence we can take one exclusive borrow `&!h` in each region. In both regions the borrow is consumed immediately by passing it to `File.write`. Affe elaborates regions automatically before type inference. Alternatively, programmers may mark regions explicitly (See Section 2.1).

This example demonstrates three features of our system:

(1) type and region inference without annotations in user code (Fig. 1b),
(2) types carry multiplicity annotations in the form of kinds,
(3) resource APIs can be written in direct style as linearity is a property of the type `File.t`.

Direct style means that there is a function like `fopen` that creates and returns a linear resource. In contrast, LH forces programmers to use resource-passing style because, in LH, linearity is a property of a function, rather than a property of a value that restricts the way that value can be handled (as in Affe). An LH API analogous to `File` might provide functions like

- `withFile : path → (handle ⊸ Unrestricted r) → r`, which creates a new file handle and takes a continuation that uses the `handle` linearly, but returns an unrestricted value[1],
- `writeFile : string → handle ⊸ handle`, which returns the transformed resource `handle`, and
- `closeFile : handle ⊸ unit`, which consumes the `handle` by closing the file.

[1]For technical reasons, LH requires the programmer to use a type like `Unrestricted` at this point.

In general, kinds can be polymorphic and constrained. Function application and composition are the archetypical examples for functions exploiting that feature.[2] For application, Affe infers the following type.

```
let app f x = f x
# app : (α -κ-> β) -> (α -κ-> β)
```

The reading of the inferred type is straightforward. If `f` is a κ-restricted function, then so is `app f`. The multiplicities of α and β play no role. As usual in ML-like languages, we implicitly assume prenex quantification by $\forall\kappa\forall\alpha\forall\beta$. Internally, the type checker also quantifies over the kinds of α and β, but the full prefix $\forall\kappa\kappa_1\kappa_2\forall(\alpha : \kappa_1)\forall(\beta : \kappa_2)$ of the type of `app` is only revealed as much as necessary for understanding the type.

For `compose`, Affe infers this type.

```
let compose f g x = f (g x)
# compose : (κ ≤ κ1) ⇒ (β -κ-> γ) -> (α -κ1-> β) -κ-> (α -κ1-> γ)
```

Like in `app`, the multiplicities of the type variables α, β, γ do not matter. However, the multiplicity κ of `f` reappears on the second to last arrow because `compose f` is a closure that inherits `f`'s multiplicity. The multiplicities of `g` and `f` both influence the multiplicity of the last arrow, so we would expect its annotation to be the least upper bound $\kappa \sqcup \kappa_1$. Thanks to subsumption of multiplicities, it is sufficient to assume $\kappa \leq \kappa_1$ and `g`'s actual multiplicity gets subsumed to κ_1. This constraint simplification is part of our type inference algorithm. As before, printing the type scheme only mentions the non-trivial constraint $\kappa \leq \kappa_1$ and omits the prenex quantification over κ, κ_1 as well as the kinds of α, β, γ.

1.2 Contributions

- A polymorphic type system that encodes linearity and affinity with borrowing in lexical regions. Polymorphism covers types and kinds that express multiplicity constraints on the number of uses of a value. This type system is a conservative extension of systems for existing ML-like languages.
- Expressive type soundness theorem with respect to a big-step linearity-aware semantics.
- An extension of the HM(X) framework [26] for constrained type inference to equip the type system with full, principal type inference.
- Soundness proof of the inference algorithm.
- Automatic inference of regions for borrows.
- A prototype implementation of the type inference algorithm, including all constraint simplification and extended with algebraic datatypes and pattern matching, available at https://affe.netlify.com/.

As Affe is built on top of the HM(X) framework, which is a general framework for expressing constraint-based typing and type inference, the extension of our work with features like typeclasses, ad-hoc overloading, traits, etc is possible and orthogonal to the presentation in this paper. While the system is geared towards type inference, it is nonetheless compatible with type annotations and thereby amenable to extensions where type inference may no longer be possible.

2 LINEARITY, AFFINITY, AND BORROWS AT WORK

Affe supports the resource-passing style common in functional encodings of session types (e.g., [29]; see also Appendix A.1 in the supplement) as well as other functional resource handling. But it really shines when manipulating mutable resources like buffers or connection pools using a mix of functional and imperative programming styles. To support this usage pattern of linearity,

[2] Compared to Quill [24] the signatures of application and composition are simpler because Affe supports kind subsumption.

```
module Array : sig
  type (α : κ) t : lin
  val create : (α : un) ⇒ int × α → α t
  val free : (α : aff) ⇒ α t → unit
  val length : &(α t) → int
  val get : (α : un) ⇒ &(α t) × int → α
  val set : (α : aff) ⇒ &!(α t) × int × α → unit
  val map : (&α → β) × &(α t) → β t
  val iter : (α → unit) × α t → unit
end
```

Fig. 2. Linear arrays

Affe relies on the notion of borrowing [8]. Our first example of linear arrays demonstrates simple borrowing and imperative programming; the second example demonstrates reborrowing and the interaction between closures and borrowing by implementing a Sudoko solver based on a hybrid copy-on-write data structure; the third example demonstrates advanced uses of regions with iterators on linear values and the low-level primitives needed to implement them. Further examples are available in Appendix A.

2.1 Imperative programming with linear arrays

The API for mutable linear arrays (Fig. 2) aims to safely handle manual allocation and deallocation of arrays that may contain affine or linear elements. A program would first use `create (n, v)` to create an array of size `n` initialized with value `v`. The value `v` must be unrestricted as it is duplicated to initialize all array elements. To `free` an array the elements must be affine. Thanks to subkinding, the type of `free` is pleasingly simple: any type α whose kind is less than or equal to **`aff`** is acceptable. The `length` function is always applicable. The `get` function is only applicable if the element type is unrestricted as one element is duplicated. To `set` an array element displaces the previous content, which must be at least affine.

The `map` function can transform arrays with arbitrary elements. In particular, it can turn unrestricted elements into linear (affine) ones. It takes a borrow of the input array and returns a newly created output array. As `free`ing requires affine elements, we provide the `iter` function which takes a suitable finalizer and an array with arbitrary elements, which is consumed. Indeed such an iteration is the only way to free an array with linear elements. A real-life API would provide a combination of `map` and `iter` as a "destructive" map that consumes its input array. Assuming a uniform representation, such a destructive map might be implemented by in-place update.

To manage the different kinds of accessing the array we distinguish between constructors, destructors, observers, and mutators. Constructors and destructors like `create` and `free` manipulate the whole array. The constructor `create` yields a linear resource which is consumed by `free`. During the lifetime of the array resource `a`, we can split off *shared borrows* **`&`**`a` that provide a read-only view or *exclusive borrows* **`&!`**`a` for read-write views. Observer functions such as `length` and `get` expect a shared borrow argument, mutator functions such as `set` expect an exclusive borrow.

Each borrow is tied to a region whose lifetime is properly contained in the lifetime of the resource. In a region, we can split off as many shared borrows of a resource as we like, but we can take only one exclusive borrow. In a subsidiary region, we can take shared borrows of any borrow or we can take an exclusive borrow of an exclusive borrow from an enclosing region. Borrows are confined to their regions. Inside the region, shared borrows are unrestricted (**`un`**) whereas exclusive borrows are affine (**`aff`**).

Using the API we can create an array of Fibonacci numbers in an imperative coding style:

```
let mk_fib_array n =
  let a = create (n, 1) in
  for i = 2 to n - 1 do
```

```
4    let x = get (&a, i-1) + get (&a, i-2) in
5    set (&!a, i, x)
6  done;
7  a
8 # mk_fib_array : int ⟶ int Array.t
```

After creation of the array, the presence of a borrow in the for loop prevents access to the "raw" resource inside the loop's body. In particular, the resource cannot be freed through a borrow. Line 4 contains two shared borrows in the same expression which forms a region by itself (recall that shared borrows are unrestricted and may thus be duplicated). These borrows are split off the exclusive borrow used in Line 5 which belongs to the next enclosing region corresponding to the loop body. The whole array can be returned in Line 7 because the borrows are no longer in scope. More precisely, here is is an annotated excerpt with regions explicitly marked by braces `{| ... |}`:

```
3  for i = 2 to n - 1 do {|
4    let x = {| get (&a, i-1) + get (&a, i-2) |} in
5    set (&!a, i, x)
6  |} done;
```

One region consists of the header expression of the `let` in Line 4. It is contained in another region spanning the body of the `for` loop. Affe guarantees that borrows never escape the smallest enclosing region. It employs a system of *indexed kinds* like $\texttt{aff}_r$ and $\texttt{un}_r$ where r is a positive integer that corresponds to the lexical nesting depth of regions. For instance, the type of `&!a` in Line 5 has kind $\texttt{aff}_1$ whereas the type of `&a` in Line 4 has kind $\texttt{un}_2$ and the typing of the inner region is such that types with kind indexes greater than or equal to 2 cannot escape. In the example, borrows cannot escape because they are consumed immediately by `get` and `set`.

2.2 Solving sudokus with hybrid data-structures

This section presents an implementation of a backtracking Sudoku solver using a safe API for persistent arrays that supports both mutable updates and immutable versioning. The implementation showcases safe mixing of resource allocation and deallocation in the presence of exclusive (mutable) and immutable borrows. It also demonstrates two new aspects: the interaction between closures and borrows and the notion of reborrowing.

Recently introduced persistent data structures permit transient mutations where non-linear uses lead to degraded performance [9] or to dynamic and static checks [32]. In particular, persistent Hash-Array-Mapped-Tries (HAMT) have been used with similar APIs in several non-pure functional languages (OCaml, Clojure, …). Affine types help formalize the performance contract between the programmer and the library, while borrows avoid the need to thread state explicitly, as usually required by an API for immutable data types.

Our implementation of a backtracking Sudoku solver abstracts this scenario. The solver maintains a two-dimensional array to represent the state of the game and uses backtracking when there are several choices to proceed. As choice points may be revisited several times, it seems advantageous to select a persistent data structure for the array. However, local changes between choice points may be implemented as cheap in-place mutations.

Fig. 3 contains an API `HYBARRAY` along with an implementation `CowArray` that enables using mutable and immutable modifications to the board through affine types and borrows. The signature differs slightly from the `Array` signature. As our application requires the `get` function, the array elements must be unrestricted, but the structure itself remains linear so as to be implemented in terms of `Array`. The in-place mutation function `set_mut` with type $\texttt{\&!}(\alpha\ \texttt{t}) \times \texttt{int} \times \alpha \to \texttt{unit}$ works on an exclusive borrow whereas the persistent `set` operation has type $\texttt{\&}(\alpha\ \texttt{t}) \times \texttt{int} \times \alpha \to \alpha\ \texttt{t}$. It takes a shared borrow because it only reads from the argument array, but returns a fresh, modified

```
module type HYBARRAY = sig
  include ARRAY
  val set : &(α t) × int × α → α t
  val set_mut : &!(α t) × int × α → unit
end
```

```
module CowArray : HYBARRAY = struct
  include Array
  let set (a, i0, x0) =
    Array.mapi ((fun (i, x) → if i = i0 then x0 else x), a)
  let set_mut = Array.set
end
```

Fig. 3. Signature and Implementation of hybrid arrays

```
val propagate : int → int → &!Matrix.t → int → unit
val solve : int → int → Matrix.t → unit
```

```
1  type board = IntSet.t Matrix.t
2
3  let propagate_line i0 j0 board n =
4    for j = j0+1 to 8 do
5      let cell = Matrix.get (&board, i0 , j) in
6      let cell' = IntSet.remove n cell in
7      Matrix.set_mut (&!board, i0, j, cell')
8    done
9
10 let propagate i j board n =
11   propagate_line i j &!board n;
12   propagate_column i j &!board n;
13   propagate_square i j &!board n
```

```
14 let rec solve i j board =
15   begin if is_solved &board then Matrix.print &board else
16     let (new_i, new_j) = next_pos (i,j) in
17     let try_solution n =
18       let new_board =
19         Matrix.set (&board, i, j, IntSet.singleton n) in
20       propagate i j &!new_board n;
21       if is_valid &new_board then
22         solve new_i new_j new_board
23     in
24     let cell = Matrix.get (&board, i, j) in
25     IntSet.iter try_solution cell;
26   end;
27   free board
```

Fig. 4. Excerpt of the Sudoku solver

structure. The module `CowArray`, also in Fig. 3, contains a very simple implementation of `HYBARRAY` that represents hybrid arrays as regular arrays and uses copy-on-write for persistent modifications. The function `mapi: (int×&`$\alpha \to \beta$`)×&(`α `t) →` β `t` is a simple variation on `Array.map` where the mapping function also takes the position of the element. Recall that `Array.map` always creates a new array for the result.

In the example code, we make use of a two-dimensional version `Matrix` of the `CowArray` data-structure. The only difference is that the API functions `get`, `set`, and `set_mut` now take two index parameters of type `int` instead of one. The internal working is exactly the same.

Our implementation of a Sudoku solver (Fig. 4) represents the board as a 2D-matrix (Line 1). Each cell contains an integer set of type `IntSet.t` that represents admissible solutions so far. This type is immutable, i.e., `IntSet.remove` produces a new value.

The main functions are `solve` and `propagate` with the typings shown on top of Fig. 4. The types of `propagate_line` etc are the same as for `propagate`. From the types, we can see that `solve` takes ownership of the board, whereas `propagate` only takes a mutable borrow. Hence, the `propagate` functions can only modify or read the board, whereas `solve` has full control.

The Sudoku solver `solve` iterates over the cells and tries each possible solution (Line 17). When a value is picked for the current cell, it creates a choice point in `new_board`, where the current cell is updated with an immutable modification (Line 19), and propagate the changes with the `propagate` function. The `propagate` function uses direct mutation through an exclusive borrow of the matrix as it need not preserve the previous version of the board. The implementation of `propagate` is split into three parts for lines, columns, and square, which are all very similar to function `propagate_lines` (Line 3).

As the `board` parameter to the `propagate` function is an exclusive borrow, it should be handled in an affine manner. To pass it safely to the three helper functions, the body of `propagate`

reborrows (i.e., it takes a borrow of a borrow) the board three times in Line 11-Line 13. The function `propagate_line` also contains two reborrows of the exclusive borrow argument `board`, an immutable one (Line 5) and an exclusive one (Line 7). It demonstrates the facility to take immutable borrows from exclusive ones.

The typing ensures that the mutations do not compromise the state at the choice point, because they operate on a new state `new_board` created for one particular branch of the choice. As the `set` function only requires an unrestricted shared borrow, the closure `try_solution` remains unrestricted even though it captures the borrow `&board`. The price is that `try_solution` cannot escape from `&board`'s region. In this example, the inferred region corresponds to the **`begin`**/**`end`** scope. Hence, `try_solution` can be used in the iteration in Line 25. As `board` is linear we must free it outside of the region before returning (Line 27).

While presented for copy-on-write arrays, the API can easily be adapted to other persistent data structures with transient mutability such as Relaxed-Radix Balance Vectors (RRB) [32] or persistent HAMTs [4, 16] to provide a convenient programming style without compromising performance.

2.3 Iterators and regions

In the examples so far, regions do not appear in type signatures. But for certain programming idioms, we want to extend the scope of a region across a function boundary. For instance, how should we fold on an array of linear objects? Here is a first attempt at the type of a fold function:

$$\texttt{val fold} : (\alpha{:}\kappa) \Rightarrow (\alpha \to \beta \xrightarrow{\kappa} \beta) \to \alpha\ \texttt{Array.t} \to \beta \xrightarrow{\texttt{lin}} \beta$$

This type puts no restrictions on the element type of the array, but it requires the `fold` function to consume the array and all its elements. The last function arrow is linear because the array type (from Section 2.1) is linear.

If we want to work on borrows of linear and affine resources, then the typing gets more involved because we must make sure those borrows are not leaked in the result. We obtain the following signature for `bfold`, the borrowing fold operation:

$$\texttt{val bfold} : (\beta{:}\kappa), (\kappa \leq \texttt{lin}_r) \Rightarrow (\texttt{\&}(\texttt{aff}_{r+1}, \alpha) \to \beta \xrightarrow{\texttt{aff}_{r+1}} \beta) \to \texttt{\&}(\kappa_1, \alpha\ \texttt{Array.t}) \to \beta \xrightarrow{\kappa_1} \beta$$

The folded function receives a shared borrow of the element in the array. The typing of the callback ensures that this borrow is neither captured nor returned by the function. This encapsulation is implemented with a universally quantified *kind index variable* r. The signature prescribes the type $\texttt{\&}(\texttt{aff}_{r+1}, \alpha)$ for the shared borrow of the resource with an affine kind at region nesting $r + 1$. The return type of the callback is constrained to kind $\kappa \leq \texttt{lin}_r$. The important part of this constraint is the $_r$ index, which ensures that the callback cannot return the borrowed argument from the more deeply nested scope. The input of the fold is a shared borrow of the array, which ensures that we have the right to share borrows of the inner content and make multiple concurrent folds.

As an easy example, we fold over an array of files `all_files : File.t Array.t` to compute the sum of their sizes:

```
let total_size_of files = bfold (fun f s → File.size f + s) files 0
let total_size = total_size_of &all_files
```

This approach is not sufficient if we want to mutate the elements of the array during iteration. To this end, we need to take an exclusive borrow of the structure to iterate on:

$$\texttt{val iter\_mut} : (\texttt{\&!}(\texttt{aff}_{r+1}, \alpha) \to \texttt{unit}) \to \texttt{\&!}(\kappa_1, \alpha\ \texttt{Array.t}) \to \texttt{unit}$$

While the distinction between mutable and immutable iteration functions seems unfortunate, it is typical of programming with borrows and is also present in the Rust standard library. It enables the programmer to explicitly state how different iterations may be composed and optimized. It also

enables different implementations such as using parallel iterations in the immutable case. Affe's region variables ensure that the content iterated on can never be leaked outside of the iteration function. This pattern is essential in many use cases of linearity such as pools of linear objects (see Appendix A.2).

To close this discussion, let's see which primitives are needed to implement functions like `bfold` and `iter_mut`. The naive sequential implementations of both functions boil down to a loop over the index range of the array:

```
let rec bfold_helper f a z i =
  if i < 0 then z
  else bfold_helper f &&a (f (get_sb &&a i) z) (i-1)
let bfold f a z =
  let l = length &&a - 1 in
  bfold_helper f &&a z l
```

```
let rec iter_helper f a i =
  if i < 0 then ()
  else (f (get_eb &&!a i); iter_helper f &&!a (i-1))
let iter_mut f a =
  let l = length &&a - 1 in
  iter_helper f &&!a l
```

Observe that we are proposing two different primitives

- `get_sb : &(κ, α Array.t) → int` $\xrightarrow{\kappa}$ `&(κ, α)`
  to get a shared borrow from a shared borrow of an array and
- `get_eb : &!(κ, α Array.t) → int` $\xrightarrow{\kappa}$ `&!(κ, α)` `(* Unsafe! *)`
  to get an exclusive borrow from an exclusive borrow of an array.

These primitives have the same underlying implemention (the same as `get` from Section 2.1). Their types arise from the intuition that the borrow of a structure should entitle to borrows of its substructures, roughly, the borrow of an array could be considered as array of borrows. However, only `get_sb` is safe: the shared borrow of the array entitles us to obtain shared borrows for the elements *in the same region* as the shared borrow for the array freezes the array inside the region. This freeze extends to the elements because lifetime of the array fully overlaps with the lifetime of its elements. Considering `get_eb`, we see that we may obtain *two different exclusive borrows* of the same array element inside a region. Clearly, the exclusive borrow for the element should live in a nested region where the array is not accessible. Hence, the safe alternative is to use a different function for obtaining borrows of elements

$$\texttt{with\_eborrow} : (\beta{:}\kappa), (\kappa \leq \mathbf{lin}_r) \Rightarrow \texttt{\&!}(\kappa_1, \alpha\ \texttt{Array.t}) \to \texttt{int} \to (\texttt{\&!}(\mathbf{aff}_{r+1}, \alpha) \to \beta) \xrightarrow{\kappa_1} \beta$$

Hence, the helper function for `iter_mut` should read like this

```
let rec iter_helper f a i =
  if i < 0 then ()
  else (with_eborrow &!a i f; iter_helper f &!a (i-1))
```

which also explains the occurrence of $\mathbf{aff}_{r+1}$ in the type of `iter_mut`.

We conclude that borrows of datastructures create the need for differently typed access functions that are tailored for different use cases. The argumentation whether such an access function is safe is sometimes subtle and gives rise to nonobvious types.

3 THE AFFE LANGUAGE

Affe is a minimal ML-like language with let-polymorphism and abstract types. Its type system manages linearity and borrowing using kinds and (kind-) constrained types. For ease of presentation, we consider a simplified internal language.

- Pattern matching is demonstrated on pairs rather than algebraic datatypes.
- There are separate operators for borrowing and reborrowing (taking the borrow of a borrow); the surface language unifies these operators using ad-hoc polymorphism / typeclasses.
- Regions are explicit in the code and must be annotated using the algorithms presented in Section 3.2.
- Regions are identified using nesting levels instead of region variables.

Expressions

$$
\begin{aligned}
e ::= {} & c \mid x \mid (e\ e') \mid \lambda x.e \mid \texttt{let}\ x = e\ \texttt{in}\ e' \\
& \mid (e, e') \mid \texttt{match}_\phi\ x, y = e\ \texttt{in}\ e' && \text{(Pairs)} \\
& \mid \{\!|e|\!\}^n_{\{x \mapsto b\}} && \text{(Region)} \\
& \mid \&^b x \mid \&\&^b x && \text{(Borrows)} \\
& \mid \texttt{create} \mid \texttt{observe} \\
& \mid \texttt{update} \mid \texttt{destroy} && \text{(Resources)} \\
b ::= {} & \mathbf{U} \mid \mathbf{A} && \text{(Borrow specification)} \\
\phi ::= {} & \mathrm{id} \mid \&^b && \text{(Match Specification)}
\end{aligned}
$$

Types

$$
\begin{aligned}
\tau ::= {} & \alpha \mid \tau \times \tau' \mid \mathrm{T}\ \overline{\tau} && \text{(ML types)} \\
& \mid \tau \xrightarrow{\mathrm{k}} \tau' && \text{(Function types)} \\
& \mid \&^b(k, \tau) && \text{(Borrowed Type)}
\end{aligned}
$$

Kinds

$$
\begin{aligned}
k ::= {} & \kappa \mid Q_n \qquad \forall n \in \mathbb{N} \cup \{\infty\} && \text{(Kinds)} \\
Q ::= {} & \mathbf{U} \mid \mathbf{A} \mid \mathbf{L} && \text{(Quality)}
\end{aligned}
$$

Constrained type and kind schemes

$$
\begin{aligned}
C ::= {} & \overline{(k \le k')} && \text{(Constraints)} \\
\sigma ::= {} & \forall\overline{\kappa}\forall\overline{(\alpha : k)}.(C \Rightarrow \tau) && \text{(Type scheme)} \\
\theta ::= {} & \forall\overline{\kappa}.(C \Rightarrow \overline{k} \to k) && \text{(Kind scheme)}
\end{aligned}
$$

Fig. 5. Syntax

The rest of this section formalizes Affe: the syntax (Section 3.1), the statics in terms of a region annotation pass (Section 3.2) and syntax-directed typing (Section 3.3), and a dynamics that is linearity- and resource-aware (Section 3.4).

3.1 Syntax

Figure 5 defines the syntax of Affe. Expressions are as usual in an ML-like language. The novel aspects are match specifications, regions, and borrows.

A borrow $\&^b x$ is always taken from a variable x. The *borrow annotation* b indicates whether the borrow is exclusive/affine (**A**) or shared/unrestricted (**U**). If x is already borrowed, then we need to use the reborrow expression $\&\&^b x$. A region $\{\!|e|\!\}^n_{\{x \mapsto b\}}$ is annotated with its nesting n, the variable x that may be borrowed in e, and the kind of borrow b. A match is indexed with a *match specification* ϕ that indicates whether the match operates on borrows ($\phi = \&^b$) or not ($\phi = id$). We consider four primitive operations to manipulate resources: `create`, `observe`, `update` and `destroy`. They serve a prototypes to demonstrate the typing and handling of resources. For a concrete type of resource, there are further arguments and perhaps different versions of the operations. But the typing and behavior of the operations is analogous to the prototype operations. Moreover, `observe` and `update` serve as eliminators for borrow types.

Many types are indexed with kinds. A kind k is either a kind variable κ or a constant (linear **L**, affine **A**, or unrestricted **U**) indexed by a nesting level $n \in \mathbb{N} \cup \{\infty\}$.

A type τ is either a type variable, a pair type, a function type indexed by a kind, a type application $\mathrm{T}\ \overline{\tau}$ of an abstract type constructor T, or a borrowed type $\&^b(k, \tau)$.

Type schemes σ add quantification over kind variables κ, kinded type variables $(\alpha : k)$, and constraints C to a type, where a constraint is a list of inequalities over kinds.

Abstract type constructors possess kind schemes θ which relate the kinds of the type constructors' arguments to the kind of the constructed type.

Generally, we write lists with an overbar and (sometimes) an index as in $\overline{\tau_i}$.

3.2 Automatic region annotation

In the surface language (Section 2) region annotations are optional. In the internal language, regions must be syntactically explicit and annotated with a nesting index and a scoped variable. This

section defines a transformation $p \rightsquigarrow p'$ which automatically inserts region annotations in programs. The input p is a program with optional region annotations of the form $\{\!|e|\!\}$. The output p' is a program with explicit annotations of the form $\{\!|e|\!\}^n_{\{x \mapsto b\}}$ such that no borrow occurs outside a region. We give an informal presentation of our code transformation and defer the complete definition to Appendix B. This code transformation aims to find, for each borrow $\&^b x$, the biggest region satisfying the following rules:

(1) The region should contain at least $\&^b x$.
(2) The region must be contained in the scope of x.
(3) An exclusive borrow $\&^{\mathrm{A}} x$ should never share a region with any other borrow of x.
(4) The variable x cannot occur in the region of $\&^b x$.

The transformation starts from each borrow $\&^b x$ and grows its associated region until enlarging it would include the binding for x or lead to a conflicting use of x or a conflicting borrow for x. As an example, consider the following program:

$$\begin{array}{l} \lambda a.\ \mathtt{let}\ x = (f\ \&^{\mathrm{A}} a)\ \mathtt{in} \\ \quad g\ (\&^{\mathrm{A}} x); \\ \quad f\ (\&^{\mathrm{U}} x)\ (\&^{\mathrm{U}} x) \end{array} \quad \rightsquigarrow \quad \begin{array}{l} \lambda a. \{\!|\ \mathtt{let}\ x = (f\ \&^{\mathrm{A}} a)\ \mathtt{in} \\ \quad \{\!|g\ (\&^{\mathrm{A}} x)|\!\}^2_{\{x \mapsto \mathrm{A}\}}; \\ \quad \{\!|f\ (\&^{\mathrm{U}} x)\ (\&^{\mathrm{U}} x)|\!\}^2_{\{x \mapsto \mathrm{U}\}} \\ |\!\}^1_{\{a \mapsto \mathrm{A}\}} \end{array}$$

As variable a only has one borrow, its region covers its whole lexical scope. Variable x has multiple conflicting borrows and requires more consideration. We place a first region around the exclusive borrow and its function application, and a second region around both shared borrows. This placement of region is optimal: making any region bigger would cause a borrowing error. In particular, it is essential to place the second borrow around *both* occurrence of $(\&^{\mathrm{U}} x$: we want the function to receive both borrows, but its result must not contain any borrow (otherwise it would escape). Region indices are assigned after placement, so it is trivial to ensure well-nested regions where the inner regions have higher indices than the outer ones.

Programmers may also annotate regions explicitly. The transformation considers an annotation as an upper bound on the contained regions. In the following program, a manual annotation has been inserted to ensure no borrow enters the reference r:

$$\begin{array}{l} \mathtt{let}\ r = \mathrm{ref}\ \ 0\ \mathtt{in} \\ \lambda a.\ set\ r\ \{\!|g\ (\&^{\mathrm{U}} a)|\!\}; \\ \quad f\ (\&^{\mathrm{U}} a) \end{array} \quad \rightsquigarrow \quad \begin{array}{l} \mathtt{let}\ r = \mathrm{ref}\ \ 0\ \mathtt{in} \\ \lambda a.\ set\ r\ \{\!|g\ (\&^{\mathrm{U}} a)|\!\}^1_{\{a \mapsto \mathrm{U}\}}; \\ \quad \{\!|f\ (\&^{\mathrm{U}} a)|\!\}^1_{\{a \mapsto \mathrm{U}\}} \end{array}$$

The rules allow merging the two regions around the borrows $\&^{\mathrm{U}} a$. However the explicit annotation indicates that the region should stay around the closure passed as argument of *set*. This feature is useful to control captures by imperative APIs.

The code transformation is purely syntactic and must be used before typing. It only produces well-nested annotations: if $\{\!| \ldots |\!\}^n_b$ is nested inside $\{\!| \ldots |\!\}^{n'}_{b'}$, then $n > n'$. Furthermore, there is at most one region per borrow, and exactly one region per exclusive borrow. In the rest of this article, we assume that all terms have been annotated by this code transformation and respect these properties.

3.3 Typing

To avoid distraction, this section focuses on the essential and novel parts of the type system. A complete description is available in Appendix D. Here we only discuss the following judgments:

$$\textsc{Lat-UAL}\quad \mathbf{U} \le \mathbf{A} \le \mathbf{L}$$

$$\textsc{Lat-Level}\quad \frac{Q \le Q' \qquad n \le n'}{Q_n \le_{\mathcal{L}} Q'_{n'}}$$

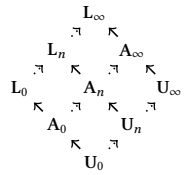

Fig. 6. Lattice ordering – $k \le_{\mathcal{L}} k'$

$$\begin{array}{lll}
\Gamma ::= \cdot \mid \Gamma; B & \text{(Environments)} \\
B ::= \emptyset & \text{(Empty)} \\
\quad \mid (\alpha : \theta) & \text{(Types)} \\
\quad \mid (x : \sigma) & \text{(Variables)} \\
\quad \mid [x : \sigma]^n_b & \text{(Suspended)} \\
\quad \mid (x \div \sigma)^k_b & \text{(Borrows)}
\end{array}$$

Fig. 7. Type environments

$$\begin{array}{rrcll}
(\sigma \le \mathbf{U}_\infty) \vdash_e & (x : \sigma) & = & (x : \sigma) \ltimes (x : \sigma) & \text{(Both)} \\
\cdot \vdash_e & (x \div \sigma)^k_{\mathrm{U}} & = & (x \div \sigma)^k_{\mathrm{U}} \ltimes (x \div \sigma)^k_{\mathrm{U}} & \text{(Borrow)} \\
\cdot \vdash_e & B_x & = & B_x \ltimes \emptyset & \text{(Left)} \\
\cdot \vdash_e & B_x & = & \emptyset \ltimes B_x & \text{(Right)} \\
\cdot \vdash_e & (x : \sigma) & = & [x : \sigma]^n_b \ltimes (x : \sigma) & \text{(Susp)} \\
(b' \le b) \vdash_e & (x \div \sigma)^k_b & = & [x : \sigma]^n_{b'} \ltimes (x \div \sigma)^k_b & \text{(SuspB)} \\
\cdot \vdash_e & [x : \sigma]^{n'}_b & = & [x : \sigma]^n_{\mathrm{U}} \ltimes [x : \sigma]^{n'}_b & \text{(SuspS)}
\end{array}$$

Fig. 8. Splitting rules for bindings – $C \vdash_e B = B_l \ltimes B_r$

$C \mid \Gamma \vdash_{\mathrm{s}} e : \tau$ — Expression e has type τ in environment Γ under constraints C.
$C \mid \Gamma \vdash_{\mathrm{s}} \tau : k$ — Type τ has kind k in environment Γ under constraints C.
$D \vdash_{\mathrm{e}} C$ — Constraint D entails constraint C.
$D =_{\mathrm{e}} C$ — Constraints C and D are equivalent.

Kinds and constraints. Affe uses kinds and constrained types to indicate linear and affine types. A kind k is either a kind variable, κ, or a constant Q_n. The quality Q describes the use pattern of the type: unrestricted **U**, affine **A**, or linear **L**. The level $n \in \mathbb{N} \cup \{\infty\}$ describes the nested regions in which the value can be used. Level 0 refers to the top-level scope outside any region; we often elide it and write **A** for $\mathbf{A}_0$. Level ∞ refers to an empty region that is infinitely nested. For instance, the constraint $(\kappa \le \mathbf{U}_\infty)$ indicates that κ must be unrestricted, but can be local to a region. Kinds form a lattice described in Fig. 6. Unrestricted values may be used where affine ones are expected and affine ones are less restricted than linear ones as reflected in Lat-UAL. Values usable at level n may be used at any more deeply nested level n' as defined in the Lat-Level axioms. Constraints are conjunctions of inequality constraints over this kind lattice, i.e., they specify upper or lower bounds for kind variables or relate two kind variables.

Environments and bindings. Affe controls the use of variables by supporting new modes of binding in type environments Γ, as defined in Fig. 7. Environments contain standard bindings of type variables to kind schemes, $(\alpha : \theta)$, value bindings $(x : \sigma)$, but also suspended and borrow bindings. A suspended binding, $[x : \sigma]^n_b$, indicates that x is earmarked for a borrowing use in a nested region marked with x but cannot be used directly. A borrow binding, $(x \div \sigma)^k_b$, replaces such a suspended binding on entry to the x-region. It indicates that the borrow $\&^b x$ can be used directly. Kind k restricts the lifetime of the borrow to the region (see rules Region and BorrowBinding in Fig. 9 and the upcoming discussion of these rules).

Constraints on an environment control substructural properties by restricting the types of variables. The constraint $(\Gamma \le k)$ stands for the conjunction of $(\sigma \le k)$, for all $(x : \sigma)$ in Γ, which in turn means that $k' \le k$ where k' is the kind of the variable's type scheme σ. Borrow bindings follow the same rules, but suspended bindings are forbidden in an environment Γ constrained like that. This intuitive explanation is sufficient to understand the Var and Abs rules shown in Fig. 9.

$$\textsc{Instance}\ \frac{\sigma = \forall\overline{\kappa_i}\forall(\overline{\alpha_j : k_j}).\, C \Rrightarrow \tau \qquad \psi = [\overline{\kappa_i \mapsto k_i}, \overline{\alpha_j \mapsto \tau_j}]}{\psi(C), \psi(\tau) = \mathrm{Inst}(\Gamma, \sigma)}$$

$$\textsc{Var}\ \frac{(x : \sigma) \in \Gamma \qquad C_x, \tau_x = \mathrm{Inst}(\Gamma, \sigma) \qquad C \vdash_{\mathrm{e}} C_x \wedge (\Gamma \backslash \{x\} \leq \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\mathrm{s}} x : \tau_x}$$

$$\textsc{Abs}\ \frac{C \mid \Gamma; (x : \tau_2) \vdash_{\mathrm{s}} e : \tau_1 \qquad C \vdash_{\mathrm{e}} (\Gamma \leq k)}{C \mid \Gamma \vdash_{\mathrm{s}} \lambda x.e : \tau_2 \xrightarrow{\mathrm{k}} \tau_1}$$

$$\textsc{App}\ \frac{C \mid \Gamma_1 \vdash_{\mathrm{s}} e_1 : \tau_2 \xrightarrow{\mathrm{k}} \tau_1 \qquad C \mid \Gamma_2 \vdash_{\mathrm{s}} e_2 : \tau_2' \qquad C \vdash_e \Gamma = \Gamma_1 \bowtie \Gamma_2 \qquad C \vdash_{\mathrm{e}} (\tau_2' \leq \tau_2)}{C \mid \Gamma \vdash_{\mathrm{s}} (e_1\ e_2) : \tau_1}$$

$$\textsc{Region}\ \frac{[x : \tau_x]_b^n \in \Gamma \qquad C \vdash_e \Gamma \rightsquigarrow_n^x \Gamma' \qquad C \mid \Gamma' \vdash_{\mathrm{s}} e : \tau \qquad C \vdash_{\mathrm{e}} (\tau \leq \mathbf{L}_{n-1})}{C \mid \Gamma \vdash_{\mathrm{s}} \{\!|e|\!\}^n_{\{x \mapsto b\}} : \tau}$$

$$\textsc{Borrow}\ \frac{(x \div \sigma)_b^k \in \Gamma \qquad C_x, \tau_x = \mathrm{Inst}(\Gamma, \sigma) \qquad C \vdash_{\mathrm{e}} C_x \wedge (\Gamma \backslash \{x\} \leq \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\mathrm{s}} \&^b x : \&^b(k, \tau_x)}$$

$$\textsc{Reborrow}\ \frac{C \mid \Gamma \vdash_{\mathrm{s}} x : \&^b(k, \tau)}{C \mid \Gamma \vdash_{\mathrm{s}} \&\&^b x : \&^b(k, \tau)}$$

$$\textsc{BorrowBinding}\ \frac{C \vdash_{\mathrm{e}} (b_n \leq k) \wedge (k \leq b_\infty) \qquad b \in \{\mathbf{U}, \mathbf{A}\}}{C \vdash_e [x : \sigma]_b^n \rightsquigarrow_n^x (x \div \sigma)_b^k}$$

Fig. 9. Selected typing rules ($C \mid \Gamma \vdash_{\mathrm{s}} e : \tau$) and borrowing rules ($C \vdash_e \Gamma \rightsquigarrow_n^x \Gamma'$)

Rule Var looks up the type scheme of the variable x in the environment Γ and instantiates it with $\mathrm{Inst}(\Gamma, \sigma)$. Instantiation follows the HM(X) formulation and takes as input a scheme σ and an environment Γ and returns a constraint C and a type τ. The rule also checks that the other bindings in Γ can be safely discarded by imposing the constraint $(\Gamma \backslash \{x\} \leq \mathbf{A}_\infty)$. It enforces that all remaining bindings (except x) are affine or unrestricted and can therefore be discarded.

Rule Abs ensures the kind annotation on the arrow type $(\tau_2 \xrightarrow{\mathrm{k}} \tau_1)$ reflects the restrictions on captured variables via the constraint $(\Gamma \leq k)$. If, for instance, any binding in Γ is affine, it gives rise to the constraint $(\mathbf{A}_n \leq k)$ and the arrow kind is at least affine at nesting level n. Capturing a borrow is perfectly fine: the kind of the borrow is also a lower bound of the arrow kind k which restricts the closure to the region of the borrow. Capturing a suspended binding is forbidden.

Copying and Splitting. The App typing rule in Fig. 9 demonstrates how Affe deals with duplication and dropping of values. The splitting $C \vdash_e \Gamma = \Gamma_1 \bowtie \Gamma_2$ in the rule decomposes the type environment Γ in two parts, Γ_1 and Γ_2, which are used to typecheck the components of the application.

Fig. 8 shows the action of splitting rules on single bindings. If x's type is unrestricted, rule Both indicates that we can duplicate it. Similarly, unrestricted borrows can be duplicated with rule Borrow. Left and Right rules are always applicable and move a binding either to the left or right environment. The rules Susp, SuspB and SuspS split off suspended bindings to the left while conserving access to the binding on the right. A suspended binding can later be turned into a borrow inside a region. Splitting of suspended bindings is asymmetric. It must follow the order of execution from left to right, which means that a resource can be used first as a borrow on the left and then later as a full resource on the right. The Susp rule works with a full resource, rule SuspB with a borrow, and rule SuspS with a suspended binding.

Splitting applies whenever an expression has multiple subexpressions: function applications, let bindings and pairs. In the expression `let` a = create 8 x `in` f $\{\!|\text{length } \&^{\mathrm{U}} a|\!\}_{\{a \mapsto \mathrm{U}\}}$ a, the rule Susp splits off a borrow from the resource a to use it in the left argument. As usual, a borrow cannot be active in the same scope as its resource. The *region* around its use ensures that the borrow in the left argument does not escape, which brings us to the next topic.

Regions. Borrowing is crucial to support an imperative programming style. To guarantee the validity of a borrow, its lifetime must be properly contained in its ancestor's lifetime. Affe ensures proper nesting of lifetimes by using regions. The expression $\{\!|e|\!\}^n_{\{x \mapsto b\}}$ indicates a region at nesting level n in which a b-borrow can be taken of x.

The typing for a region (rule Region in Fig. 9) replaces suspended bindings by borrow bindings (rule BorrowBinding), typechecks the body of the region, and ensures that the borrow does not leak outside. This last check is done with indices that correspond to the nesting level of the region. The kind k of the borrow is indexed with the level n corresponding to its region, thanks to the constraint $(b_n \leq k)$. The constraint $(\tau \leq \mathbf{L}_{n-1})$ ensures that the return type of the region must live at some enclosing, lower level.

As an example, consider the expression $\{\!|f(\&^{\mathrm{U}} c)|\!\}^n_{\{x \mapsto \mathrm{U}\}}$ where c is a linear channel in an environment Γ. The first step is to check that $[c : \mathrm{channel}]_{\mathrm{U}}$ is in Γ. When entering the region, rule Region imposes $C \vdash_e \Gamma \rightsquigarrow^x_n \Gamma'$, which defines Γ' corresponding to Γ where the suspended binding is replaced by the borrow binding $(c \div \mathrm{channel})^k_{\mathrm{U}}$. To constrain the borrow to this region we impose the constraint $C \vdash_{\mathrm{e}} (\mathbf{U}_n \leq k) \wedge (k \leq \mathbf{U}_\infty)$, which affirms that the borrow is unrestricted, but can only be used in nesting levels n and higher. Rule Region also imposes the constraint $(\tau \leq \mathbf{L}_{n-1})$, which prevents the borrow, having kind k of level $\geq n$, from escaping the region's body of type τ.

Pattern matching. Elimination of pairs is done using a matching construct ($\mathtt{match}_\phi\ x, x' = e_1\ \mathtt{in}\ e_2$). This construct is mostly standard, except it can operate both on a normal pair and a borrow of a pair. The intuition is as follows: A syntactic marker ϕ indicates if it applies to a pair ($\phi = \mathrm{id}$) or a borrow ($\phi = \&^b$). If $\phi = \mathrm{id}$, the typing simplifies to the usual elimination of a pair. Otherwise, e_1 is expected to be a borrow of type $\&^b(k, \tau_1 \times \tau_1')$ and the variables x and x' have type $\&^b(k, \tau_1)$ and $\&^b(k, \tau_1')$, respectively. Thus, the borrow of a pair is considered as a pair of borrows of its components.

Resource management. To demonstrate how Affe deals with resources, we introduce an abstract type $\mathrm{R}\ \tau$ whose content of type τ must be unrestricted ($\mathbf{U}_0$) and which is equipped with the four operations introduced in Section 3.1:

- `create`: $\forall \kappa_\alpha(\alpha : \kappa_\alpha).\ (\kappa_\alpha \leq \mathbf{U}_0) \Rightarrow \alpha \to \mathrm{R}\ \alpha$
- `observe`: $\forall \kappa \kappa_\alpha(\alpha : \kappa_\alpha).\ (\kappa_\alpha \leq \mathbf{U}_0) \Rightarrow \&^{\mathrm{U}}(\kappa, \mathrm{R}\ \alpha) \to \alpha$
- `update`: $\forall \kappa \kappa_\alpha(\alpha : \kappa_\alpha).\ (\kappa_\alpha \leq \mathbf{U}_0) \Rightarrow \&^{\mathbf{A}}(\kappa, \mathrm{R}\ \alpha) \to \alpha \xrightarrow{\mathbf{A}} \mathtt{Unit}$
- `destroy`: $\forall \kappa_\alpha(\alpha : \kappa_\alpha).\ (\kappa_\alpha \leq \mathbf{U}_0) \Rightarrow \mathrm{R}\ \alpha \to \mathtt{Unit}$

3.4 Semantics

It is straightforward to give a standard semantics for Affe, but such a semantics would not be very informative. In this section, we give a big-step semantics that performs explicit bookkeeping of the number of times a value is used and of the mode in which a reference to a resource is used (e.g., borrowed or not). This bookkeeping is based on a set of permissions that regulate the currently allowed mode of access to resources and closures. It enables us to state and prove a highly informative type soundness result (see Section 5) with expressive invariants that ensure proper resource usage.

The dynamics of Affe is given in big-step functional style [3, 28, 34]. A function eval manipulates the semantic objects defined in Fig. 10. The semantics is defined in terms of *elaborated expressions* e with kind, constraint, and splitting annotations inserted by the typechecker. A splitting sp is evidence of the splitting relation for type environments used in the typing rules.

Elaborated expressions

$$
\begin{aligned}
e ::= {} & x \mid x[\overline{k};\overline{\tau}] \mid \lambda^k x.e \mid (e\ e')_{sp} \\
& \mid (e,e')^k_{sp} \mid \texttt{match}_\phi\ x,y =_{sp} e\ \texttt{in}\ e' && \text{(Pairs)} \\
& \mid \texttt{let}\ x =_{sp} e\ \texttt{in}\ e' && \text{(Mono let)} \\
& \mid \texttt{letfun}\ (x:\sigma) =_{sp} \lambda^k y.e\ \texttt{in}\ e' && \text{(Poly let)} \\
& \mid \{\!|e|\!\}^n_{\{x\mapsto b\}} && \text{(Region)} \\
& \mid \&^b x \mid \&\&^b x && \text{(Borrows)} \\
& \mid \texttt{create} \mid \texttt{observe} \\
& \mid \texttt{update} \mid \texttt{destroy} && \text{(Resources)}
\end{aligned}
$$

Splittings

$$sp : (C \vdash_e \Gamma = \Gamma_l \ltimes \Gamma_r)$$

Storables

$$
\begin{aligned}
w ::= {} & \text{STPOLY}(\gamma,\overline{\kappa},C,k,x,e) && \text{(Poly Closures)} \\
& \mid \text{STCLOS}(\gamma,k,x,e) && \text{(Closures)} \\
& \mid \text{STPAIR}(k,r,r') && \text{(Pairs)} \\
& \mid \text{STRSRC}(r) && \text{(Resources)} \\
& \mid \bullet && \text{(Freed Resource)}
\end{aligned}
$$

Environment

$$
\begin{aligned}
\rho ::= {} & \overline{\mathbf{U}}\,\overline{\mathbf{A}}\,\ell && \text{(Locations)} \\
\pi ::= {} & \{\} \mid \pi + \rho && \text{(Permissions)} \\
r ::= {} & \rho \mid c && \text{(Results)} \\
\gamma ::= {} & \cdot \mid \gamma(x \mapsto r) && \text{(Enviroments)} \\
\delta ::= {} & \cdot \mid \delta(\ell \mapsto w) && \text{(Stores)}
\end{aligned}
$$

Fig. 10. Syntax of internal language

Let-polymorphism in the surface language gives rise to elaborated `letfun` expressions annotated with a type scheme σ and a kind k indicating their usage restriction (linear, affine, etc) relative to the variables and constraints of σ. Their use gives rise to explicit instantiation of the kind and type variables. Pairs come with a kind tag k indicating the usage restriction.

Addresses ρ are composed of a raw location ℓ, which is just a pointer into a store, and a stack of modifiers that indicates the borrows and reborrows that have been taken from ℓ. Once we have taken an unrestricted borrow (from a raw location or a borrowed one), then we can take further unrestricted borrows from it, but no more affine ones.

A permission π is a set of addresses that may be accessed during evaluation. A well-formed permission contains at most one address for each raw location.

Non-trivial results are boxed in the semantics. So, a result r is either an address or a primitive constant (e.g., a number).

A value environment γ maps variables to results.

A storable w describes the content of a location in the store. There are five kinds of storables. A *poly closure* represents a polymorphic function. It consists of an environment and the components of an elaborated abstraction. A *closure* represents a monomorphic function in the usual way. A *resource* contains a result and the *hole* $\bullet$ fills a released location.

A store δ is a partial map from raw locations to storables. The function `salloc : store → storable → (loc * store)` is such that `salloc delta w` allocates an unused location in `delta` and fills it with `w`. It returns the location and the extended store.

The evaluation function is indexed by a step count `i` so that each invocation is guaranteed to terminate either with an error, a timeout, or a result. Its return type is a monad α `sem` which combines error reporting and timeout:

```
type α sem = Error of string | TimeOut | Ok of α
val eval: store→perm→venv→int→exp→(store *perm *result) sem
```

Function `eval` evaluates the given expression in the context of an initial store, a permission to use addresses in the store, a value environment, and a step count. If successful, it returns the final store, the remaining permissions, and the actual result.

We give some excerpts of the definition of `eval` in Fig. 11 and leave the full definition for Appendix F. The definition uses OCaml syntax with extensive pretty printing. The pervasive **let***

```
let rec eval
  (δ:store) (π:perm) (γ:venv) i e
  : (store × perm × result) sem =
  if i=0 then TimeOut else
  let i' = i − 1 in
  match e with
  | App (e1, e2, sp) →
    let (γ1, γ2) = vsplit γ sp in
    let* (δ1, π1, r1) = eval δ π γ1 i' e1 in
    let* ℓ1 = getloc r1 in
    let*? () = ℓ1 ∈ π1 in
    let* w = δ1(ℓ1) in
    let* (γ', k', x', e') = getstclos w in
    let π1' = if k' ≤ U then π1 else π1 − ℓ1 in
    let* δ1' = δ1(ℓ1) ← (if k' ≤ U then w else •) in
    let* (δ2, π2, r2) = eval δ1' π1' γ2 i' e2 in
    let* (δ3, π3, r3) = eval δ2 π2 γ'(x' ↦ r2) i' e' in
    Ok (δ3, π3, r3)

  | Borrow (b, x) →
      let+ ρ = γ(x) in
      let*? () = ρ ? b && ρ ∈ π in
      Ok (δ, π, ρ)

  | Varinst (x, k̄) →
    let* rx = γ(x) in
    let* ℓ = getloc rx in
    let*? () = ℓ ∈ π in
    let* w = δ(ℓ) in
    let* (γ', κ̄', C', k', x', e') = getstpoly w in
    let π' =
      if C'{k̄\>κ̄'} =e [(k' ≤ U)]{k̄\>κ̄'}
      then π else π − ℓ
    in
    let w = STCLOS (γ', k'{k̄\>κ̄'}, x', e'{k̄\>κ̄'}) in
    let (ℓ', δ') = salloc δ w in
    Ok (δ', π' + ℓ', ℓ')

  | Region (e, n, x, τx, b) →
    let+ ρ = γ(x) in
    let* ρ' = b.ρ in
    let* π' = reach ρ τx δ in
    let* π'' = b.π' in
    let γ' = γ(x ↦ ρ') in
    let π = (π ∪ π'') \ π' in
    let* (δ1, π1, r1) = eval δ π γ' i' e in
    let π1 = (π1 \ π'') ∪ π' in
    Ok (δ1, π1, r1)
```

Fig. 11. Big-step interpretation

operator acts as monadic bind for the sem monad. The operator **let***?* : bool → (unit → *α* sem) → *α* sem converts a boolean argument into success or failure in the monad.

```
let (let*?) : bool → (unit → β sem) → β sem =
  fun b f → if b then f () else Error ("test failed")
```

The function header of eval checks whether time is up and otherwise proceeds processing the expression.

The Varinst case corresponds to instantiation. It obtains the variable's value, checks that it is a location, checks the permission (the **let***?* clause), obtains the storable *w* at that location, and checks that it is a poly closure (STPOLY). Next, it updates the permission: if the poly closure is unrestricted, then the location remains in the permission set, otherwise it is removed. Finally, we allocate a new monomorphic closure, add it to the permissions, and return the pointer as the result along with the updated store and permissions.

The App case implements (elaborated) function application. We first apply the splitting sp to gamma and evaluate subterm e_1 with its part of the environment and the decremented timer i'. The result must be a location that we are permitted to use. Moreover, there must be a monomorphic STCLOS stored at that location. The permission to further use this closure remains in force only if the closure is unrestricted. Finally, we evaluate the argument, then the function body, and return its result.

The Region case implements a region. It obtains the address for x, the suspended binding, and extends it with the intended borrow *b*. This extension may fail if we try to take an affine borrow of an unrestricted borrow. Next, we rebind x to the borrow's address, extend the permission accordingly, and execute the region's body. Finally, we withdraw the permission and return the result.

The Borrow case obtains the address for x, checks that it is a borrow of the correct mode b and whether it is permitted to use it. It just returns the address.

4 INFERENCE

An important contribution of Affe is its principal type inference. Our type inference algorithm is based on the HM(X) framework [26], a Hindley-Milner type system for a language with constrained types where constraints are expressed in an arbitrary theory X. If X has certain properties, then HM(X) guarantees principal type inference. We apply HM(X) to a concrete constraint language which we name $C_{\mathcal{L}}$. We adapt and extend HM(X)'s rules to support kind inference, track linearity, and handle borrows and regions. We formulate constraint solving and simplification algorithms for $C_{\mathcal{L}}$. Finally, we prove that the inference algorithm computes principal types.

4.1 Preliminaries

In the context of inference, it is critical to know which elements are input and output of inference judgments. In the following, when presenting a new judgment, we write input parameters in **bold green**. The remaining parameters are output parameters.

Usage Environments. To determine if a variable is used in an affine manner, we track its uses and the associated kinds. In the expression $f\ x\ x$, x is used twice. If x is of type τ, which is of kind k, we add the constraint $(k \leq \mathbf{U})$. To infer such constraints, our inference judgment not only takes an environment as parameter but also returns a *usage environment*, denoted Σ, which summarizes usages of variables and borrows. Usage environments are defined like normal environments. In Section 3.3, we use relations to split environments and to transform suspended bindings into borrows inside a region. These relations take a constraint parameter which validates the transformations. In the context of inference, we define new judgments which *infer* the constraints.

- $C \Subset \Sigma = \boldsymbol{\Sigma_1} \ltimes \boldsymbol{\Sigma_2}$. Given two usage environments $\boldsymbol{\Sigma_1}$ and $\boldsymbol{\Sigma_2}$, we return Σ, the merged environment, and C, a set of constraints that must be respected.
- $C \Subset \Sigma \rightsquigarrow^{\boldsymbol{x}}_{\boldsymbol{n}} \boldsymbol{\Sigma'}$. Given a usage environment $\boldsymbol{\Sigma'}$, a nesting level $\boldsymbol{n}$, and a variable name $\boldsymbol{x}$, we return Σ where the borrow binding of x in Σ', if it exists, is replaced by a suspended binding. We also return the constraints C.

Both relations are total and non-ambiguous in term of their input (i.e., functions), and use the rules presented in Sections 3.3 and 3.3. The relations used for syntax-directed typing can trivially be defined in terms of these new relations by using constraint entailment. All relations are fully described in Appendix D.2.

Constraint Normalization. The HM(X) framework assumes the existence of a function "normalize" which takes a constraint $\boldsymbol{C}$ and a substitution $\boldsymbol{\psi}$ and returns a simplified constraint C' and an updated substitution ψ'. Normalization returns a normal form such that ψ' is a most general unifier. For now, we simply assume the existence of such a function for our constraint system and defer details to Section 4.3.

4.2 Type Inference

We write $\Sigma \mid (C, \psi) \mid \boldsymbol{\Gamma} \vdash_{\mathrm{w}} \boldsymbol{e} : \tau$ when $\boldsymbol{e}$ has type τ in $\boldsymbol{\Gamma}$ under the constraints C and unifier ψ with a usage environment Σ. $\boldsymbol{\Gamma}$ and $\boldsymbol{e}$ are the input parameters of our inference algorithm. Unlike in the syntax-directed version, Γ contains only regular and type bindings. Suspended and borrow bindings can only be present in Σ. We revisit some of the syntax-directed rules presented in Section 3.3 to highlight the novelties of our inference algorithm and the differences with the syntax-directed system in Fig. 12. The complete type inference rules are shown in Appendix E.

$$\textsc{Var}_I \quad \frac{(x : \forall\overline{\kappa_i}\forall(\overline{\alpha_j : k_j}).\, C \Rrightarrow \tau) \in \Gamma \qquad \overline{\kappa_i'}, \overline{\alpha_j'} \text{ fresh} \qquad (C, \psi) = \text{normalize}(C_x, [\overline{\kappa_i \mapsto \kappa_i'}, \overline{\alpha_j \mapsto \alpha_j'}])}{(x : \sigma) \,|(C, \psi|_{\text{fv}(\Gamma)}) \,|\, \Gamma \vdash_{\text{w}} x : \psi\tau}$$

$$\textsc{Abs}_I \quad \frac{\alpha, \kappa \text{ fresh} \qquad \Sigma_x \,|(C', \psi') \,|\, \Gamma; (x : \alpha) \vdash_{\text{w}} e : \tau \qquad D = C' \wedge (\Sigma_x \backslash \{x\} \leq \kappa) \wedge \text{Weak}_{(x:\alpha)}(\Sigma_x) \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma_x \backslash \{x\} \,|(C, \psi \backslash \{\alpha, \kappa\}) \,|\, \Gamma \vdash_{\text{w}} \lambda x.e : \psi(\alpha) \xrightarrow{\psi(\kappa)} \tau}$$

$$\textsc{Region}_I \quad \frac{\Sigma' \,|(C', \psi') \,|\, \Gamma \vdash_{\text{w}} e : \tau \qquad C_r \Lleftarrow \Sigma \rightsquigarrow_n^x \Sigma' \qquad (C_\tau, \psi_\tau) \,|\, \Gamma \vdash_{\text{w}} \tau : k_\tau \qquad D = C' \wedge C_\tau \wedge (k_\tau \leq \mathbf{L}_{n-1}) \wedge C_r \qquad (C, \psi) = \text{normalize}(D, \psi' \sqcup \psi_\tau)}{\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} \{\!|e|\!\}^n_{\{x \mapsto b\}} : \tau}$$

$$\textsc{App}_I \quad \frac{\alpha, \kappa \text{ fresh} \qquad \Sigma_1 \,|(C_1, \psi_1) \,|\, \Gamma \vdash_{\text{w}} e_1 : \tau_1 \qquad C_s \Lleftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2 \qquad \Sigma_2 \,|(C_2, \psi_2) \,|\, \Gamma \vdash_{\text{w}} e_2 : \tau_2 \qquad D = C_1 \wedge C_2 \wedge (\tau_1 \leq \tau_2 \xrightarrow{\kappa} \alpha) \wedge C_s \qquad \psi' = \psi_1 \sqcup \psi_2 \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} (e_1\ e_2) : \psi(\alpha)}$$

$$\text{Weak}_{(x:\sigma)}(\Sigma) = \text{if } (x \in \Sigma) \text{ then True else } (\sigma \leq \mathbf{A}_\infty)$$

Fig. 12. Selected inference rules – $\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} e : \tau$

Environments and Bindings. In the syntax-directed system, the Var rule ensure that linear variables are not discarded at the *leaves*. In the inference algorithm, we operate in the opposite direction: we collect data from the leaves and enforce linearity at *binders*. This policy is reflected in the $\textsc{Var}_I$ and $\textsc{Abs}_I$ rules. Typing for variables is very similar to traditional Hindley-Milner type inference. To keep track of linearity, we record that x was used with the scheme σ by returning a usage environment $\Sigma = \{(x : \sigma)\}$. This usage environment is in turn used at each binder to enforce proper usage of linear variable via the Weak property as shown for lambda expressions in the $\textsc{Abs}_I$ rule. First, we typecheck the body of the lambda and obtain a usage environment Σ_x. As in the syntax-directed type system, we introduce the constraint $(\Sigma\backslash\{x\} \leq \kappa)$ which properly accounts for captures in the body of the lambda expression. We then introduce the constraint $\text{Weak}_{(x:\sigma)}(\Sigma)$, which fails if we try to abandon a linear variable. The Weak constraint is introduced at each binding construct. Finally, we normalize constraints to ensure that the inference algorithm always return the simplest possible constraints and unifiers.

Splitting and Regions. Inference versions of the App and Region rules are similar to the original ones, but now *return* the usage environment Σ. As such, we use the "inference" version of the relations on the environment, $C \Lleftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2$ and $C \Lleftarrow \Sigma \rightsquigarrow_n^x \Sigma'$, which returns the necessary constraints. We then collect all constraints and normalize them.

4.3 Constraints

To properly define our type system, we need to define $C_{\mathcal{L}}$, a constraint system equipped with an entailment relation noted $\vdash_{\text{e}}$ and a normalizing function. For concision, we first demonstrate the constraint solving algorithm with an example. We then state the various properties that make it suitable for use in the HM(X) framework. The complete constraint system is defined in Appendix C.

4.3.1 Constraints normalization by example. Consider the expression $\lambda f.\lambda x.((f\ x), x)$. The inference algorithm yields the following constraints:

$$\Gamma = (\alpha_f : \kappa_f)(\alpha_x : \kappa_x) \ldots$$
$$C = (\alpha_f \leq \gamma \xrightarrow{\kappa_1} \beta) \wedge (\gamma \leq \alpha_x) \wedge (\beta \times \alpha_x \leq \alpha_r) \wedge (\kappa_x \leq \mathbf{U})$$

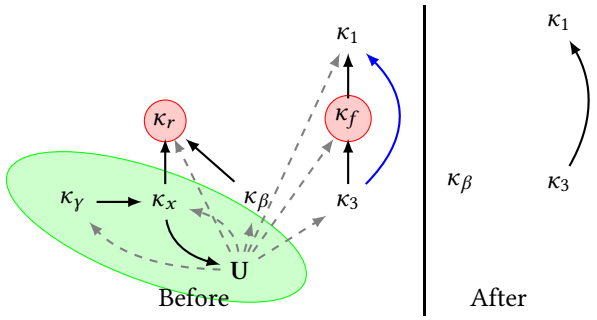

Fig. 13. Graph representing the example constraints

The first step of the algorithm uses Herbrand unification to obtain a type skeleton.

$$(\gamma \xrightarrow{\kappa_3} \beta) \xrightarrow{\kappa_2} \gamma \xrightarrow{\kappa_1} \beta \times \gamma$$

In addition, we obtain the following kind constraints:

$$(\kappa_x \leq \mathbf{U}) \wedge (\kappa_\gamma \leq \kappa_x) \wedge (\kappa_x \leq \kappa_r) \wedge (\kappa_\beta \leq \kappa_r) \wedge (\kappa_3 \leq \kappa_f) \wedge (\kappa_f \leq \kappa_1)$$

We translate these constraints into a relation whose graph is shown in Fig. 13. The algorithm then proceeds as follow:

- From the constraints above, we deduce the graph shown with plain arrows on the left of Fig. 13.
- We add all the dashed arrows by saturating lattice inequalities. For clarity, we only show $\mathbf{U}$.
- We identify the connected component circled in green. We deduce $\kappa_\gamma = \kappa_x = \mathbf{U}$.
- We take the transitive closure, which adds the arrow in blue from κ_3 to κ_1.
- We remove the remaining nodes not present in the type skeleton (colored in red): κ_r and κ_f.
- We clean up the graph (transitive reduction, remove unneeded constants, ...), and obtain the graph shown on the right. We deduce $\kappa_3 \leq \kappa_1$.

The final constraint is thus

$$\kappa_\gamma = \kappa_x = \mathbf{U} \wedge \kappa_3 \leq \kappa_1$$

If we were to generalize, we would obtain the type scheme:

$$\forall \kappa_\beta \kappa_1 \kappa_2 \kappa_3 (\gamma : \mathbf{U})(\beta : \kappa_\beta).\ (\kappa_3 \leq \kappa_1) \Rightarrow (\gamma \xrightarrow{\kappa_3} \beta) \xrightarrow{\kappa_2} \gamma \xrightarrow{\kappa_1} \beta \times \gamma$$

We can further simplify this type by exploiting variance. As κ_1 and κ_2 are only used in covariant position, they can be replaced by their lower bounds, κ_3 and $\mathbf{U}$. By removing the unused quantifiers, we obtain a much simplified equivalent type:

$$\forall \kappa (\gamma : \mathbf{U}).(\gamma \xrightarrow{\kappa} \beta) \rightarrow \gamma \xrightarrow{\kappa} \beta \times \gamma$$

4.3.2 Properties of the constraint system. To apply HM(X) to $C_{\mathcal{L}}$, normalize must compute principal normal forms and $C_{\mathcal{L}}$ must be regular.

Property 1 (Principal normal form). *Normalization computes principal normal forms for $C_{\mathcal{L}}$, i.e. given a constraint $D \in C_{\mathcal{L}}$, a substitution ϕ and $(C, \psi) = \text{normalize}(D, \phi)$, then $\phi \leq \psi$, $C =_e \psi D$ and $\psi C = C$.*

Property 2 (Regular constraint system). *$C_{\mathcal{L}}$ is regular, ie, for x, x' two types or kinds, $\vdash_e (x = x')$ implies $\text{fv}(x) = \text{fv}(x')$*

These properties are sufficient to state that $HM(C_{\mathcal{L}})$ provides principal type inference. The next section shows that these properties carry over to the inference algorithm for our extension of HM(X) with kind inference, affine types, and borrows. This algorithm includes sound and complete constraint simplification. In addition, we may add "best-effort" simplification rules which help reduce the size of inferred signatures [36].

4.4 Soundness and Principality

The extended inference algorithm is sound and complete with respect to our extension of HM(X). The first theorem states that inference is sound with respect to the syntax-directed type system.

Theorem 4.1 (Soundness of inference). *Given a type environment Γ containing only value bindings, $\Gamma|_\tau$ containing only type bindings, and a term e:*
if $\Sigma \,|(C, \psi) \,|\, \Gamma; \Gamma_\tau \vdash_{\mathrm{w}} e : \tau$
then $C \,|\, \psi(\Sigma; \Gamma_\tau) \vdash_{\mathrm{s}} e : \tau$, $\psi C = C$ *and* $\psi\tau = \tau$

The syntax-directed derivation holds with the usage environment Σ instead of the originally provided environment Γ. Indeed, Γ does not contain suspended and borrow bindings. Those are discovered on the fly and recorded in Σ. Type bindings are taken directly from the syntax-directed derivation.

The second theorem states that inference is complete: for any given syntax-directed typing derivation, our inference algorithm can find a derivation that gives a type at least as general.

Definition 4.2 (Instance relation). Given a constraint C and two schemes $\sigma = \forall\overline{\alpha}.D \Rightarrow \tau$ and $\sigma' = \forall\overline{\alpha'}.D' \Rightarrow \tau'$. Then $C \vdash_{\mathrm{e}} \sigma \preceq \sigma'$ iff $C \vdash_{\mathrm{e}} D[\alpha \to \tau'']$ and $C \wedge D' \vdash_{\mathrm{e}} (\tau[\alpha \to \tau''] \preceq \tau')$

Definition 4.3 (Flattened Environment). A flattened environment, written as $\Downarrow\Gamma$, is the environment where all the binders are replaced by normal ones. More formally:

$$\Downarrow\Gamma = \left\{(x : \tau) \in \Gamma \mid \vee(\&^b x : \&^b(k, \tau)) \in \Gamma \vee [x : \tau]^n_b \in \Gamma\right\} \cup \{(\alpha : k) \mid (\alpha : k) \in \Gamma\}$$

Theorem 4.4 (Principality). *Let* $\mathrm{True} \,|\, \Gamma \vdash_{\mathrm{s}} e : \sigma$ *a closed typing judgment. Then* $\Sigma \,|(C, \psi) \,|\, \Downarrow \Gamma \vdash_{\mathrm{w}} e : \tau$ *such that:*

$$(\mathrm{True}, \sigma_o) = gen(C, \psi\Gamma, \tau) \qquad\qquad \vdash_{\mathrm{e}} \sigma_o \preceq \sigma$$

5 METATHEORY

There are several connections between the type system and the operational semantics, which we state as a single type soundness theorem. The theorem relies on several standard notions like store typing $\vdash \delta : \Delta$ and agreement of the results in the value environment with the type environment $\Delta \vdash \gamma : \Gamma$ that we define formally in Appendix G where we also present selected cases of the proofs. The non-standard part is the handling of permissions. With $\mathrm{getloc}(\pi)$ we extract the underlying raw locations from the permissions as in $\mathrm{getloc}(\overline{\mathbf{U}}\,\overline{\mathbf{A}}\,\ell) = \ell$ and with $\mathrm{reach}_\delta(\gamma)$ we transitively trace the addresses reachable from γ in store δ. We write $\Delta \preceq \Delta'$ and $\delta \preceq \delta'$ for extending the domain of the store type and of the store, respectively. The permission set contains the set of addresses that can be used during evaluation. It is managed by the region expression as well as by creation and use of resources as shown in Section 3.4. We distinguish several parts of the value environment γ that correspond to the different kinds of bindings in the type environment: $\gamma^{\mathbf{L}}$ for active entries of direct references to linear resources, closures, etc; $\gamma^{\mathbf{A}}$ for affine borrows or resources; $\gamma^{\mathbf{U}}$ for unrestricted values including unrestricted borrows; and $\gamma_{\#}$ for suspended entries. The judgment $\Delta \vdash \gamma : \Gamma$ is defined in terms of this structure. We treat $\mathrm{reach}_\delta(\gamma)$ as a multiset to properly discuss linearity and affinity. We use the notation $M(x)$ for the number of times x occurs in multiset M.

Theorem 5.1 (Type Soundness). *Suppose that*

(A1) $C \mid \Gamma \vdash_s e : \tau$
(A2) $\Delta \vdash \gamma : \Gamma$
(A3) $\vdash \delta : \Delta$
(A4) π *is wellformed and* $\mathrm{getloc}(\pi) \subseteq \mathrm{dom}(\delta) \setminus \delta^{-1}(\bullet)$
(A5) $\mathrm{reach}_0(\gamma) \subseteq \pi$, $\mathrm{reach}_\delta(\gamma) \subseteq \downarrow\pi$.
(A6) $\mathrm{getloc}(\gamma^{\mathbf{L}})$, $\mathrm{getloc}(\gamma^{\mathbf{A}})$, $\mathrm{getloc}(\gamma^{\mathbf{U}})$, *and* $\mathrm{getloc}(\gamma_\#)$ *are all disjoint*
(A7) Incoming Resources:
  (a) $\forall \ell \in \mathrm{getloc}(\mathrm{reach}_\delta(\gamma)), \delta(\ell) \neq \bullet$.
  (b) $\forall \ell \in \Theta = \mathrm{getloc}(\mathrm{reach}_\delta(\gamma^{\mathbf{L}}, \gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})), \Theta(\ell) = 1$.

For all $i \in \mathbb{N}$, *if* R' = eval δ π γ i e *and* $R' \neq$ TimeOut, *then* $\exists\, \delta', \pi', r', \Delta'$ *such that*

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$
(R2) $\Delta \leq \Delta'$, $\delta \leq \delta'$, $\vdash \delta' : \Delta'$
(R3) $\Delta' \vdash r' : \tau$
(R4) π' *is wellformed and* $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
(R6) Frame:
  For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ *it must be that*
  - $\delta'(\ell) = \delta(\ell)$ *and*
  - *for any* ρ *with* $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

(R7) Unrestricted values, resources, and borrows:
  For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{U}}, \gamma_\#^{\mathbf{U}})$ *with* $\mathrm{getloc}(\rho) = \{\ell\}$, *it must be that* $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ *and* $\rho \in \pi'$.
(R8) Affine borrows and resources:
  For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})$ *with* $\mathrm{getloc}(\rho) = \{\ell\}$, *it must be that* $\ell \in \mathrm{dom}(\delta)$. *If* $\rho \neq \ell$, *then* $\delta'(\ell) \neq \bullet$. *If* $\rho \in \mathrm{reach}_{\delta'}(\gamma_\#^{\mathbf{A}})$, *then* $\rho \in \pi'$.
(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. *Let* $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathbf{L}})$.
  For all $\ell \in \Theta$ *it must be that* $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, *and if* $\delta(\ell)$ *is a resource, then* $\delta'(\ell) = \bullet$.
(R10) No thin air permission:
  $\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.

The proof of the theorem is by functional induction on the evaluation judgment, which is indexed by the strictly decreasing counter i.

The assumptions A1-A3 and results R1-R3 state the standard soundness properties for lambda calculi with references.

The rest of the statement accounts for the substructural properties and borrowing in the presence of explicit resource management. Incoming resources are always active (i.e., not freed). Linear and affine resources as well as suspended affine borrows have exactly one pointer in the environment. The Frame condition states that only store locations reachable from the current environment can change and that all permissions outside the reachable locations remain the same. Unrestricted values, resources, and borrows do not change their underlying resource and do not spend their permission. Affine borrows and resources may or may not spend their permission. Borrows are not freed, but resources may be freed. Incoming suspended borrows have no permission attached to them and their permission has been retracted on exit of their region. A linear resource is always freed. Outgoing permissions are either inherited from the caller or they refer to newly created values.

6 LIMITATIONS AND EXTENSIONS

6.1 Flow sensitivity

The type system defined so far does not support any form of flow sensitivity. Therefore, code patterns that rely on subtle flow-sensitive usage of permissions and linearity will most likely not typecheck in Affe. For example, the following merge function on linear lists cannot be expressed directly, because matching against `l1` and `l2` consumes both lists.

```
let rec merge l1 l2 = match l1, l2 with
 | h1::t1, h2::t2 →
   if &h1 < &h2
   then h1::(merge t1 l2) (* Must expand l2 to h2::t2 here *)
   else h2::(merge l1 t2)
 | ....
```

Patterns like this require a richer logic, such as provided by Mezzo [31]. However, Weiss et al. [42] formalize Rust's notion of non-lexical lifetimes which partially supports such code patterns. We believe this notion can be adapted to Affe's notion of regions.

Non-Lexical Regions. The notion of non-lexical lifetimes is a recent addition to Rust. With this feature code is acceptable even if borrowing does not respect lexical scoping as in this example:

```
let a = &x in (f a; g &!x)
```

This code pattern is dynamically safe because a is not used after the function call `f a`. Here, this can be made explicit by transforming the code to `(let a = &x in f a); g &!x`. However, this is not possible in programs with branches who uses different dynamic patterns. Non-lexical lifetimes (NLL) handle such a pattern by removing expressions that do not mention a from its region; in this example, NLL removes the last expression. In Affe, regions are lexical and marked by the expression $\{\!|e|\!\}^n_b$. During inference, kind constraints prevent escaping from a region.

To add support for non-lexical lifetimes, we could replace the lexical region by an annotation on each expression indicating which borrows are live in this expression. When exploring a subexpression, we would compare the annotations, and automatically apply the Region rule when they differ. This approach is equivalent to inlining the Region rule in all the other rules.

Applied to the program above, only the first two expressions would be annotated to be "in the region associated with &x", but not the last expression. Thanks to these annotations, type checking the sequence would check that the borrow does not escape the left-hand side (i.e., the second expression `f a`).

6.2 Capabilities and Identity

In Affe the tracking of linearity does not rely on any notion of "identity": the type system cannot specify that two objects are the same, simply that they share the same usage pattern with regards to linearity. A language like Alms [40], on the other hand, often relies on a notion of identity to express capabilities. For instance, the Alms typing `Array.create : int → α → ∃ β. (α, β) array` uses β as a unique identification of the array. Functions such as `Array.acquire : (α, β) array → β cap` are used to obtain capabilities to operate on the array.

While such uses are partially covered by borrows and regions, a notion of identity associated to regions would enable us to express regions directly in type signatures. For instance, the `get_eb` function shown in Section 2.3 could be made safe by creating a restricted inner region on function application, with the signature: `&!(κ, α Array.t) → int → ∃ (κ' < κ) &!(κ', α)`

This approach relies on existential types to model identities. At present, Affe does not support existentials as it would forgo principal type inference. However, existentials are compatible with the HM(X) framework [35] and would make a very desirable addition to Affe. Work on GADTs in

OCaml and Haskell demonstrates that existential types can be put to use without compromising inference in the rest of the language, by integrating unpacking and pattern matching.

6.3 Ad-hoc Polymorphism and Borrows

In our formalization, we use two operators, $\&^b x$ and $\&\&^b x$ to distinguish between borrows and borrows of borrows. Such a distinction is inconvenient for programming. Using a typeclass-like mechanism, we can replace these operators by a single overloaded operator, $\&^b x$, which expects x to be `Borrowable` and would then desugar to the more precise operators. A similar solution is used in Rust through the `Borrow` and `Defer` traits. This approach also enables method calls on objects without explicit borrows, such as `foo.len()` where len expects a shared borrow.

Ad-hoc polymorphism fits demonstrably in the HM(X) framework of constrained types and preserves all properties of our language such as principal type inference. Its soundness is orthogonal to linear types and has been explored in the literature [27].

6.4 A Richer Region System

Affe requires that each region is identified by an index drawn from a partial order that is compatible with the nesting of regions. This order can be implemented in many ways, including region variables as often used in algebraic effects systems, existentials, etc.

For simplicity, the formalization uses the concrete implementation with natural numbers for indices. The proofs only rely on the existence of a partial order and could be adapted to one of the more abstract approaches. In particular, Affe could reuse regions variables provided by the ongoing work on effect systems for OCaml [11].

6.5 Standard Features

Algebraic Datatypes. Algebraic data types are a staple of functional programming and fit nicely in our paradigm. Indeed, it is sufficient to ensure that the kinds of the constructor arguments are less than or equal to the kind of the datatype. Hence, it is forbidden to include affine elements in an unrestricted datatype, whereas the elements in a linear list may be linear or unrestricted. Our prototype implements non-recursive algebraic datatypes with pattern matching.

Branching constructs. Our formalization of Affe does not cover conditionals. In the typing rules, a conditional is supported as usual by checking all branches with the same constraint and typing environment and requiring the return types to match. It materializes in the inference algorithm as a straightforward (symmetric) join relation which is used for all elimination rules on sum types. This extension is implemented in our prototype.

6.6 Concurrency

While our present semantic model does not consider concurrency, some design decisions were taken with a possible extension to concurrency in mind. The main impact of this foresight is the distinction between exclusive borrows and shared borrows, which materializes in the metatheory. The intended contract of the shared borrow is that it can be duplicated and that the program consistently observes the same state of the underlying resource inside its region even in the presence of concurrency.

The exclusive borrow, on the other hand, is propagated according to the evaluation order with the intention that any suspended binding split of from an exclusive borrow has finished its action on the resource before the borrow gets exercised. In the presence of concurrency, this intended semantics of the exclusive borrow should guarantee freedom of data races.

Language	UAL	State	Borrows	Multiplicity Subsumption	Multiplicity Polymorphism	Identity	Concurrency	Escape hatch	Inference	Formalisation	Basis
System F° [23]	UL	✘	✘	✘	✘	✘	✘	✘	✘	Coq	System F
Alms [40]	UA	✔	✘	~	✔	✔	✔	✔	Local	✘	ML
Quill [24]	UL	✘	✘	✘	✔	✘	✘	✘	Principal	Manual	Qual. types
Lin. Haskell [7]	UL	~	✘	✘	✔	✘	~	~	Non-pr.	Manual	Haskell
Mezzo [5, 31]	UA	✔	~	~	✔	✔	✔	✔	Local	Coq	ML
Rust [19, 21]	UA	✔	✔	✔	~	✘	✔	✔	Local	Coq	—
Plaid [2, 12]	UA	✔	✘	✔	✔	✔	✘	✔	✘	Manual	Java
Affe	UAL	✔	✔	✔	✔	✘	✘	~	Principal	Manual	ML/HM(X)

Fig. 14. Comparison matrix

That is, if a thread closes over a borrow, that thread should have terminated before the parent thread leaves the borrow's region. Rust addresses this lifetime issue with the `move` qualification for a thread which transfers ownership of the free variables to the thread. However, moving (in Rust) only applies to the resource itself, but not to borrows. A more discerning kind system would be needed for Affe to enable safe sharing of synchronizable resources or borrows analogously to Rust's Sync trait.

7 RELATED WORK

The comparison matrix in Fig. 14 gives an overview over the systems discussed in this section. Each column indicates whether a feature is present (✔), absent (✘), or partially supported (~), i.e., if the feature is limited or can only be obtained through a non-trivial encoding. Features are selected according to their relevance for type-based resource management and programmer convenience.

The column UAL specifies the substructural features supported (Unrestricted, Affine, Linear). The columns "State" and "Borrows" indicate support for the respective feature. In an ideal world, the presence of linearity and state indicates that the system is able to support safe manual memory management as linearity enforces manual deallocation. True affinity and state only works with garbage collection, which eventually automatically finalizes an object no longer referenced. In practice, this distinction is often watered down. For example, Rust automatically destructs objects at the end of their lifetime, creating the illusion of affinity while the low-level code is strictly linear. However, there are ways to consume an object at the source level without invoking its destructor (using `mem::forget`)[3] where the high-level code exhibits linearity, but the low-level code is affine.

"Multiplicity Subsumption" indicates that unrestricted elements can be promoted to affine and then linear. This promotion applies to objects, resources, borrows, and closures. "Multiplicity Polymorphism" refers to polymorphism over substructural features: a function can be parameterized over the multiplicity restriction of an object. For instance, the type of function composition should

[3]See https://doc.rust-lang.org/nomicon/leaking.html which contains further examples and discussion. Thanks to Derek Dreyer and Ralf Jung for pointing this out.

express that applies to functions with linear, affine, and unrestricted multiplicity and returns a function with the same multiplicity. "Identity" indicates that the language supports a notion of identity, usually through existential types, as described in Section 6.2. "Concurrency" indicates whether the language supports concurrency. For example, the implementation of Linear Haskell supports state and concurrency, but its theory covers neither. "Escape hatch" indicates whether a programmer can (locally) opt out of resource management through language-integrated means such as Rust's `unsafe`. Partial support "~", in the case of Affe for instance, indicates that this feature is available, but not formalized. Type "inference" can be local, principal, or non-principal (if the inferred type is not necessarily the most general one). "Formalization" refers to the existence of a formal semantics and type soundness proof. "Basis" indicates the heritage or inspiration of the language.

7.1 Substructural type-systems in functional languages

Many systems propose combinations of functional programming and linear types in a practical setting. The goal of Affe is to combine key ingredients from these proposals while still preserving complete type inference. Many of the following languages support linear or affine types, but rarely both. In many cases, it is easy to adapt a system to support both, as Affe does. None of the following languages support borrows.

System F° [23] extends System F with kinds to distinguish between linear and unrestricted types. The authors provide a linearity-aware semantics with a soundness proof. Unlike Affe, System F° does not allow quantification over kinds which limits its expressivity. For instance, it does not admit a most general type for function composition. Being based on System F, it does not admit principal type inference.

Quill [24] is a Haskell-like language with linear types. Quill does not expose a kind language, but uses the framework of qualified types to govern linearity annotations on arrows. Its type inference algorithm is proven sound and complete. Affe infers type signatures for all Quill examples, but often with simpler types because Quill does not support subkinding. Quill comes with a linearity-aware semantics and soundness proof. Quill does not support borrows.

Alms [40] is an ML-like language with rich, kind-based affine types and ML modules, similar to Affe. Alms examples often rely on existential types to track the identity of objects. For instance, consider the signature `Array.create : int →`α`→∃`β`. (`α`, `β`)array` where β uniquely identifies the array. Due to the reliance on existentials, Alms does not support complete type inference. Furthermore, Alms does not support borrows and often relies on explicit capability passing. In our experience, Affe's limited support for existential types through regions is sufficient to express many of Alms' examples and leads to a more convenient programming style for imperative code. Alms kind structure features unions, intersections and dependent kinds while Affe uses constrained types. We believe most of Alms' kind signatures can be expressed equivalently in our system: for instance the pair type constructor has kind $\Pi\alpha\Pi\beta.\langle\alpha\rangle \sqcup \langle\beta\rangle$ (where α and β are types and Π is the dependent function) in Alms compared to $\kappa \to \kappa \to \kappa$ in Affe thanks to subkinding. Finally, Alms provides excellent support for abstraction through modules by allowing to keep some type unrestricted inside a module, but exposing it as affine. Affe supports such programming style thanks to subsumption.

The goal of Linear Haskell [7] (LH) is to retrofit linear types to Haskell. Unlike the previously discussed approaches, LH relies on "linear arrows", written $\multimap$ as in linear logic, which are functions that *use* their argument exactly once. This design is easy to retrofit on top of an existing compiler such as GHC, but has proven quite controversial[4]. Most relevant to Affe:

[4] See the in-depth discussion attached to the GHC proposal for LH on GitHub: https://github.com/ghc-proposals/ghc-proposals/pull/111#issuecomment-403349707.

- LH does not admit subtyping for arrows and requires η-expansion to pass unrestricted functions in linear contexts. This approach is acceptable in a non-strict language such as Haskell but changes the semantics in a strict setting.
- While the LH paper specifies a full type system along with a linearity-aware soundness proof, there is neither formal description of the type inference algorithm nor a proof of the properties of inference. Subsequent work [22] formalizes the inference for rank 1 qualified-types. However, there is an implementation of the inference as part of GHC.
- LH promotes a continuation-passing style with functions such as `withFile : path →(file ⊸Unrestricted r)⊸r` to ensure linear use of resources. This style leads to problems with placing the annotation on, e.g., the IO monad. Affe follows System F°, Quill, and Alms, all of which support resource handling in direct style, where types themselves are described as affine or linear. (Of course, continuation-passing style is also supported.) We expect that the direct approach eases modular reasoning about linearity. In particular, using abstraction through modules, programmers only need to consider the module implementation to ensure that linear resources are properly handled.

Mezzo [5, 31] is an ML-like language with a rich capability system which is able to encode numerous properties akin to separation logic [33]. Mezzo explores the boundaries of the design space of type systems for resources. Hence, it is more expressive than Affe, but much harder to use. The Mezzo typechecker relies on explicit annotations and it is not known whether type inference for Mezzo is possible.

Munch-Maccagnoni [25] presents an extension of OCaml for resource management in the style of C++'s RAII and Rust's lifetimes. This system assumes the existence of a linear type system and develops the associated compilation and runtime infrastructure. We believe our approach is complementary and aim to combine them in the future.

7.2 Other substructural type-systems

Affe uses borrows and regions which were initially developed in the context of linear and affine typing for imperative and object-oriented programming [8, 15].

Rust [21] is the first mainstream language that builds on the concepts of borrowing and ownership to enable safe low-level programming. Affe is inspired by Rust's borrowing system and transfers some of its ideas to a functional setting with type inference, garbage collection, and an ML-like module system. Everything is affine in Rust and marker traits like `Copy`, `Send`, and `Sync` are used to modulate the characteristics of types. Affe relies on kinds to express substructural properties of types and marker traits may be considered as implementing a fine-grained kind structure on Rust types. Rust's lifetime system is more explicit and more expressive than Affe's regions. While Rust provides partial lifetime inference, it does not support full type inference. Moreover, Rust programmers have full control over memory allocation and memory layout of objects; they can pass arguments by value or by reference. These features are crucial for the efficiency goals of Rust. In contrast, Affe is garbage collected, assumes a uniform object representation, and all arguments are passed by reference. This choice forgoes numerous issues regarding interior mutability and algebraic data types. In particular, it allows us to easily nest mutable references inside objects, regardless whether they are linear or unrestricted.

In Rust, programmers can implement their low-level abstractions by using unsafe code fragments. Unsafe code is not typechecked with the full force of the Rust type system, but with a watered down version that ignores ownership and lifetimes. This loophole is needed to implement datastructures like doubly-linked lists or advanced concurrency abstractions. When unsafe code occurs as part of a function body, the Rust typechecker leaves the adherence of the unsafe

code to the function's type signature as a proof obligation to the programmer. The RustBelt project [19] provides a formal foundation for creating such proofs by exhibiting a framework for semantic soundness of the Rust type system in terms of a low-level core language that incorporates aspects of concurrency (i.e., data-race freedom). Similar proof obligations would be needed in Affe to check that an implementation of the module types or the type of fold shown in Section 2 matches the semantics of the typings. We aim to develop a suitable framework for this task for Affe. At present, the metatheory of Affe does not cover concurrency.

Weiss et al. [42] formalize Rust's ownership discipline from a source-level perspective. Their approach is purely syntactic and is therefore not able to reason about unsafe fragments of Rust code. However, their flow-sensitive type discipline enables soundness proofs for non-lexical lifetimes, which have been adopted in Rust, but cannot be expressed in Affe at present.

Vault [10] and Plaid [2, 12] leverage typestate and capabilities to express rich properties in objects and protocols. These systems are designed for either low-level or object-oriented programming and do not immediately lend themselves to a more functional style. While these systems are much more powerful than Affe's, they require programmer annotations and do not support inference. It would be interesting to extend Affe with limited forms of typestate as a local, opt-in feature to provide more expressivity at the cost of inference.

7.3 Type-system features

Affe relies on constrained types to introduce the kind inequalities required for linear types. HM(X) [26] allows us to use constrained types in an ML-like language with complete type inference. HM(X) has been shown to be compatible with subtyping, bounded quantification and existentials [35], GADTs [37], and there exists a syntactic soundness proof [38]. These results make us confident that the system developed in Affe could be applied to larger and more complex languages such as OCaml and the full range of features based on ad-hoc polymorphism.

Affe's subtyping discipline is similar to structural subtyping, where the only subtyping (or here, subkinding) is at the leaves. Such a discipline is known to be friendly to inference and has been used in many contexts, including OCaml, and has been combined with constraints [26, 41]. It also admits classical simplification rules [30, 36] which we partially use in our constraint solving algorithm. Affe's novelty is a kind language sufficiently simple to make all simplification rules complete, which allows us to keep type signatures simple.

8 CONCLUSIONS

Affe is an ML-like language extended with sound handling of linear and affine resources. Its main novel feature is the combination of full type inference and a practically useful notion of shared and exclusive borrowing of linear and affine resources. Although the inferred types are much richer internally than plain ML types, most of that complexity can be hidden from user-level programmers. On the other hand, programmers of libraries dealing with resources have sufficient expressiveness at their fingertips to express many resource management schemes.

The main restriction of the current system is that the lifetime of borrows is determined by lexical scoping. Overcoming this restriction is subject of future work and will probably require extending the type system by some notion of effect, which is currently discussed in the OCaml community. Moreover, other systems rely on existential types for extra expressiveness. We chose not to include existentials to preserve complete type inference, but our design can be extended in this direction. Finally, our matching construct is very simplistic. Our implementation supports full algebraic data types and we believe it can be further extended to support manipulating borrows of data-structures and internal mutability.

ACKNOWLEDGMENTS

This material is based upon work supported by the German Research Council, DFG, project reference number TH 665/11-1. We are indebted to the anonymous reviewers for their thoughtful and constructive comments.

REFERENCES

[1] Peter Achten and Marinus J. Plasmeijer. 1995. The Ins and Outs of Clean I/O. *J. Funct. Program.* 5, 1 (1995), 81–110. https://doi.org/10.1017/S0956796800001258

[2] Jonathan Aldrich, Joshua Sunshine, Darpan Saini, and Zachary Sparks. 2009. Typestate-oriented programming. In *Companion to the 24th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2009, October 25-29, 2009, Orlando, Florida, USA*, Shail Arora and Gary T. Leavens (Eds.). ACM, 1015–1022. https://doi.org/10.1145/1639950.1640073

[3] Nada Amin and Tiark Rompf. 2017. Type Soundness Proofs With Definitional Interpreters. In *POPL*. ACM, 666–679.

[4] Phil Bagwell. 2001. Ideal Hash Trees.

[5] Thibaut Balabonski, François Pottier, and Jonathan Protzenko. 2016. The Design and Formalization of Mezzo, a Permission-Based Programming Language. *ACM Trans. Program. Lang. Syst.* 38, 4 (2016), 14:1–14:94. http://dl.acm.org/citation.cfm?id=2837022

[6] Erik Barendsen and Sjaak Smetsers. 1995. Uniqueness Type Inference. In *Programming Languages: Implementations, Logics and Programs, 7th International Symposium, PLILP'95, Utrecht, The Netherlands, September 20-22, 1995, Proceedings (Lecture Notes in Computer Science)*, Manuel V. Hermenegildo and S. Doaitse Swierstra (Eds.), Vol. 982. Springer, 189–206. https://doi.org/10.1007/BFb0026821

[7] Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. 2018. Linear Haskell: Practical Linearity in a Higher-Order Polymorphic Language. *PACMPL* 2, POPL (2018), 5:1–5:29. https://doi.org/10.1145/3158093

[8] John Tang Boyland and William Retert. 2005. Connecting Effects and Uniqueness with Adoption. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, Jens Palsberg and Martín Abadi (Eds.). ACM, 283–295. https://doi.org/10.1145/1040305.1040329

[9] Sylvain Conchon and Jean-Christophe Filliâtre. 2007. A Persistent Union-Find Data Structure. In *Proceedings of the ACM Workshop on ML, 2007, Freiburg, Germany, October 5, 2007*, Claudio V. Russo and Derek Dreyer (Eds.). ACM, 37–46. https://doi.org/10.1145/1292535.1292541

[10] Robert DeLine and Manuel Fähndrich. 2001. Enforcing High-Level Protocols in Low-Level Software. In *PLDI*. ACM, 59–69.

[11] Stephen Dolan, Spiros Eliopoulos, Daniel Hillerström, Anil Madhavapeddy, K. C. Sivaramakrishnan, and Leo White. 2017. Concurrent System Programming with Effect Handlers. In *Trends in Functional Programming - 18th International Symposium, TFP 2017, Canterbury, UK, June 19-21, 2017, Revised Selected Papers (Lecture Notes in Computer Science)*, Meng Wang and Scott Owens (Eds.), Vol. 10788. Springer, 98–117. https://doi.org/10.1007/978-3-319-89719-6_6

[12] Ronald Garcia, Éric Tanter, Roger Wolff, and Jonathan Aldrich. 2014. Foundations of Typestate-Oriented Programming. *ACM Trans. Program. Lang. Syst.* 36, 4 (2014), 12:1–12:44. https://doi.org/10.1145/2629609

[13] Simon J. Gay and Vasco Thudichum Vasconcelos. 2010. Linear type theory for asynchronous session types. *J. Funct. Program.* 20, 1 (2010), 19–50. https://doi.org/10.1017/S0956796809990268

[14] Jean-Yves Girard. 1987. Linear Logic. *Theor. Comput. Sci.* 50 (1987), 1–102. https://doi.org/10.1016/0304-3975(87)90045-4

[15] Dan Grossman, J. Gregory Morrisett, Trevor Jim, Michael W. Hicks, Yanling Wang, and James Cheney. 2002. Region-Based Memory Management in Cyclone. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany, June 17-19, 2002*, Jens Knoop and Laurie J. Hendren (Eds.). ACM, 282–293. https://doi.org/10.1145/512529.512563

[16] Rich Hickey. 2017. clojure/PersistentHashMap.java. https://github.com/richhickey/clojure/blob/master/src/jvm/clojure/lang/Persist

[17] Kohei Honda. 1993. Types for Dyadic Interaction. In *Proceedings of 4th International Conference on Concurrency Theory (LNCS)*, Eike Best (Ed.). Springer Verlag, 509–523.

[18] Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *Programming Languages and Systems - ESOP'98, 7th European Symposium on Programming, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings (Lecture Notes in Computer Science)*, Chris Hankin (Ed.), Vol. 1381. Springer, 122–138. https://doi.org/10.1007/BFb0053567

[19] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: Securing the Foundations of the Rust Programming Language. *Proc. ACM Program. Lang.* 2, POPL (2018), 66:1–66:34. https://doi.org/10.1145/3158154

[20] Sam Lindley and J. Garrett Morris. 2017. Lightweight Functional Session Types. In *Behavioral Types: From Theory to Tools*, Simon Gay and António Ravara (Eds.). River Publishers.

[21] Nicholas D. Matsakis and Felix S. Klock II. 2014. The Rust Language. In *Proceedings of the 2014 ACM SIGAda annual conference on High integrity language technology, HILT 2014, Portland, Oregon, USA, October 18-21, 2014*, Michael Feldman and S. Tucker Taft (Eds.). ACM, 103–104. https://doi.org/10.1145/2663171.2663188

[22] Kazutaka Matsuda. 2019. A Modular Inference of Linear Types for Multiplicity-Annotated Arrows. *CoRR* abs/1911.00268 (2019). arXiv:1911.00268 http://arxiv.org/abs/1911.00268

[23] Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. 2010. Lightweight Linear Types in System F°. In *Proceedings of TLDI 2010: 2010 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, Madrid, Spain, January 23, 2010*, Andrew Kennedy and Nick Benton (Eds.). ACM, 77–88. https://doi.org/10.1145/1708016.1708027

[24] J. Garrett Morris. 2016. The Best of Both Worlds: Linear Functional Programming Without Compromise. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 448–461. https://doi.org/10.1145/2951913.2951925

[25] Guillaume Munch-Maccagnoni. 2018. Resource Polymorphism. *CoRR* abs/1803.02796 (2018). arXiv:1803.02796 http://arxiv.org/abs/1803.02796

[26] Martin Odersky, Martin Sulzmann, and Martin Wehr. 1999. Type Inference with Constrained Types. *TAPOS* 5, 1 (1999), 35–55.

[27] Martin Odersky, Philip Wadler, and Martin Wehr. 1995. A Second Look at Overloading. In *Proceedings of the seventh international conference on Functional programming languages and computer architecture, FPCA 1995, La Jolla, California, USA, June 25-28, 1995*, John Williams (Ed.). ACM, 135–146. https://doi.org/10.1145/224164.224195

[28] Scott Owens, Magnus O. Myreen, Ramana Kumar, and Yong Kiam Tan. 2016. Functional Big-Step Semantics. In *ESOP (Lecture Notes in Computer Science)*, Vol. 9632. Springer, 589–615.

[29] Luca Padovani. 2017. A Simple Library Implementation of Binary Sessions. *J. Funct. Program.* 27 (2017), e4. https://doi.org/10.1017/S0956796816000289

[30] François Pottier and Vincent Simonet. 2002. Information Flow Inference for ML. In *Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, OR, USA, January 16-18, 2002*, John Launchbury and John C. Mitchell (Eds.). ACM, 319–330. https://doi.org/10.1145/503272.503302

[31] Jonathan Protzenko. 2014. *Mezzo: a typed language for safe effectful concurrent programs. (Mezzo: un langage typé pour programmer de manière concurrent et sure en présence d'effets).* Ph.D. Dissertation. Paris Diderot University, France. https://tel.archives-ouvertes.fr/tel-01086106

[32] Juan Pedro Bolívar Puente. 2017. Persistence for the Masses: RRB-Vectors in a Systems Language. *PACMPL* 1, ICFP (2017), 16:1–16:28. https://doi.org/10.1145/3110260

[33] John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*. IEEE Computer Society, 55–74. https://doi.org/10.1109/LICS.2002.1029817

[34] Jeremy Siek. 2013. Type Safety in Three Easy Lemmas. http://siek.blogspot.de/2013/05/type-safety-in-three-easy-lemmas.html.

[35] Vincent Simonet. 2003. An extension of HM(X) with bounded existential and universal data-types. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming, ICFP 2003, Uppsala, Sweden, August 25-29, 2003*, Colin Runciman and Olin Shivers (Eds.). ACM, 39–50. https://doi.org/10.1145/944705.944710

[36] Vincent Simonet. 2003. Type Inference with Structural Subtyping: A Faithful Formalization of an Efficient Constraint Solver. In *Programming Languages and Systems, First Asian Symposium, APLAS 2003, Beijing, China, November 27-29, 2003, Proceedings (Lecture Notes in Computer Science)*, Atsushi Ohori (Ed.), Vol. 2895. Springer, 283–302. https://doi.org/10.1007/978-3-540-40018-9_19

[37] Vincent Simonet and François Pottier. 2007. A Constraint-Based Approach to Guarded Algebraic Data Types. *ACM Trans. Program. Lang. Syst.* 29, 1 (2007), 1. https://doi.org/10.1145/1180475.1180476

[38] Christian Skalka and François Pottier. 2002. Syntactic Type Soundness for HM(X). *Electr. Notes Theor. Comput. Sci.* 75 (2002), 61–74. https://doi.org/10.1016/S1571-0661(04)80779-5

[39] Martin Sulzmann. 1997. *Proofs of soundness and completeness of type inference for HM (X).* Technical Report. Research Report YALEU/DCS/RR-1102, Yale University, Department of Computer Science.

[40] Jesse A. Tov and Riccardo Pucella. 2011. Practical affine types. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, Thomas Ball

and Mooly Sagiv (Eds.). ACM, 447–458. https://doi.org/10.1145/1926385.1926436

[41] Valery Trifonov and Scott F. Smith. 1996. Subtyping Constrained Types. In *Static Analysis, Third International Symposium, SAS'96, Aachen, Germany, September 24-26, 1996, Proceedings (Lecture Notes in Computer Science)*, Radhia Cousot and David A. Schmidt (Eds.), Vol. 1145. Springer, 349–365. https://doi.org/10.1007/3-540-61739-6_52

[42] Aaron Weiss, Daniel Patterson, Nicholas D. Matsakis, and Amal Ahmed. 2019. Oxide: The Essence of Rust. *CoRR* abs/1903.00982 (2019). arXiv:1903.00982 http://arxiv.org/abs/1903.00982

A FURTHER EXAMPLES

A.1 A session on linearity

Session typing [17, 18] is a type discipline for checking protocols statically. A session type is ascribed to a communication channel and describes a sequence of interactions. For instance, the type `int!int!int?end` specifies a protocol where the program must send two integers, receive an integer, and then close the channel. In this context, channels must be used linearly, because every use of a channel "consumes" one interaction and changes the type of the channel. That is, after sending two integers on the channel, the remaining channel has type `int?end`. Here are some standard type operators for session types $\mathcal{S}$:

$\tau!\mathcal{S}$	Send a value of type τ then continue with $\mathcal{S}$.
$\tau?\mathcal{S}$	Receive a value of type τ then continue with $\mathcal{S}$.
$\mathcal{S} \oplus \mathcal{S}'$	Internal choice between protocols $\mathcal{S}$ and $\mathcal{S}'$.
$\mathcal{S} \,\&\, \mathcal{S}'$	Offer a choice between protocols $\mathcal{S}$ and $\mathcal{S}'$.

Padovani [29] has shown how to encode this style of session typing in ML-like languages, but his implementation downgrades linearity to a run-time check for affinity. Building on that encoding we can provide a safe API in Affe that statically enforces linear handling of channels:

```
1 type S st : lin
2 val receive: (α ? S) st → α × S st
3 val send : (α ! S) st → α -lin-> S st
4 val create : unit → S st × (dual S) st
5 val close : end st → unit
```

Line 1 introduces a parameterized abstract type `st` which is linear as indicated by its kind `lin`. Its low-level implementation would wrap a handle for a socket, for example. The `receive` operation in Line 2 takes a channel that is ready to receive a value of type α and returns a pair of the value a the channel at its residual type $\mathcal{S}$. It does not matter whether α is restricted to be linear, in fact `receive` is polymorphic in the kind of α. This kind polymorphism is the default if no constraints are specified. The `send` operation takes a linear channel and returns a single-use function that takes a value of type α suitable for sending and returns the channel with updated type. The `create` operation returns a pair of channel endpoints. They follow dual communication protocols, where the `dual` operator swaps sending and receiving operations. Finally, `close` closes the channel.

In Fig. 15 we show how to use these primitives to implement client and server for an addition service. No linearity annotations are needed in the code, as all linearity properties can be inferred from the linearity of the `st` type.

The inferred type of the server, `add_service`, is `(int ! int ! int ? end)st →unit`. The client operates by sending two messages and receiving the result. This code is polymorphic in both argument and return types, so it could be used with any binary operator. Moreover, the `op_client` function can

```
1 let add_service ep =
2   let x, ep = receive ep in
3   let y, ep = receive ep in
4   let ep = send ep (x + y) in
5   close ep
6 # add_service : (int ! int ! int ? end) st → unit
```

(a) Addition server

```
1 let op_client ep x y =
2   let ep = send ep x in
3   let ep = send ep y in
4   let result, ep = receive ep in
5   close ep;
6   result
7 # op_client :
8     (α1 ? α2 ? β ! end) st → α1 -lin-> α2 -lin-> β
```

(b) Binary operators client

Fig. 15. Corresponding session type programs in Affe

be partially applied to a channel. Since the closure returned by such a partial application captures the channel, it can only be used once. This restriction is reflected by the arrow of kind `lin`, $\xrightarrow{\texttt{lin}}$, which is the type of a single-use function. The general form of arrow types in Affe is $\xrightarrow{k}$, where k is a kind that restricts the number of uses of the function. For convenience, we shorten $\xrightarrow{\texttt{un}}$ to $\to$. Affe infers the single-use property of the arrows without any user annotation. In fact, the only difference between the code presented here and Padovani's examples [29] is the kind annotation on the type definition of `st`.

To run client and server, we can create a channel and apply `add_service` to one end and `op_client` to the other. Failure to consume either channel endpoints (`a` or `b`) would result in a type error.

```
1 let main () =
2   let (a, b) = create () in
3   fork add_service a;
4   op_client b 1 2
5 # main : unit ⟶ int
```

A.2 Pool of linear resources

We present an interface and implementation of a pool of linear resources where the extended scope of the region enforces proper use of the resources.

Fig. 16a contains the interface of the `Pool` module. A pool is parameterized by its content. The kind of the pool depends on the content: linear content implies a linear pool while unrestricted content yields an unrestricted pool. The functions `Pool.create` and `Pool.consume` build/destroy a pool given creators/destructors for the elements of the pool. The function `Pool.use` is the workhorse of the API, which borrows a resource from the pool to a callback. It takes a shared borrow of a pool (to enable concurrent access) and a callback function. The callback receives a exclusive borrow of an arbitrary resource from the pool. The typing of the callback ensures that this borrow is neither captured nor returned by the function.

This encapsulation is implemented with a universally quantified *kind index variable* r. The signature prescribes the type `&!(`$\texttt{aff}_{r+1}$`,`α_1`)` for the exclusive borrow of the resource with an affine kind at region nesting $r + 1$. The return type of the callback is constrained to kind $\kappa_2 \leq \texttt{aff}_r$ so that the callback certainly cannot return the borrowed argument. In a specific use of `Pool.use`, the index r gets unified with the current nesting level of regions so that the region for the callback effectively gets "inserted" into the lexical nesting at the callsite. Fig. 16b shows a simple example using the `Pool` module.

The implementation in Fig. 16c represents a bag of resources using a concurrent queue with atomic add and remove operations. The implementation of the `Pool.create` and `Pool.consume` functions is straightforward. The function `Pool.use` first draws an element from the pool (or creates a fresh element), passes it down to the callback function `f`, and returns it to the pool afterwards. For clarity, we explicitly delimit the region in Line 11 to ensure that the return value of `f &!o` does not capture `&!o`. In practice, the type checker inserts this region automatically.

B AUTOMATIC REGION ANNOTATION

We now define our automatic region annotation which is presented in Section 3.2. First, we extend the region annotation to $\{\!|E|\!\}^n_S$ where S is a map from variables to borrow indicator b. This annotation, defined below, is equivalent to nested region annotations for each individual variable.

$$\{\!|e|\!\}^n_{\{x\mapsto b\};S} = \{\!|\{\!|e|\!\}^n_S|\!\}^n_{\{x\mapsto b\}} \qquad\qquad \{\!|e|\!\}^n_\emptyset = e$$

Figure 17 define a rewriting relation $e \rightsquigarrow e'$ which indicates that an optionally annotated term

```
type (α:κ) pool : κ
create : (unit → α) → α pool
consume : (α → unit) → α pool → unit
use : (α1:κ1),(α2:κ2),(κ2 ≤ aff_r) ⇒
  &(α1 pool) → (&!(aff_{r+1},α1) -lin-> α2) -κ1-> α2
```

(a) Signature

```
(*Using the pool in queries.*)
let create_user pool name =
  Pool.use &pool (fun connection →
    Db.insert "users" [("name", name)] connection)

let uri = "postgresql://localhost:5432"
let main users =
  (*Create a database connection pool.*)
  let pool = Pool.create (fun _ → Db.connect uri) in
  List.parallel_iter (create_user &pool) users;
  Pool.consume (fun c → Db.close c)
```

(b) Example of use

```
type (α:κ) pool : κ =
  { spawn: unit → α; queue: α CQueue.t }
let create spawn =
  { spawn ; queue = CQueue.create () }
let consume f c = CQueue.iter f c.queue
let use { spawn ; queue } f =
  let o = match CQueue.pop &queue with
    | Some x → x
    | None () → spawn ()
  in
  let r = {| f &!o |} in
  CQueue.push o &queue;
  r
```

(c) Implementation

Fig. 16. The Pool module

$$
\begin{array}{rcll}
\{x \mapsto b\} \oplus \cdot & = & \cdot, \{x \mapsto b\}, \cdot & \text{AnnotRegion-Left} \\
\cdot \oplus \{x \mapsto b\} & = & \cdot, \{x \mapsto b\}, \cdot & \text{AnnotRegion-Right} \\
\{x \mapsto \mathbf{U}\} \oplus \{x \mapsto \mathbf{U}\} & = & \cdot, \{x \mapsto \mathbf{U}\}, \cdot & \text{AnnotRegion-Immut} \\
\{x \mapsto \mathbf{U}\} \oplus \{x \mapsto \mathbf{A}\} & = & \{x \mapsto \mathbf{U}\}, \{x \mapsto \mathbf{A}\}, \cdot & \text{AnnotRegion-MutLeft} \\
\{x \mapsto \mathbf{A}\} \oplus \{x \mapsto b\} & = & \{x \mapsto \mathbf{A}\}, \cdot, \{x \mapsto b\} & \text{AnnotRegion-MutRight}
\end{array}
$$

$$
\frac{e = \&^b x \mid \&\&^b x}{e \rightsquigarrow_n e, \{x \mapsto b\}} \qquad \frac{e = c \mid x}{e \rightsquigarrow_n e, \cdot} \qquad \frac{\forall i,\ e_i \rightsquigarrow_{n+1} e_i', B_i \qquad B_1 \oplus B_2 = S_1, S, S_2}{(e_1\ e_2) \rightsquigarrow_n (\{\!|e_1'|\!\}_{S_1}^{n+1}\ \{\!|e_2'|\!\}_{S_2}^{n+1}), S}
$$

$$
\frac{\forall i,\ e_i \rightsquigarrow_{n+1} e_i', B_i \qquad B_1 \oplus (B_2 \backslash \{x\}) = S_1, S, S_2 \qquad S_2' = S_2 \cup B_2\big|_x}{\texttt{let}\ x = e_1\ \texttt{in}\ e_2 \rightsquigarrow_n \texttt{let}\ x = \{\!|e_1'|\!\}_{S_1}^{n+1}\ \texttt{in}\ \{\!|e_2'|\!\}_{S_2'}^{n+1}, S}
$$

Rewrite-Lam

$$
\frac{e \rightsquigarrow_{n+1} e', B \qquad B_x = B\big|_x}{\lambda x.e \rightsquigarrow_n \lambda x.\{\!|e'|\!\}_{B_x}^{n+1}, B \backslash \{x\}}
$$

$$
\frac{\forall i,\ e_i \rightsquigarrow_{n+1} e_i', B_i \qquad B_1 \oplus (B_2 \backslash \{x, y\}) = S_1, S, S_2 \qquad S_2' = S_2 \cup B_2\big|_{x,y}}{\texttt{match}_\phi\ x, y = e_1\ \texttt{in}\ e_2 \rightsquigarrow_n \texttt{match}_\phi\ x, y = \{\!|e_1'|\!\}_{S_1}^{n+1}\ \texttt{in}\ \{\!|e_2'|\!\}_{S_2'}^{n+1}, S}
$$

Rewrite-Region

$$
\frac{e \rightsquigarrow_{n+1} e', B}{\{\!|e|\!\} \rightsquigarrow_n \{\!|e'|\!\}_B^{n+1}, \cdot}
$$

Rewrite-Pair

$$
\frac{\forall i,\ e_i \rightsquigarrow_{n+1} e_i', B_i \qquad B_1 \oplus B_2 = S_1, S, S_2}{(e_1, e_2) \rightsquigarrow_n (\{\!|e_1'|\!\}_{S_1}^{n+1}, \{\!|e_2'|\!\}_{S_2}^{n+1}), S}
$$

AnnotRegion

$$
\frac{\forall i,\ b_i \oplus b_i' = S_i^l, S_i^m, S_i^r}{\overline{b_i} \oplus \overline{b_i'} = \cup_i S_i^l, \cup_i S_i^m, \cup_i S_i^r}
$$

Rewrite-Top

$$
\frac{e \rightsquigarrow_1 e', S}{e \rightsquigarrow \{\!|e'|\!\}_S^1}
$$

Fig. 17. Automatic region annotation — $e \rightsquigarrow e'$

e can be rewritten as a fully annotated term e'. Through the rule Rewrite-Top, this is defined in term of an inductively defined relation $e \rightsquigarrow_n e', S$ where n is the current nesting and S is a set of variable that are not yet enclosed in a region. The base cases are constants, variables and borrows. The general idea is to start from the leaves of the syntax tree, create a region for each borrow, and enlarge the region as much as possible. This is implemented by a depth-first walk of the syntax tree

$$C ::= (\tau_1 \le \tau_2) \mid (k_1 \le k_2) \mid C_1 \wedge C_2 \mid \exists \alpha.C \mid \exists \kappa.C$$

Fig. 18. The constraint language

$$\frac{l \le_{\mathcal{L}} l'}{\vdash_{\mathrm{e}} (l \le l')} \qquad \frac{\forall i,\ C \vdash_{\mathrm{e}} (l_i \le k)}{C \vdash_{\mathrm{e}} (\wedge_i l_i \le k)} \qquad \frac{\forall i,\ C \vdash_{\mathrm{e}} (k \le l_i)}{C \vdash_{\mathrm{e}} (k \le \vee_i l_i)} \qquad \frac{C \vdash_{\mathrm{e}} (k \le k') \wedge (\tau \le \tau')}{C \vdash_{\mathrm{e}} (\&^b(k,\tau) \le \&^b(k',\tau'))}$$

$$\frac{C \vdash_{\mathrm{e}} (\tau_1' \le \tau_1) \quad C \vdash_{\mathrm{e}} (\tau_2 \le \tau_2') \quad C \vdash_{\mathrm{e}} (k \le k')}{C \vdash_{\mathrm{e}} (\tau_1 \xrightarrow{\mathrm{k}} \tau_2 \le \tau_1' \xrightarrow{\mathrm{k'}} \tau_2')} \qquad \frac{\forall i,\ C \vdash_{\mathrm{e}} (\tau_i = \tau_i)}{C \vdash_{\mathrm{e}} (\mathrm{T}\ \overline{\tau_i} \le \mathrm{T}\ \overline{\tau_i'})} \qquad \le \text{ transitive, reflexive}$$

Fig. 19. Base entailment rules – $C \vdash_{\mathrm{e}} D$

which collects each variable that has a corresponding borrow. At each step, it rewrites the inner subterms, consider which borrow must be enclosed by a region now, and return the others for later enclosing. Binders force immediate enclosing of the bound variables, as demonstrated in rule Rewrite-Lam. For nodes with multiple children, we use a scope merge operator to decide if regions should be placed and where. This is shown in rule Rewrite-Pair. The merge operator, written $B_l \oplus B_r = (S_l, S, S_r)$, takes the sets B_l and B_r returned by rewriting the subterms and returns three sets: S_l and S_r indicates the variables that should be immediately enclosed by a region on the left and right subterms and S indicates the set of the yet-to-be-enclosed variables. As an example, the rule AnnotRegion-MutLeft is applied when there is an shared borrow and a exclusive borrow. In that case, a region is created to enclose the shared borrow, while the exclusive borrow is left to be closed later. This is coherent with the rules for environment splitting and suspended bindings from Section 3.3. Explicitly annotated regions are handled specially through rule Rewrite-Region. In that case, we assume that all inner borrows should be enclosed immediately.

C CONSTRAINTS

We place our constraint system in a more general setting. We define the constraint solver in terms of an arbitrary commutative bounded lattice $(\mathcal{L}, \le_{\mathcal{L}})$, i.e., a lattice which has a minimal and a maximal element ($l^\top$ and $l^\bot$) and where meet and joins are commutative. We write lattice elements as l and $\bigwedge_i l_i$ (resp. $\bigvee_i l_i$) for the greatest lower bound (resp. least upper bound) in $\mathcal{L}$. The lattice for Affe (see Section 3.3) is a bounded lattice with $l^\top = \mathbf{L}_\infty$ and $l^\bot = \mathbf{U}_0$.

Let $\mathcal{C}_{\mathcal{L}}$ be the set of constraints in such a lattice $\mathcal{L}$. The full grammar of constraints is shown in Fig. 18. Constraints are made of kind inequalities, conjunctions and projections along with type unification constraints. Since types might contain kinds (for instance, on the arrows), type unification is oriented and written as $\le$. For simplicity, we consider all type constructors invariant in their parameters and define $(\tau = \tau')$ as $(\tau \le \tau') \wedge (\tau' \le \tau)$.

Entailment is denoted by $C \vdash_{\mathrm{e}} D$, where D is a consequence of the constraints C. We say that C and D are equivalent, $C =_{\mathrm{e}} D$, when $C \vdash_{\mathrm{e}} D$ and $D \vdash_{\mathrm{e}} C$.

We directly reuse the following definitions from HM(X).

Property 3 (Cylindric constraint system). *A cylindric constraint system is a constraint system such that, for any constraint C:*

$$C \vdash_{\mathrm{e}} \exists x.C \qquad\qquad C \vdash_{\mathrm{e}} D \implies \exists x.C \vdash_{\mathrm{e}} \exists x.D$$
$$\exists x.(C \wedge \exists x.D) =_{\mathrm{e}} \exists x.C \wedge \exists x.D \qquad\qquad \exists x.\exists y.D =_{\mathrm{e}} \exists y.\exists x.D$$

Property 4 (Term rewriting system). *A term rewriting system is a system where, for every types τ,τ', there exists an equality predicates $(\tau = \tau')$ which is symmetric, reflexive, transitive, stable under substitution and such that, for any predicate P:*

$$(x = y) \wedge \exists x.C \wedge (x = y) \vdash_{\mathrm{e}} C$$
$$P[x \to \tau] =_{\mathrm{e}} \exists x.P \wedge (x = \tau) \text{ where } x \notin \mathrm{fv}(\tau)$$

Definition C.1 (Constraint system with lattice). $C_{\mathcal{L}}$ is defined as the smallest cylindric term constraint system that satisfies the axiom shown in Fig. 19.

We define the set of solved formed $\mathcal{S}$ as the quotient set of $C_{\mathcal{L}}$ by $=_{\mathrm{e}}$. We will show later that such constraints are in fact only composed of kind inequalities, and thus correspond to the syntactic constraints used in type and kind schemes. We now define our normalization procedure normalize(C_0, ψ_0), where $C_0 \in C_{\mathcal{L}}$ is a set of constraints and ψ_0 is a substitution. It returns a constraint $C \in \mathcal{S}$ in solved form and a unifier ψ. The main idea of the algorithm is to first remove all the type equalities by using regular Herbrand unification. After that, we only have a set of inequalities among kinds, which we can consider as a relation. We can then saturate the relation, unify all kinds that are in the same equivalence classes to obtain a most general unifier on kind variables, remove all existentially quantified variables and then minimize back the relation and apply various simplification rules to make the resulting type easier to understand to users.

More precisely, we apply the following steps:

(1) Solve all type equality constraints through Herbrand unification and gather all existential quantifications at the front of the constraint. We obtain a constraint $C^k = \exists \overline{\kappa}.\ (k_j \le k'_j)_j$ and a substitution ψ_τ.
   We write $\mathcal{R}$ for the relation $(k_j \le k'_j)_j$, $\mathcal{G}$ the underlying directed graph and V its vertices.
(2) Saturate the lattice equalities in $\mathcal{R}$.
   More precisely, for each kind variable $\kappa \in V$, for each constant l_i (resp. l_j) such that there is a path from l_i to κ (resp. from κ to l_j) in $\mathcal{G}$, add an edge from $\bigvee l_i$ to κ (resp. from κ to $\bigwedge l_j$). This step is well defined since $\mathcal{L}$ is a bounded lattice and $\bigvee \emptyset$ and $\bigwedge \emptyset$ are well defined.
   We also complement $\mathcal{R}$ with $(\le)$ by adding an edge between related constants.
(3) At this point, we can easily check for satisfiability: A constraint is satisfiable (in the given environment) if and only if, for any constants l_1 and l_2 such that there is a path from l_1 to l_2 in $\mathcal{G}$, then $l_1 \le_{\mathcal{L}} l_2$. If this is not the case, we return **fail**.
(4) For each strongly connected component in $\mathcal{G}$, unify all its vertices and replace it by a representative. We write ψ_k for the substitution that replaces a kind variable by its representative. The representative of a strongly connected component g can be determined as follows:
   - If g does not contain any constant, then the representative is a fresh kind variable.
   - If g contains exactly one constant, it is the representative.
   - Otherwise, the initial constraint C_0 is not satisfiable.

   Note that this step will also detect all unsatisfiable constraints.
(5) Take the transitive closure of $\mathcal{R}$.
(6) Remove all the vertices corresponding to the kind variables $\overline{\kappa}$ that are existentially quantified in C^k.
(7) Take the transitive reduction of $\mathcal{R}$.
(8) Remove the extremums of $\mathcal{L}$ and the edges of $(\le)$ from $\mathcal{R}$.
(9) Return $C = \{k \le k' \mid k \,\mathcal{R}\, k'\}$ and $\psi = \psi_\tau \sqcup \psi_k$.

An example of this algorithm in action is shown in Section 4.3.1. Our algorithm is complete, computes principal normal forms, and already simplifies constraints significantly (thanks to steps

6, 7 and 8). It can be extended with further simplification phases. In particular, our implementation and all the signatures presented in Section 2 use a variance-based simplification where all covariant (resp. contravariant) variables are replaced by their lower (resp. upper) bounds. All the simplification mechanisms presented here, including the variance-based one, are complete. It is also possible to add "best-effort" simplification rules which help reduce the size of inferred signatures even further [36].

C.1 Principal constraint system

We now prove that $C_{\mathcal{L}}$ supports all the properties necessary for principal type inference, as defined by HM(X). We first prove that constraint solving does compute normal forms, and that such normal forms are unique.

Lemma C.2 (Principal normal form). *Given a constraint $D \in C_{\mathcal{L}}$, a substitution ϕ and $(C, \psi) =$ normalize(D, ϕ), then $\phi \leq \psi$, $C =_{\mathrm{e}} \psi D$ and $\psi C = C$.*

Proof. Let us partition ϕ into a part which affects type variables, ϕ_τ, and a part which affects kind variables, ϕ_k.

We write (C^k, ψ_τ) for the result of the modified Herbrand unification on (D, ϕ) in step (1). Herbrand unification computes the most general unifier. Our modified Herbrand unification only output additional kind constraints for kind on the arrows and does not change the result of the unification. Thus, we have $\phi_\tau \leq \psi_\tau$, $C^k =_{\mathrm{e}} \psi_\tau D$ and $\psi_\tau C^k = C^k$.

Let C^{k+} be the result after step (2), we trivially have that $\mathrm{fv}(C^{k+}) = \mathrm{fv}(C^k)$ and that $C^{k+} =_{\mathrm{e}} C^k$.

Let C^A and ψ_k be the results after step (4). By definition, we have $\psi_k C^{k+} =_{\mathrm{e}} C^A$ and $\psi_k C^A = C^A$. Since ϕ_k has already be applied to C before unifying the strongly connected components, we have that $\phi_k \leq \psi_k$.

Let $\psi = \psi_\tau \sqcup \psi_k$. Since ψ_τ and ψ_k have disjoint supports, we have $C^A = \psi_\tau C^A =_{\mathrm{e}} \psi C^{k+} =_{\mathrm{e}} \psi D$ and $\psi C^A = C^A$. Furthermore, $\phi_\tau \sqcup \phi_k \leq \psi_\tau \sqcup \psi_k$.

Steps (5) to (9) all preserve the free variables and the equivalence of constraints, which concludes. □

Lemma C.3 (Uniqueness). *Given (C_1, ψ_1) and (C_2, ψ_2) such that $\psi_1 C_1 =_{\mathrm{e}} \psi_2 C_2$, then* normalize$(C_1, \psi_1)$ *and* normalize(C_2, ψ_2) *are identical up to α-renaming.*

Proof. In Lemma C.2, we have showed that all the steps of the normalization procedure preserve equivalence. Since $\psi_1 C_1 =_{\mathrm{e}} \psi_2 C_2$, equivalence between the two results of the normalization procedures is preserved for all steps.

We write $P(C_a)$ if for all $C = (k, k)'$ such that $C_a \vdash_{\mathrm{e}} C$ and $\nvdash_e C$, we have $C \in \mathcal{R}_a$.

Let us write C_1' and C_2' for the constraints after step (4). $P(C_1')$ and $P(C_2')$ hold. Indeed, since C_1' and C_2' are only composed of existential quantifications and kind inequalities, the only rules that applies are transitivity and lattice inequalities. After step (2) and (5), the associated relations are fully saturated for these two rules, hence all inequalities that can be deduced from C_a' are already present in the relation.

The property P is preserved by step (6) since we only remove inequalities that involve existentially quantified variables. Such inequalities could not be picked in P.

Let us write C_a'' for $a \in \{1, 2\}$ the constraints after step (5). Since there are no more existential variables, we have $C_a'' = (k_i, k_i')_i = \mathcal{R}_a''$. For any $C = (k, k')$ such that $\vdash_{\mathrm{e}} C$ and $C_a'' \vdash_{\mathrm{e}} C$, then $C \in (\leq) \subset \mathcal{R}_a''$. Indeed, the only trivial inequalities in our system are equalities of the form (κ, κ), which were removed in step (4) and the lattice inequalities.

$$\textsc{KVar}\ \frac{(\alpha : k) \in \Gamma}{C \mid \Gamma \vdash_{\mathrm{s}} \alpha : k} \qquad \textsc{KApp}\ \frac{(\mathrm{T} : \forall \overline{\kappa_i}.\, D \Rightarrow (\overline{k'_j}) \to k') \in \Gamma \quad \psi = [\overline{\kappa_i} \to \overline{k_i}] \quad C \vdash_{\mathrm{e}} \psi D \quad \forall j \quad C \mid \Gamma \vdash_{\mathrm{s}} \tau_j : k_j \quad C \mid \Gamma \vdash_{\mathrm{s}} k_j \leq \psi k'_j}{C \mid \Gamma \vdash_{\mathrm{s}} \mathrm{T}\ \overline{\tau_j} : \psi k'}$$

$$\textsc{KPair}\ \frac{\forall i \quad C \mid \Gamma \vdash_{\mathrm{s}} \tau_i : k_i \quad C \mid \Gamma \vdash_{\mathrm{s}} k_i \leq k}{C \mid \Gamma \vdash_{\mathrm{s}} \tau_1 \times \tau_2 : k} \qquad \textsc{KBorrow}\ C \mid \Gamma \vdash_{\mathrm{s}} \&^b(k, \tau) : k \qquad \textsc{KArr}\ C \mid \Gamma \vdash_{\mathrm{s}} \tau_1 \xrightarrow{\mathrm{k}} \tau_2 : k$$

Fig. 20. Syntax-directed kinding rule – $C \mid \Gamma \vdash_{\mathrm{s}} \tau : k$

Let us consider $C = (k, k') \in \mathcal{R}''_1$. Since $C''_1 =_{\mathrm{e}} C''_2$, we have $C''_2 \vdash_{\mathrm{e}} C$. If $\nvdash_e C$, by $P(C''_2)$ we have that $C \in R''_2$. If $\vdash_{\mathrm{e}} C$, then $C \in (\leq) \subset R''_2$. We conclude that $R''_1 \subset R''_2$. By symmetry, $R''_1 = R''_2$ and $C''_1 = C''_2$.

This equality is preserved by step (7) and (8) since the transitive reduction of a directed acyclic graph is unique, which concludes. □

We can now prove all the necessary high level properties.

Lemma C.4. *For all $C \in \mathcal{S}$, $C \vdash_{\mathrm{e}} x = x$ implies $\vdash_{\mathrm{e}} x = x$.*

Proof. By definition of normalize, We have $C = \overline{(k \leq k')}$ such that the underlying relation has no cycles. Thus, we can not deduce neither kind nor type equalities from C. □

Property 5 (Regular constraint system). *$C_{\mathcal{L}}$ is regular, ie, for x, x' two types or kinds, $\vdash_{\mathrm{e}}(x = x')$ implies $\mathrm{fv}(x) = \mathrm{fv}(x')$*

Proof. The only equalities possibles are between variables (via symmetry) or between constants. □

Finally, we can conclude with all the properties we need for HM(X):

Theorem C.5 (Principal constraints). *$C_{\mathcal{L}}$ has the principal constraint property,* normalize *computes principal normal forms for $C_{\mathcal{L}}$ and $C_{\mathcal{L}}$ is regular.*

This is sufficient to show that $HM(C_{\mathcal{L}})$ is principal. However, we do not use HM(X) directly but an extended version with kind inference, linear and affine types, and borrow. We extend the proofs of HM(X) to such a system in Appendix E.

D SYNTAX-DIRECTED TYPING

D.1 Kinding

We write $C \mid \Gamma \vdash_{\mathrm{s}} \tau : k$ if τ has kind k in environment Γ under constraints C. The rules are shown in Fig. 20. Kinds and types follow a small calculus with variables (α,…), functions (type constructors t), application ($\mathrm{T}\ \overline{\tau}$) and primitives such as types for arrows ($\tau \xrightarrow{\mathrm{k}} \tau'$) and borrows ($\&^b(k, \tau)$). Kind checking can thus be done in a fairly straightforward, syntax-directed fashion by simply following the syntax of the types. Kind arrows can only appear when looking up the kind scheme of a type constructor t. Kind arrows are forbidden in any other contexts.

$$\frac{D \Leftarrow \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad C \vdash_{\mathrm{e}} D}{C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2}\ \textsc{ESplit-Check} \qquad \cdot \Leftarrow \cdot = \cdot \ltimes \cdot\ \textsc{ESplit-Empty} \qquad \frac{C_1 \Leftarrow \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad C_2 \Leftarrow b = b_1 \ltimes b_2}{C_1 \wedge C_2 \Leftarrow \Gamma; b = \Gamma_1; b_1 \ltimes \Gamma_2; b_2}\ \textsc{ESplit-Nonempty}$$

$$\begin{aligned}
(\sigma \le \mathbf{U}_\infty) \vdash_e (x:\sigma) &= (x:\sigma) \ltimes (x:\sigma) && \text{(Both)}\\
\cdot \vdash_e (x \div \sigma)^k_{\mathrm{U}} &= (x \div \sigma)^k_{\mathrm{U}} \ltimes (x \div \sigma)^k_{\mathrm{U}} && \text{(Borrow)}\\
\cdot \vdash_e B_x &= B_x \ltimes \emptyset && \text{(Left)}\\
\cdot \vdash_e B_x &= \emptyset \ltimes B_x && \text{(Right)}\\
\cdot \vdash_e (x:\sigma) &= [x:\sigma]^n_b \ltimes (x:\sigma) && \text{(Susp)}\\
(b' \le b) \vdash_e (x \div \sigma)^k_b &= [x:\sigma]^n_{b'} \ltimes (x \div \sigma)^k_b && \text{(SuspB)}\\
\cdot \vdash_e [x:\sigma]^{n'}_b &= [x:\sigma]^n_{\mathrm{U}} \ltimes [x:\sigma]^{n'}_b && \text{(SuspS)}
\end{aligned}$$

Fig. 21. Splitting — environments $C \vdash_e \Gamma = \Gamma_l \ltimes \Gamma_r$; inference $C \Leftarrow \Gamma = \Gamma_l \ltimes \Gamma_r$; binders $C \Leftarrow b = b_r \ltimes b_l$

$$\frac{C_r \Leftarrow [x:\tau]^n_b \rightsquigarrow^x_n b}{C_r \Leftarrow \Gamma; [x:\tau]^n_b \rightsquigarrow^x_n \Gamma; b}\ \textsc{EBorrow} \qquad \frac{C \vdash_{\mathrm{e}} D \qquad D \Leftarrow \Gamma; [x:\tau]^n_b \rightsquigarrow^x_n \Gamma; b}{C \vdash_e \Gamma; [x:\tau]^n_b \rightsquigarrow^x_n \Gamma; b}\ \textsc{EBorrow-Check}$$

$$\frac{b \in \{\mathbf{U}, \mathbf{A}\} \qquad C = (b_n \le k) \wedge (k \le b_\infty)}{C \Leftarrow [x:\sigma]^n_b \rightsquigarrow^x_n (x \div \sigma)^k_b}\ \textsc{EBorrow-Binder}$$

Fig. 22. Borrowing — environments $C \vdash_e \Gamma \rightsquigarrow^x_n \Gamma'$; inference $C \Leftarrow \Gamma \rightsquigarrow^x_n \Gamma'$; binders $C \Leftarrow b \rightsquigarrow^x_n b'$

D.2 Environments

In Section 3.3, we only gave a partial description of the splitting and borrowing relations on environments, $C \vdash_e \Gamma = \Gamma \ltimes \Gamma$ and $C \vdash_e \Gamma \rightsquigarrow^x_n \Gamma$. The complete definitions are shown on Figs. 21 and 22. All the definitions are made in term of the inference version, which returns fresh constraints. The solving version then simply uses entailment, as shown in rule ESplit-Check and EBorrow-Check. The remaining new rules are dedicated to iterating over the environment.

Fig. 23 defines the rewriting relation on environment constraints, $(\Gamma \le k) \rightsquigarrow C$, which rewrites a constraint of the form $(\Gamma \le k)$ into C. It proceeds by iterating over the environment and expanding the constraints for each binding. Suspended bindings are rejected (ConstrSusp). Borrow bindings directly use the annotated kind (ConstrBorrow). Other bindings use the underlying type scheme (ConstrBinding). Type schemes are constrained by first inferring the kind, and then emitting the constraint (ConstrSD and ConstrI).

D.3 Typing

The rules for syntax-directed typing are shown in Fig. 24 and follow the presentation given in Section 3.3. As usual in HM type systems, introduction of type-schemes is included in the Let rule via generalization. We define $\mathrm{gen}(C, \Gamma, \tau) = (\exists \overline{\kappa}, \overline{\alpha}.C, \forall \overline{\kappa}, \overline{\alpha}.C \Rightarrow \tau)$ where $\overline{\kappa}, \overline{\alpha} = (\mathrm{fv}(\tau) \cup \mathrm{fv}(C)) \setminus \mathrm{fv}(\Gamma)$. The typing rules specific to the internal language are shown in Fig. 25.

$$(\cdot \le k) \rightsquigarrow \cdot \qquad \frac{(\Gamma \le k) \rightsquigarrow C \qquad \Gamma \vdash (B \le k) \rightsquigarrow D}{(\Gamma; B \le k) \rightsquigarrow C \wedge D}$$

ConstrBinding
$$\frac{\Gamma \vdash (\sigma \le k) \rightsquigarrow C}{\Gamma \vdash ((x : \sigma) \le k) \rightsquigarrow C}$$

ConstrBorrow
$$\frac{}{\Gamma \vdash ((\&^b x : \&^b(k', \tau)) \le k) \rightsquigarrow (k' \le k)}$$

ConstrSusp
$$\frac{}{\Gamma \vdash ([x : \sigma]^n_b \le k) \rightsquigarrow \text{False}}$$

ConstrSD
$$\frac{C \wedge C_x \mid \Gamma \vdash_{\text{w}} \tau : k' \qquad D = C \wedge C_x \wedge (k' \le k)}{\Gamma \vdash_{\text{w}} ((\forall \kappa_i \forall (\alpha_j : k_j).\ C_x \Rightarrow \tau) \le k) \rightsquigarrow D}$$

ConstrI
$$\frac{C \wedge C_x \mid \Gamma \vdash_{\text{s}} \tau : k' \qquad D = C \wedge C_x \wedge (k' \le k)}{\Gamma \vdash_{\text{s}} (\forall \kappa_i \forall (\alpha_j : k_j).\ C_x \Rightarrow \tau \le k) \rightsquigarrow D}$$

Fig. 23. Rewriting constraints on environments — $(\Gamma \le k) \rightsquigarrow C$

Instance
$$\frac{\sigma = \forall \overline{\kappa_i} \forall (\overline{\alpha_j : k_j}).\ C \Rightarrow \tau \qquad \psi = [\overline{\kappa_i \mapsto k_i}, \overline{\alpha_j \mapsto \tau_j}]}{\psi(C), \psi(\tau) = \text{Inst}(\Gamma, \sigma)}$$

Var
$$\frac{(x : \sigma) \in \Gamma \qquad C_x, \tau_x = \text{Inst}(\Gamma, \sigma) \qquad C \vdash_{\text{e}} C_x \wedge (\Gamma \backslash \{x\} \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} x : \tau_x}$$

Pair
$$\frac{C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad C \mid \Gamma_1 \vdash_{\text{s}} e_1 : \tau_1 \qquad C \mid \Gamma_2 \vdash_{\text{s}} e_2 : \tau_2}{C \mid \Gamma \vdash_{\text{s}} (e_1, e_2) : \tau_1 \times \tau_2}$$

Region
$$\frac{[x : \tau_x]^n_b \in \Gamma \qquad C \vdash_e \Gamma \rightsquigarrow^x_n \Gamma' \qquad C \mid \Gamma' \vdash_{\text{s}} e : \tau \qquad C \vdash_{\text{e}} (\tau \le \mathbf{L}_{n-1})}{C \mid \Gamma \vdash_{\text{s}} \{\!|e|\!\}^n_{\{x \mapsto b\}} : \tau}$$

Const
$$\frac{C \vdash_{\text{e}} (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} c : \text{CType}(c)}$$

Abs
$$\frac{C \mid \Gamma; (x : \tau_2) \vdash_{\text{s}} e : \tau_1 \qquad C \vdash_{\text{e}} (\Gamma \le k)}{C \mid \Gamma \vdash_{\text{s}} \lambda x.e : \tau_2 \xrightarrow{\text{k}} \tau_1}$$

Borrow
$$\frac{(x \div \sigma)^k_b \in \Gamma \qquad C_x, \tau_x = \text{Inst}(\Gamma, \sigma) \qquad C \vdash_{\text{e}} C_x \wedge (\Gamma \backslash \{x\} \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} \&^b x : \&^b(k, \tau_x)}$$

Reborrow
$$\frac{C \mid \Gamma \vdash_{\text{s}} x : \&^b(k, \tau)}{C \mid \Gamma \vdash_{\text{s}} \&\&^b x : \&^b(k, \tau)}$$

App
$$\frac{C \mid \Gamma_1 \vdash_{\text{s}} e_1 : \tau_2 \xrightarrow{\text{k}} \tau_1 \qquad C \mid \Gamma_2 \vdash_{\text{s}} e_2 : \tau_2' \qquad C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad C \vdash_{\text{e}} (\tau_2' \le \tau_2)}{C \mid \Gamma \vdash_{\text{s}} (e_1\ e_2) : \tau_1}$$

Let
$$\frac{C \wedge D \mid \Gamma_1 \vdash_{\text{s}} e_1 : \tau_1 \qquad (C_\sigma, \sigma) = \text{gen}(D, \Gamma, \tau_1) \qquad C \vdash_{\text{e}} C_\sigma \qquad C \mid \Gamma; (x : \sigma) \vdash_{\text{s}} e_2 : \tau_2 \qquad C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2}{C \mid \Gamma \vdash_{\text{s}} \texttt{let}\ x = e_1\ \texttt{in}\ e_2 : \tau_2}$$

MatchPair
$$\frac{C \mid \Gamma_1 \vdash_{\text{s}} e_1 : \phi(\tau_1 \times \tau_1') \qquad C \mid \Gamma_2; (x : \phi(\tau_1)); (x' : \phi(\tau_1')) \vdash_{\text{s}} e_2 : \tau_2 \qquad C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2}{C \mid \Gamma \vdash_{\text{s}} \texttt{match}_\phi\ x, x' = e_1\ \texttt{in}\ e_2 : \tau_2}$$

Update
$$\frac{C \mid \Gamma \vdash_{\text{s}} \tau : k \qquad C \vdash_{\text{e}} (k \le \mathbf{U}_0) \wedge (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} \texttt{update} : \&^{\text{A}}(k', \text{R}\ \tau) \rightarrow \tau \xrightarrow{\text{A}} \texttt{Unit}}$$

Create
$$\frac{C \mid \Gamma \vdash_{\text{s}} \tau : k \qquad C \vdash_{\text{e}} (k \le \mathbf{U}_0) \wedge (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} \texttt{create} : \tau \rightarrow \text{R}\ \tau}$$

Observe
$$\frac{C \mid \Gamma \vdash_{\text{s}} \tau : k \qquad C \vdash_{\text{e}} (k \le \mathbf{U}_0) \wedge (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} \texttt{observe} : \&^{\text{U}}(k', \text{R}\ \tau) \rightarrow \tau}$$

Destroy
$$\frac{C \mid \Gamma \vdash_{\text{s}} \tau : k \qquad C \vdash_{\text{e}} (k \le \mathbf{U}_0) \wedge (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\text{s}} \texttt{destroy} : \text{R}\ \tau \rightarrow \texttt{Unit}}$$

Fig. 24. Syntax-directed typing rules – $C \mid \Gamma \vdash_{\text{s}} e : \tau$

$$\textsc{Var}\ \frac{(x:\tau)\in\Gamma \qquad C\vdash_{\mathrm{e}}(\Gamma\backslash\{x\}\leq \mathbf{A}_\infty)}{C\mid\Gamma\vdash_{\mathrm{s}} x:\tau}$$

$$\textsc{VarInst}\ \frac{(x:\forall\overline{\kappa_i}\forall(\overline{\alpha_j:k_j}).\, C_x\Rightarrow\tau)\in\Gamma \qquad \psi=[\overline{\kappa_i\mapsto k_i},\overline{\alpha_j\mapsto\tau_j}] \qquad C\vdash_{\mathrm{e}}\psi(C_x)\wedge(\Gamma\backslash\{x\}\leq\mathbf{A}_\infty)}{C\mid\Gamma\vdash_{\mathrm{s}} x[\overline{k_i};\overline{\tau_j}]:\psi\tau}$$

$$\textsc{App}\ \frac{sp: C\vdash_e \Gamma=\Gamma_1\ltimes\Gamma_2 \qquad C\mid\Gamma_1\vdash_{\mathrm{s}} e_1:\tau_2\xrightarrow{\mathrm{k}}\tau_1 \qquad C\mid\Gamma_2\vdash_{\mathrm{s}} e_2:\tau_2' \qquad C\vdash_{\mathrm{e}}(\tau_2'\leq\tau_2)}{C\mid\Gamma\vdash_{\mathrm{s}}(e_1\ e_2)_{sp}:\tau_1}$$

$$\textsc{Pair}\ \frac{sp: C\vdash_e \Gamma=\Gamma_1\ltimes\Gamma_2 \qquad C\mid\Gamma_1\vdash_{\mathrm{s}} e_1:\tau_1 \qquad C\mid\Gamma_2\vdash_{\mathrm{s}} e_2:\tau_2}{C\mid\Gamma\vdash_{\mathrm{s}}(e_1,e_2)^k_{sp}:\tau_1\times\tau_2}$$

$$\textsc{Abs}\ \frac{C\mid\Gamma;(x:\tau_2)\vdash_{\mathrm{s}} e:\tau_1 \qquad C\vdash_{\mathrm{e}}(\Gamma\leq k)}{C\mid\Gamma\vdash_{\mathrm{s}}\lambda^k x.e:\tau_2\xrightarrow{\mathrm{k}}\tau_1}$$

$$\textsc{Let}\ \frac{sp: C\vdash_e \Gamma=\Gamma_1\ltimes\Gamma_2 \qquad C\mid\Gamma_1\vdash_{\mathrm{s}} e_1:\tau_1 \qquad C\mid\Gamma;(x:\tau_1)\vdash_{\mathrm{s}} e_2:\tau_2}{C\mid\Gamma\vdash_{\mathrm{s}}\mathtt{let}\ x=_{sp}\ e_1\ \mathtt{in}\ e_2:\tau_2}$$

$$\textsc{PLet}\ \frac{\begin{array}{c}sp: C\vdash_e \Gamma=\Gamma_1\ltimes\Gamma_2 \qquad \sigma_1=\forall\overline{\kappa_i}(\overline{\alpha_j:k_j}).\, D\Rightarrow\tau_2\xrightarrow{\mathrm{k}}\tau_1 \qquad C\wedge D\mid\Gamma_1;(\overline{\alpha_j:k_j});(x:\tau_2)\vdash_{\mathrm{s}} e_1:\tau_1\\ C\wedge D\vdash_{\mathrm{e}}(\Gamma_1\leq k) \qquad C\vdash_{\mathrm{e}}\exists(\overline{\kappa_i},\overline{\alpha_j}).D \qquad C\mid\Gamma;(f:\sigma_1)\vdash_{\mathrm{s}} e_2:\tau_2\end{array}}{C\mid\Gamma\vdash_{\mathrm{s}}\mathtt{letfun}\,(f:\sigma_1)=_{sp}\lambda^k x.e_1\ \mathtt{in}\ e_2:\tau_2}$$

$$\textsc{MatchPair}\ \frac{sp: C\vdash_e \Gamma=\Gamma_1\ltimes\Gamma_2 \qquad C\mid\Gamma_1\vdash_{\mathrm{s}} e_1:\phi(\tau_1\times\tau_1') \qquad C\mid\Gamma_2;(x:\phi(\tau_1));(y:\phi(\tau_1'))\vdash_{\mathrm{s}} e_2:\tau_2}{C\mid\Gamma\vdash_{\mathrm{s}}\mathtt{match}_\phi\ x,x'=_{sp}\ e_1\ \mathtt{in}\ e_2:\tau_2}$$

Fig. 25. Syntax-directed typing rules for internal language – $C\mid\Gamma\vdash_{\mathrm{s}} e:\tau$

$$\textsc{KArr}\ \frac{}{(\mathrm{True}, \emptyset) \mid \Gamma \vdash_{\mathrm{w}} \tau_1 \xrightarrow{\mathrm{k}} \tau_2 : k} \qquad \textsc{KBorrow}\ \frac{}{(\mathrm{True}, \emptyset) \mid \Gamma \vdash_{\mathrm{w}} \&^{b}(k, \tau) : k} \qquad \textsc{KVar}\ \frac{(\alpha : k) \in \Gamma}{(\mathrm{True}, \emptyset) \mid \Gamma \vdash_{\mathrm{w}} \alpha : k}$$

$$\textsc{KPair}\ \frac{\forall i,\ (C_i, \psi_i) \mid \Gamma \vdash_{\mathrm{w}} \tau_i : k_i \qquad \kappa \text{ fresh} \qquad (C, \psi) = \mathrm{normalize}(\overline{C_i} \wedge \overline{(k_i \leq \kappa)_i}, \overline{\psi_i})}{(C, \psi) \mid \Gamma \vdash_{\mathrm{w}} \tau_1 \times \tau_2 : \psi\kappa}$$

$$\textsc{KApp}\ \frac{(\mathrm{T} : \forall \overline{\kappa_i}.\ C_0 \Rightarrow (\overline{k'_j}) \to k') \in \Gamma \qquad \overline{\kappa'_i} \text{ fresh} \qquad \psi_0 = [\overline{\kappa_i} \to \overline{\kappa'_i}] \qquad \forall j,\ (C_j, \psi_j) \mid \Gamma \vdash_{\mathrm{w}} \tau_j : k'_j \qquad (C, \psi) = \mathrm{normalize}(C_0 \wedge \overline{C_j} \wedge \overline{(k'_j = k_j)_j}, \psi_0 \sqcup \overline{\psi_j})}{(C, \psi|_{\mathrm{fv}(\Gamma)}) \mid \Gamma \vdash_{\mathrm{w}} \mathrm{T}\ \overline{\tau_j} : \psi k}$$

Fig. 26. Kind inference rules – $(C, \psi) \mid \Gamma \vdash_{\mathrm{w}} \tau : k$

E TYPE INFERENCE

In this appendix, we provide the complete type inference rules and show that our type inference algorithm is sound and complete. The constraints rules are already shown in Section 4. Kind inference is presented in Appendix E.1 and the detailed treatment of let-bindings in Appendix E.2. The type inference rules are shown in Fig. 27. The various theorems and their proofs are direct adaptations of the equivalent statements for HM(X) [39].

E.1 Kind Inference

We write $(C, \psi) \mid \Gamma \vdash_{\mathrm{w}} \tau : k$ when type τ has kind k in environment Γ under constraints C and unifier ψ. Γ and τ are the input parameters of our inference procedure. We present the kind inference algorithm as a set of rules in Fig. 26. Higher-kinds are not generally supported and can only appear by looking-up the kind scheme of a type constructor, for use in the type application rule KApp. Type variables must be of a simple kind in rule KVar. Kind schemes are instantiated in the KVar rules by creating fresh kind variables and the associated substitution. KArr and KBorrow simply returns the kind of the primitive arrow and borrow types. The normalize function is used every time several constraints must be composed in order to simplify the constraint and return a most general unifier.

E.2 Generalization and constraints

The Let rule combines several ingredients previously seen: since let expressions are binders, we use Weak on the bound identifier x. Since let-expressions contain two subexpressions, we use the environment splitting relation, $C_s \Leftarrow \Sigma = \Sigma_1 \ltimes (\Sigma_2 \backslash \{x\})$. We remove the x from the right environment, since it is not available outside of the expression e_2, and should not be part of the returned usage environment.

As per tradition in ML languages, generalization is performed on let-bindings. Following HM(X), we write $(C_\sigma, \sigma) = \mathrm{gen}(C, \Gamma, \tau)$ for the pair of a constraint and a scheme resulting from generalization. The definition is provided in Fig. 12. The type scheme σ is created by quantifying over all the appropriate free variables and the current constraints. The generated constraint C_σ uses a new projection operator, $\exists x.D$ where x can be either a type or a kind variable, which allow the creation

$$\textsc{Var}_I \quad \frac{(x : \forall\overline{\kappa_i}\forall(\overline{\alpha_j : k_j}).\, C \Rrightarrow \tau) \in \Gamma \qquad \overline{\kappa'_i}, \overline{\alpha'_j} \text{ fresh} \qquad (C, \psi) = \text{normalize}(C_x, [\overline{\kappa_i \mapsto \kappa'_i}, \overline{\alpha_j \mapsto \alpha'_j}])}{(x : \sigma) \,|(C, \psi|_{\text{fv}(\Gamma)}) \,|\, \Gamma \vdash_{\text{w}} x : \psi\tau}$$

$$\textsc{ReBorrow}_I \quad \frac{\Sigma \,|(C', \psi') \,|\, \Gamma \vdash_{\text{w}} x : \tau' \qquad \kappa \text{ fresh} \qquad (C, \psi) = \text{normalize}(C' \wedge (\tau' \leq \&^b(\kappa, \tau)), \psi')}{(x \div \tau)^{\kappa}_b \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} \&^b x : \&^b(\kappa, \tau)}$$

$$\textsc{Abs}_I \quad \frac{\alpha, \kappa \text{ fresh} \qquad \Sigma_x \,|(C', \psi') \,|\, \Gamma; (x : \alpha) \vdash_{\text{w}} e : \tau \qquad D = C' \wedge (\Sigma_x \backslash \{x\} \leq \kappa) \wedge \text{Weak}_{(x:\alpha)}(\Sigma_x) \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma_x \backslash \{x\} \,|(C, \psi \backslash \{\alpha, \kappa\}) \,|\, \Gamma \vdash_{\text{w}} \lambda x.e : \psi(\alpha) \xrightarrow{\psi(\kappa)} \tau}$$

$$\textsc{Borrow}_I \quad \frac{\_ \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} x : \tau \qquad \kappa \text{ fresh}}{(x \div \tau)^{\kappa}_b \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} \&^b x : \&^b(\kappa, \tau)}$$

$$\textsc{Region}_I \quad \frac{\Sigma' \,|(C', \psi') \,|\, \Gamma \vdash_{\text{w}} e : \tau \qquad C_r \Lleftarrow \Sigma \rightsquigarrow^x_n \Sigma' \qquad (C_\tau, \psi_\tau) \,|\, \Gamma \vdash_{\text{w}} \tau : k_\tau \qquad D = C' \wedge C_\tau \wedge (k_\tau \leq \mathbf{L}_{n-1}) \wedge C_r \qquad (C, \psi) = \text{normalize}(D, \psi' \sqcup \psi_\tau)}{\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} \{\!|e|\!\}^n_{\{x \mapsto b\}} : \tau}$$

$$\textsc{App}_I \quad \frac{\alpha, \kappa \text{ fresh} \qquad \Sigma_1 \,|(C_1, \psi_1) \,|\, \Gamma \vdash_{\text{w}} e_1 : \tau_1 \qquad C_s \Lleftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2 \qquad \Sigma_2 \,|(C_2, \psi_2) \,|\, \Gamma \vdash_{\text{w}} e_2 : \tau_2 \qquad D = C_1 \wedge C_2 \wedge (\tau_1 \leq \tau_2 \xrightarrow{\kappa} \alpha) \wedge C_s \qquad \psi' = \psi_1 \sqcup \psi_2 \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} (e_1\ e_2) : \psi(\alpha)}$$

$$\textsc{MatchPair}_I \quad \frac{\alpha, \kappa, \alpha', \kappa' \text{ fresh} \qquad \Sigma_1 \,|(C_1, \psi_1) \,|\, \Gamma \vdash_{\text{w}} e_1 : \tau_1 \qquad \Gamma' = \Gamma; (x : \phi(\alpha)); (\alpha : \kappa); (x' : \phi(\alpha')); (\alpha' : \kappa') \qquad \Sigma_2 \,|(C_2, \psi_2) \,|\, \Gamma' \vdash_{\text{w}} e_2 : \tau_2 \qquad \psi' = \psi_1 \sqcup \psi_2 \qquad C_s \Lleftarrow \Sigma = \Sigma_1 \ltimes (\Sigma_2 \backslash \{x, x'\}) \qquad D = C'_1 \wedge C_2 \wedge (\tau_1 \leq \phi(\alpha \times \alpha')) \wedge C_s \wedge \text{Weak}_{(x:\phi(\alpha)),(x':\phi(\alpha'))}(\Sigma_2) \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma \,|(C, \psi|_{\text{fv}(\Gamma)}) \,|\, \Gamma \vdash_{\text{w}} \texttt{match}_\phi\ x, x' = e_1\ \texttt{in}\ e_2 : \psi\tau_2}$$

$$\textsc{Pair}_I \quad \frac{\Sigma_1 \,|(C_1, \psi_1) \,|\, \Gamma \vdash_{\text{w}} e_1 : \tau_1 \qquad \Sigma_2 \,|(C_2, \psi_2) \,|\, \Gamma \vdash_{\text{w}} e_2 : \tau_2 \qquad \psi' = \psi_1 \sqcup \psi_2 \qquad C_s \Lleftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2 \qquad D = C_1 \wedge C_2 \wedge C_s \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} (e_1, e_2) : \tau_1 \times \tau_2}$$

$$\textsc{Let}_I \quad \frac{\Sigma_1 \,|(C_1, \psi_1) \,|\, \Gamma \vdash_{\text{w}} e_1 : \tau_1 \qquad (C_\sigma, \sigma) = \text{gen}(C_1, \psi_1\Gamma, \tau_1) \qquad \Sigma_2 \,|(C_2, \psi_2) \,|\, \Gamma; (x : \sigma) \vdash_{\text{w}} e_2 : \tau_2 \qquad C_s \Lleftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2 \backslash \{x\} \qquad \psi' = \psi_1 \sqcup \psi_2 \qquad D = C_\sigma \wedge C_2 \wedge C_s \wedge \text{Weak}_{(x:\sigma)}(\Sigma_2) \qquad (C, \psi) = \text{normalize}(D, \psi')}{\Sigma \,|(C, \psi|_{\text{fv}(\Gamma)}) \,|\, \Gamma \vdash_{\text{w}} \texttt{let}\ x = e_1\ \texttt{in}\ e_2 : \psi\tau_2}$$

$$\text{Weak}_{(x:\sigma)}(\Sigma) = \text{if } (x \in \Sigma) \text{ then True else } (\sigma \leq \mathbf{A}_\infty)$$

Fig. 27. Type inference rules – $\Sigma \,|(C, \psi) \,|\, \Gamma \vdash_{\text{w}} e : \tau$

of local variables inside the constraints. This allows us to encapsulate all the quantified variables in the global constraints. It also reflects the fact that there should exist at least one solution for C for the scheme to be valid. Odersky et al. [26] give a detailed account on the projection operators in HM inference.

E.3 Soundness

Lemma E.1. *Given constraints C and D and substitution ψ, if $D \vdash_e C$ then $\psi D \vdash_e \psi C$.*

Proof. By induction over the entailment judgment. □

Lemma E.2. *Given a typing derivation $C \mid \Gamma \vdash_s e : \tau$ and a constraint $D \in \mathcal{S}$ in solved form such that $D \vdash_e C$, then $D \mid \Gamma \vdash_s e : \tau$*

Proof. By induction over the typing derivation □

Lemma E.3. *Given a type environment $\Gamma, \Gamma' \subset \Gamma$, a term e and a variable $x \in \Gamma$, if $C \mid \Gamma' \vdash_s e : \tau$ then $C \wedge \mathrm{Weak}_x(\Gamma') \mid \Gamma'; (x : \Gamma(x)) \vdash_s e : \tau$*

Proof. Trivial if $x \in \Sigma$. Otherwise, by induction over the typing derivation. □

We define the flattening $\Downarrow\Gamma$ of an environment Γ, as the environment where all the binders are replaced by normal ones. More formally:

$$\Downarrow\Gamma = \left\{(x : \sigma) \mid (x : \sigma) \in \Gamma \vee (x \div \sigma)^k_b \in \Gamma \vee [x : \sigma]^n_b \in \Gamma\right\} \cup \{(\alpha : k) \mid (\alpha : k) \in \Gamma\}$$

Lemma E.4. *Given a type environment Γ and a term e such that $\Sigma \mid (C, \psi) \mid \Gamma \vdash_w e : \tau$ then $\Downarrow\Sigma \subset \Gamma$.*

Proof. By induction over the typing derivation. □

Theorem E.5 (Soundness of inference). *Given a type environment Γ containing only value bindings, $\Gamma|_\tau$ containing only a type binding and a term e, if $\Sigma \mid (C, \psi) \mid \Gamma; \Gamma_\tau \vdash_w e : \tau$ then $C \mid \psi(\Sigma; \Gamma_\tau) \vdash_s e : \tau$, $\psi C = C$ and $\psi\tau = \tau$*

Proof. We proceed by induction over the derivation of $\vdash_w$. Most of the cases follow the proofs from HM(X) closely. For brevity, we showcase three rules: the treatment of binders and weakening, where our inference algorithm differ significantly from the syntax-directed rule, and the *Pair* case which showcase the treatment of environment splitting.

$$\textbf{Case}\ \frac{\begin{array}{c}\textsc{Var}_I \\ (x : \forall\overline{\kappa_i}\forall(\overline{\alpha_j : k_j}).\ C \Rightarrow \tau) \in \Gamma \qquad \overline{\kappa'_i}, \overline{\alpha'_j}\ \text{fresh} \\ (C, \psi) = \mathrm{normalize}(C_x, [\overline{\kappa_i \mapsto \kappa'_i}, \overline{\alpha_j \mapsto \alpha'_j}])\end{array}}{(x : \sigma) \mid (C, \psi|_{\mathrm{fv}(\Gamma)}) \mid \Gamma \vdash_w x : \psi\tau}$$

We have $\Sigma = \{x \mapsto \sigma\}$. Without loss of generality, we can consider $\psi_x = \psi'|_{\mathrm{fv}(\Gamma)} = \psi'|_{\mathrm{fv}(\sigma)}$. Since $\Sigma\backslash\{x\}$ is empty and by definition of normalize, we have that $C \vdash_e \psi'(C_x) \wedge (\psi_x\Sigma\backslash\{x\} \leq \mathbf{A}_\infty)$, $\psi' \leq \psi$ and $\psi' C = C$. By definition, $\psi_x\psi' = \psi'$. By rule *Var*, we obtain $C \mid \psi_x(\Sigma; \Gamma_\tau) \vdash_s x : \psi_x\psi'\tau$, which concludes.

$$\textbf{Case}\ \frac{\begin{array}{c}\textsc{Abs}_I \\ \alpha, \kappa\ \text{fresh} \qquad \Sigma_x \mid (C', \psi') \mid \Gamma; (x : \alpha) \vdash_w e : \tau \\ D = C' \wedge (\Sigma_x\backslash\{x\} \leq \kappa) \wedge \mathrm{Weak}_{(x:\alpha)}(\Sigma_x) \\ (C, \psi) = \mathrm{normalize}(D, \psi')\end{array}}{\Sigma_x\backslash\{x\} \mid (C, \psi\backslash\{\alpha, \kappa\}) \mid \Gamma \vdash_w \lambda x.e : \psi(\alpha) \xrightarrow{\psi(\kappa)} \tau}$$

By induction, we have $C' \mid \Sigma_x; \Gamma_\tau \vdash_s e : \tau$, $\psi C = C$ and $\psi\tau = \tau$.
By definition of normalize and Lemma E.1, we have $C \vdash_e C' \wedge (\psi'\Sigma \leq \psi'\kappa) \wedge \mathrm{Weak}_x(\psi'\Sigma_x)$ and $\psi \leq \psi'$. By Lemma E.1, we have $C \vdash_e (\Sigma \leq \psi\kappa)$.
We now consider two cases:

(1) If $x \in \Sigma_x$, then $\mathrm{Weak}_x(\psi\Sigma_x) = \mathrm{True}$ and by Lemma E.4, $\Sigma_x = \Sigma; (x : \alpha)$. We can deduce $C' \wedge \mathrm{Weak}_{(x:\alpha)}(\psi\Sigma_x) \mid \psi\Sigma; \Gamma_\tau; (x : \psi(\alpha)) \vdash_\mathrm{s} e : \tau$.
(2) If $x \notin \Sigma_x$, then $\Sigma = \Sigma_x$ and $\mathrm{Weak}_{(x:\alpha)}(\psi\Sigma_x) = (\psi\alpha \leq \mathbf{A}_\infty)$. By Lemma E.3, we have $C' \wedge \mathrm{Weak}_{(x:\alpha)}(\psi\Sigma_x) \mid \psi\Sigma; \Gamma_\tau; (x : \psi(\alpha)) \vdash_\mathrm{s} e : \tau$

By Lemma E.2, we have $C \mid \psi(\Sigma; \Gamma_\tau); (x : \psi(\alpha)) \vdash_\mathrm{s} e : \tau$.

By rule *Abs*, we obtain $C \mid \psi(\Sigma; \Gamma_\tau) \vdash_\mathrm{s} \lambda x.e : \psi(\alpha) \xrightarrow{\psi(\kappa)} \tau$ which concludes.

Case $\dfrac{\begin{array}{c}\textsc{Pair}_I \\ \Sigma_1 \mid (C_1, \psi_1) \mid \Gamma \vdash_\mathrm{w} e_1 : \tau_1 \qquad \Sigma_2 \mid (C_2, \psi_2) \mid \Gamma \vdash_\mathrm{w} e_2 : \tau_2 \qquad \psi' = \psi_1 \sqcup \psi_2 \\ C_s \Leftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2 \qquad D = C_1 \wedge C_2 \wedge C_s \qquad (C, \psi) = \mathrm{normalize}(D, \psi')\end{array}}{\Sigma \mid (C, \psi) \mid \Gamma \vdash_\mathrm{w} (e_1, e_2) : \tau_1 \times \tau_2}$

By induction, we have $C_1 \mid \psi_1(\Sigma_1; \Gamma_\tau^1) \vdash_\mathrm{s} e_1 : \tau_1$, $\psi_1 C_1 = C_1$, and $\psi_1\tau_1 = \tau_1$ and $C_2 \mid \psi_2(\Sigma_2; \Gamma_\tau^2) \vdash_\mathrm{s} e_2 : \tau_2$, $\psi_2 C_2 = C_2$ and $\psi_2\tau_2 = \tau_2$. Wlog, we can rename the type Γ_τ^1 and Γ_τ^2 to be disjoint and define $\Gamma_\tau = \Gamma_\tau^1 \cup \Gamma_\tau^2$. By normalization, $C \vdash_\mathrm{e} D$, $\psi \leq \psi'$ and $\psi C = C$. By Lemma E.2 and by substitution, we have $C \mid \psi\Sigma_1 \vdash_\mathrm{s} e_1 : \psi\tau_1$ and $C \mid \psi\Sigma_2 \vdash_\mathrm{s} e_2 : \psi\tau_2$. We directly have that $\psi C_s \Leftarrow \psi\Sigma = \Sigma_1 \ltimes \psi\Sigma_2$ and by Lemma E.2, $\psi C \vdash_\mathrm{e} \psi C_s$. By rule *Pair*, we obtain $C \mid \psi(\Sigma; \Gamma_\tau) \vdash_\mathrm{s} (e_1\ e_2) : \psi(\alpha_1 \times \alpha_2))$, which concludes.

□

E.4 Completeness

We now state our algorithm is complete: for any given syntax-directed typing derivation, our inference algorithm can find a derivation that gives a type at least as general. For this, we need first to provide a few additional definitions.

Definition E.6 (More general unifier). Given a set of variable U and ψ, ψ' and ϕ substitutions. Then $\psi \leq_U^\phi \psi'$ iff $(\phi \circ \psi)|_U = \psi'|_U$.

Definition E.7 (Instance relation). Given a constraints C and two schemes $\sigma = \forall\overline{\alpha}.D \Rightarrow \tau$ and $\sigma' = \forall\overline{\alpha}'.D' \Rightarrow \tau'$. Then $C \vdash_\mathrm{e} \sigma \preceq \sigma'$ iff $C \vdash_\mathrm{e} D[\alpha \to \tau'']$ and $C \wedge D' \vdash_\mathrm{e} (\tau[\alpha \to \tau''] \leq \tau')$

We also extend the instance relation to environments Γ.
We now describe the interactions between splitting and the various other operations.

Lemma E.8. *Given $C \Leftarrow \Gamma = \Gamma_1 \ltimes \Gamma_2$, Then $\Downarrow\Gamma = \Downarrow\Gamma_1 \cup \Downarrow\Gamma_2$.*

Proof. By induction over the splitting derivation. □

Lemma E.9. *Given $C \vdash_e \Gamma_1 = \Gamma_2 \ltimes \Gamma_3$, $C' \Leftarrow \Gamma_1' = \Gamma_2' \ltimes \Gamma_3'$ and ψ such that $\Gamma_i' \subset \Gamma_i''$ and $\vdash_\mathrm{e} \psi\Gamma_i'' \preceq \Gamma_i$ for $i \in \{1; 2; 3\}$.*
Then $C \vdash_\mathrm{e} \psi C'$.

Proof. By induction over the derivation of $C' \Leftarrow \Gamma_1' = \Gamma_2' \ltimes \Gamma_3'$. □

We can arbitrarily extend the initial typing environment in an inference derivation, since it is not used to check linearity.

Lemma E.10. *Given $\Sigma \mid (C, \psi) \mid \Downarrow\Gamma \vdash_\mathrm{w} e : \tau$ and Γ' such that $\Gamma \subseteq \Gamma'$, then $\Sigma \mid (C, \psi) \mid \Downarrow\Gamma' \vdash_\mathrm{w} e : \tau$*

Proof. By induction over the type inference derivation. □

Finally, we present the completeness theorem.

Theorem E.11 (Completeness). *Given* $C' \mid \Gamma' \vdash_s e : \tau'$ *and* $\vdash_e \psi'\Gamma \preceq \Gamma'$. *Then*

$$\Sigma \mid (C, \psi) \mid \Downarrow \Gamma \vdash_w e : \tau$$

for some environment Σ, *substitution* ψ, *constraint* C *and type* τ *such that*

$$\psi \leq^{\phi}_{\mathrm{fv}(\Gamma)} \psi' \qquad C' \vdash_e \phi C \qquad \vdash_e \phi\sigma \preceq \sigma' \qquad \Sigma \subset \Gamma$$

where $\sigma' = gen(C', \Gamma', \tau')$ *and* $\sigma = gen(C, \Gamma, \tau)$

Proof. Most of the difficulty of this proof comes from proper handling of instanciation and generalization for type-schemes. This part is already proven by Sulzmann [39] in the context of HM(X). As before, we will only present few cases which highlights the handling of bindings and environments. For clarity, we will only present the part of the proof that only directly relate to the new aspect introduced by Affe.

Case $\dfrac{\textsc{Abs}\quad C' \mid \Gamma'_x; (x : \tau'_2) \vdash_s e : \tau'_1 \qquad C \vdash_e (\Gamma'_x \leq k)}{C' \mid \Gamma'_x \vdash_s \lambda^k x.e : \tau'_2 \xrightarrow{k} \tau'_1}$ and $\vdash_e \psi'\Gamma \preceq \Gamma'$.

Let us pick α and κ fresh. Wlog, we choose $\psi'(\alpha) = \tau_2$ and $\psi'(\kappa) = k$ so that $\vdash_e \psi'\Gamma_x \preceq \Gamma'_x$. By induction:

$$\Sigma_x \mid (C, \psi) \mid \Downarrow \Gamma_x; (x : \alpha) \vdash_w e : \tau \qquad \psi \leq^{\phi}_{\mathrm{fv}(\Gamma_x) \cup \{\alpha;\kappa\}} \psi'$$

$$C' \vdash_e \phi C \qquad \vdash_e \phi\sigma \preceq \sigma' \qquad \Sigma_x \subset \Gamma_x; (x : \alpha)$$

$$\sigma' = \mathrm{gen}(C', \Gamma'_x; (x : \tau'_2), \tau'_1) \qquad \sigma = \mathrm{gen}(C, \Gamma_x; (x : \alpha), \tau_1)$$

Let $C_a = C \wedge (\Sigma \leq \kappa) \wedge \mathrm{Weak}_{(x:\alpha)}(\Sigma_x)$ and By definition, $\psi_D \backslash \{\alpha; \kappa\} \leq^{\phi_D}_{\mathrm{fv}(\Gamma_x)} \psi$ which means we have $\psi_D \backslash \{\alpha; \kappa\} \leq^{\phi \circ \phi_D}_{\mathrm{fv}(\Gamma_x)} \psi'$. We also have that $\Sigma_x \backslash \{x\} \subset \Gamma_x$.

Since $C \vdash_e (\Gamma'_x \leq k)$, we have $C \vdash_e \psi'(\Sigma \leq \kappa)$. If $x \in \Sigma_x$, then $\mathrm{Weak}_{(x:\alpha)}(\Sigma_x) = \mathrm{True}$. Otherwise we can show by induction that $C' \vdash_e \psi' \mathrm{Weak}_{(x:\alpha)}(\Sigma_x)$. We also have $\psi C = C$, which gives us $C' \vdash_e \psi'(C_a)$. We can deduce $C' \vdash_e \psi'(C_a)$.
This means (C', ψ') is a normal form of C_a, so a principal normal form exists. Let $(D, \psi_D) = \mathrm{normalize}(C_a, \psi\{\alpha; \kappa\})$. By the property of principal normal forms, we have $C' \vdash_e \rho D$ and $\psi_D \leq^{\rho}_{\mathrm{fv}(\Gamma_x)} \psi'$.

By application of $\textsc{Abs}_I$, we have $\Sigma_x \backslash \{x\} \mid (C, \psi_D \backslash \{\alpha, \kappa\}) \mid \Downarrow \Gamma_x \vdash_w \lambda x.e : \psi_D(\alpha) \xrightarrow{\psi_D(\kappa)} \tau_1$.
The rest of the proof proceeds as in the original HM(X) proof.

Case $\dfrac{\textsc{Pair}\quad C' \mid \Gamma'_1 \vdash_s e_1 : \tau'_1 \qquad C' \mid \Gamma'_2 \vdash_s e_2 : \tau'_1 \qquad C \vdash_e \Gamma' = \Gamma'_1 \ltimes \Gamma'_2}{C' \mid \Gamma' \vdash_s (e_1\ e_2) : \tau'_1 \times \tau'_2}$
and $\vdash_e \psi'\Gamma \preceq \Gamma'$

The only new elements compared to HM(X) is the environment splitting. By induction:

$$\Sigma_1 \mid (C_1, \psi_1) \mid \Downarrow \Gamma_1 \vdash_w e : \tau_1 \qquad \psi_1 \leq^{\phi_2}_{\mathrm{fv}(\Gamma)} \psi'_1 \qquad C' \vdash_e \phi C_1 \qquad \vdash_e \phi_2\sigma_1 \preceq \sigma'_1$$

$$\Sigma_1 \subset \Gamma_1 \qquad \sigma'_1 = \mathrm{gen}(C', \Gamma'_1, \tau'_1) \qquad \sigma_1 = \mathrm{gen}(C, \Gamma_1, \tau_1)$$

and

$$\Sigma_2 \,|(C_2, \psi_2)\,| \Downarrow \Gamma_2 \vdash_{\mathrm{w}} e : \tau_2 \qquad \psi_2 \leq^{\phi_2}_{\mathrm{fv}(\Gamma)} \psi_2' \qquad C' \vdash_{\mathrm{e}} \phi C_2$$

$$\vdash_{\mathrm{e}} \phi_2 \sigma_2 \preceq \sigma_2' \qquad \Sigma_2 \subset \Gamma_2$$

$$\sigma_2' = \mathrm{gen}(C', \Gamma_2', \tau_2') \qquad \sigma_2 = \mathrm{gen}(C, \Gamma_2, \tau_2)$$

By Lemmas E.8 and E.10, we have

$$\Sigma_1 \,|(C_1, \psi_1)\,| \Downarrow \Gamma \vdash_{\mathrm{w}} e : \tau_1 \qquad \Sigma_2 \,|(C_2, \psi_2)\,| \Downarrow \Gamma \vdash_{\mathrm{w}} e : \tau_2$$

Let $C_s \Leftarrow \Sigma = \Sigma_1 \ltimes \Sigma_2$. We know that $\vdash_{\mathrm{e}} \psi' \Gamma \preceq \Gamma'$, $\vdash_{\mathrm{e}} \psi_i' \Gamma_i \preceq \Gamma_i'$ and $\Sigma_i \subset \Gamma_i$. By Lemma E.9, we have $C \vdash_{\mathrm{e}} \psi' C_s$. The rest of the proof follows HM(X).

□

Corollary E.12 (Principality). *Let* $\mathrm{True} \mid \Gamma \vdash_{\mathrm{s}} e : \sigma$ *a closed typing judgment. Then* $\Sigma \,|(C, \psi)\,| \Downarrow \Gamma \vdash_{\mathrm{w}} e : \tau$ *such that:*

$$(\mathrm{True}, \sigma_o) = \mathit{gen}(C, \psi\Gamma, \tau) \qquad \vdash_{\mathrm{e}} \sigma_o \preceq \sigma$$

F SEMANTICS DEFINITIONS

Fig. 28 presents the full big-step interpretation. Fig. 29 contains the cases for resources.

G PROOFS FOR METATHEORY

- For simplicity, we only consider terms in A-normal forms following the grammar:

$$e ::= \ldots \mid (x\ x') \mid (x, x')^k \mid \mathtt{match}_\phi\ x, y = z\ \mathtt{in}\ e$$

  Typing and semantics rules are unchanged.
- Borrow qualifiers $\beta ::= \mathbf{U}_n \mid \mathbf{A}_n$ where $n \geq 0$ is a region level. A vector of borrow qualifiers $\overline{\beta}$ is wellformed if all **U**s come before all **A**s in the vector.
- Borrow compatibility $\overline{\beta} \nwarrow \beta$,

$$\beta_n \overline{\beta} \nwarrow \beta_n$$

- Store typing $\vdash \delta : \Delta$,

$$\frac{(\forall \ell \in \mathrm{dom}(\delta))\ \ \Delta \vdash \delta(\ell) : \Delta(\ell)}{\vdash \delta : \Delta}$$

- Relating storables to type schemes $\Delta \vdash w : \sigma$
  We write dis(Γ) for $\gamma^{\mathbf{L}}$ and $\gamma^{\mathbf{A}}$ and $\gamma^{\mathbf{U}}$ and $\gamma_{\#}$ are all disjoint.

$$\frac{(\exists \Gamma)\ \Delta \vdash \gamma : \Gamma \qquad \mathrm{dis}(\Gamma) \qquad C \mid \Gamma; (x : \tau_2) \vdash_{\mathrm{s}} e : \tau_1 \qquad \overline{\alpha} = \mathrm{fv}(\tau_1, \tau_2) \setminus \mathrm{fv}(\Gamma)}{\Delta \vdash (\gamma, \lambda[\overline{\kappa} \mid C \Rightarrow k]x.e) : \forall \overline{\kappa} \forall \overline{(\alpha : k)}.(C \Rightarrow \tau_2 \xrightarrow{\mathrm{k}} \tau_1)}$$

- Relating storables to types $\Delta \vdash w : \tau$

$$\frac{(\exists \Gamma, C)\ \Delta \vdash \gamma : \Gamma \qquad \mathrm{dis}(\Gamma) \qquad C \mid \Gamma; (x : \tau_2) \vdash_{\mathrm{s}} e : \tau_1 \qquad C \vdash_{\mathrm{e}} (\Gamma \leq k)}{\Delta \vdash (\gamma, \lambda^k x.e) : \tau_2 \xrightarrow{\mathrm{k}} \tau_1}$$

$$\frac{\Delta \vdash r_1 : \tau_1 \qquad \Delta \vdash r_2 : \tau_2 \qquad \cdot \vdash_{\mathrm{e}} (\tau_1 \leq k) \wedge (\tau_1 \leq k)}{\Delta \vdash (r_1, r_2)^k : \tau_1 \times^k \tau_2} \qquad \frac{\Delta \vdash r : \mathrm{IType}(\mathrm{T}, \overline{\tau})}{\Delta \vdash [r] : \mathrm{T}\ \overline{\tau}} \qquad \Delta \vdash \bullet : \tau$$

- Relating results to type schemes $\Delta \vdash r : \sigma$

$$\Delta \vdash c : \mathrm{CType}(c) \qquad \Delta \vdash \ell : \Delta(\ell) \qquad \frac{\overline{\beta} \nwarrow b_n \qquad \Delta \vdash \ell : \tau}{\Delta \vdash \overline{\beta}\ell : \&^b(b_n, \tau)}$$

- We write $\mathrm{aff}_\Delta(\ell)$ to express that ℓ points to a resource that requires at least affine treatment. Borrow types do not appear in store types as the store only knows about the actual resources. Define $\mathrm{aff}_\Delta(\ell)$ if one of the following cases holds:
  - $\Delta(\ell) = \forall \overline{\kappa} \forall \overline{(\alpha : k)}.(C \Rightarrow \tau_2 \xrightarrow{\mathrm{k}} \tau_1)$ and $C \wedge (k \leq \mathbf{U}_\infty)$ is contradictory;
  - $\Delta(\ell) = \tau_2 \xrightarrow{\mathrm{k}} \tau_1$ and $(\mathbf{A} \leq k)$;
  - $\Delta(\ell) = \tau_1 \times^k \tau_2$ and $(\mathbf{A} \leq k)$;
  - $\Delta(\ell) = \mathrm{T}\ \overline{\tau}$.
- We write $\mathrm{lin}_\Delta(\ell)$ to express that ℓ points to a linear resource.
  Define $\mathrm{lin}_\Delta(\ell)$ if one of the following cases holds:
  - $\Delta(\ell) = \forall \overline{\kappa} \forall \overline{(\alpha : k)}.(C \Rightarrow \tau_2 \xrightarrow{\mathrm{k}} \tau_1)$ and $C \wedge (k \leq \mathbf{A}_\infty)$ is contradictory;
  - $\Delta(\ell) = \tau_2 \xrightarrow{\mathrm{k}} \tau_1$ and $(\mathbf{L} \leq k)$;
  - $\Delta(\ell) = \tau_1 \times^k \tau_2$ and $(\mathbf{L} \leq k)$;

```
let rec eval
  (δ:store) (π:perm) (γ:venv) i e
  : (store × perm × result) sem =
  if i=0 then TimeOut else
  let i' = i − 1 in
  match e with
  | Const (c) →
    Ok (δ, π, c)

  | Var (x) →
    let* r = γ(x) in
    Ok (δ, π, r)

  | Varinst (x, k̄) →
    let* rx = γ(x) in
    let* ℓ = getloc rx in
    let*? () = ℓ ∈ π in
    let* w = δ(ℓ) in
    let* (γ', κ̄', C', k', x', e') = getstpoly w in
    let π' =
      if C'{k̄\>κ̄'} =e [(k' ≤ U)]{k̄\>κ̄'}
      then π else π − ℓ
    in
    let w = STCLOS (γ', k'{k̄\>κ̄'}, x', e'{k̄\>κ̄'}) in
    let (ℓ', δ') = salloc δ w in
    Ok (δ', π' + ℓ', ℓ')

  | Lam (k, x, e) →
    let w = STCLOS (γ, k, x, e) in
    let (ℓ', δ') = salloc δ w in
    let π' = π + ℓ' in
    Ok (δ', π', ℓ')

  | App (e1, e2, sp) →
    let (γ1, γ2) = vsplit γ sp in
    let* (δ1, π1, r1) = eval δ π γ1 i' e1 in
    let* ℓ1 = getloc r1 in
    let*? () = ℓ1 ∈ π1 in
    let* w = δ1(ℓ1) in
    let* (γ', k', x', e') = getstclos w in
    let π1' = if k' ≤ U then π1 else π1 − ℓ1 in
    let* δ1' = δ1(ℓ1) ← (if k' ≤ U then w else •) in
    let* (δ2, π2, r2) = eval δ1' π1' γ2 i' e2 in
    let* (δ3, π3, r3) = eval δ2 π2 γ'(x'↦ r2) i' e' in
    Ok (δ3, π3, r3)

  | Let (x, e1, e2, sp) →
    let (γ1, γ2) = vsplit γ sp in
    let* (δ1, π1, r1) = eval δ π γ1 i' e1 in
    let* (δ2, π2, r2) = eval δ1 π1 γ2(x ↦ r1) i' e2 in
    Ok (δ2, π2, r2)

  | LetFun (f, σ, k, x, e, e', sp) →
    let (γ1, γ2) = vsplit γ sp in
    let ∀(κ̄, _, C, τ) = σ in
    let w = STPOLY (γ1, κ̄, C, k, x, e') in
    let (ℓ', δ') = salloc δ w in
    let π' = π + ℓ' in
    let* (δ1, π1, r1) = eval δ' π' γ2(f ↦ ℓ') i' e' in
    Ok (δ1, π1, r1)

  | Pair (k, e1, e2, sp) →
    let (γ1, γ2) = vsplit γ sp in
    let* (δ1, π1, r1) = eval δ π γ1 i' e1 in
    let* (δ2, π2, r2) = eval δ1 π1 γ2 i' e2 in
    let w = STPAIR (k, r1, r2) in
    let (ℓ', δ') = salloc δ2 w in
    Ok (δ', π2 + ℓ', ℓ')

  | Match (x, x', e1, e2, sp) →
    let (γ1, γ2) = vsplit γ sp in
    let* (δ1, π1, r1) = eval δ π γ1 i' e1 in
    let* ℓ = getloc r1 in
    let* w = δ1(ℓ) in
    let* (k', r1', r2') = getstpair w in
    let π1' = (if k' ≤ U then π1 else π1 − ℓ) in
    let γ2' = γ2(x ↦ r1)(x' ↦ r1') in
    let* (δ2, π2, r2) = eval δ1 π1' γ2' i' e2 in
    Ok (δ2, π2, r2)

  | Matchborrow (x, x', e1, e2, sp) →
    let (γ1, γ2) = vsplit γ sp in
    let* (δ1, π1, r1) = eval δ π γ1 i' e1 in
    let* ρ = getaddress r1 in
    let* (b, _, ℓ) = getborrowed_loc r1 in
    let* w = δ1(ℓ) in
    let* (k', r1', r2') = getstpair w in
    let* ρ1 = getaddress r1' in
    let* ρ2 = getaddress r2' in
    let* ρ1' = b.ρ1 in
    let* ρ2' = b.ρ2 in
    let π1'' = (if k' ≤ U then π1 else π1 − ρ) in
    let γ2'' = γ2(x ↦ ρ1')(x' ↦ ρ2') in
    let* (δ2, π2, r2) = eval δ1 π1'' γ2'' i' e2 in
    Ok (δ2, π2, r2)
```

Fig. 28. Big-step interpretation

| Region (e, n, x, τ_x, b) $\longrightarrow$
  let+ $\rho = \gamma$(x) **in**
  let* ρ' = b.ρ **in**
  let* π' = reach ρ τ_x δ **in**
  let* π'' = b.π' **in**
  let γ' = γ(x $\mapsto \rho'$) **in**
  let π = ($\pi \cup \pi''$) $\setminus$ π' **in**
  let* (δ_1, π_1, r_1) = eval δ π γ' i' e **in**
  let π_1 = ($\pi_1 \setminus \pi''$) $\cup$ π' **in**
  Ok (δ_1, π_1, r_1)

| Borrow (b, x) $\longrightarrow$
  let+ $\rho = \gamma$(x) **in**
  let*? () = ρ ? b **&&** $\rho \in \pi$ **in**
  Ok (δ, π, ρ)

| Destroy (e_1) $\longrightarrow$
  let* (δ_1, π_1, r_1) = eval δ π γ i' e_1 **in**
  let* ρ = getaddress r_1 **in**
  let* ℓ = getloc r_1 **in**
  let* w = $\delta_1(\ell)$ **in**
  let* r = getstrsrc w **in**
  let*? () = $\rho \in \pi_1$ **in**
  let* δ_1' = $\delta_1(\ell) \leftarrow \bullet$ **in**
  let π_1' = $\pi_1 - \ell$ **in**
  Ok (δ_1', π_1', ())

| Create $\longrightarrow$
  let w = STRSRC (0) **in**
  let (ℓ_1, δ_1) = salloc δ w **in**
  let π_1 = $\pi + \ell_1$ **in**
  Ok (δ_1, π_1, ℓ_1)

| Observe (e_1) $\longrightarrow$
  let* (δ_1, π_1, r_1) = eval δ π γ i' e_1 **in**
  let* ρ = getaddress r_1 **in**
  let*? () = $\rho \in \pi_1$ **in**
  let* (b, _, ℓ) = getborrowed_loc r_1 **in**
  let*? () = (b = **U**) **in**
  let* w = $\delta_1(\ell)$ **in**
  let* r = getstrsrc w **in**
  Ok (δ_1, π_1, r)

| Update (e_1, e_2, sp) $\longrightarrow$
  let (γ_1, γ_2) = vsplit γ sp **in**
  let* (δ_1, π_1, r_1) = eval δ π γ_1 i' e_1 **in**
  let* ρ = getaddress r_1 **in**
  let* (b, _, ℓ) = getborrowed_loc r_1 **in**
  let*? () = (b = **A**) **in**
  let* (δ_2, π_2, r_2) = eval δ_1 π_1 γ_2 i' e_2 **in**
  let* w = $\delta_2(\ell)$ **in**
  let* r = getstrsrc w **in**
  let*? () = $\rho \in \pi_2$ **in**
  let* δ_2' = $\delta_2(\ell) \leftarrow$ STRSRC (r_2) **in**
  let π_2' = $\pi_2 - \rho$ **in**
  Ok (δ_2', π_2', ())

Fig. 29. Big-step interpretation (resources)

  - $\Delta(\ell) = \mathrm{T}\,\overline{\tau}$.
- It remains to characterize unrestricted resources. Define $\mathrm{unr}_\Delta(\ell)$ if neither $\mathrm{aff}_\Delta(\ell)$ nor $\mathrm{lin}_\Delta(\ell)$ holds.
- Relating environments to contexts
  $\Delta \vdash \gamma^{\mathbf{L}}, \gamma^{\mathbf{A}}, \gamma^{\mathbf{U}}, \gamma^{\mathbf{A}}_{\#}, \gamma^{\mathbf{U}}_{\#} : \Gamma$.
  Here we consider an environment $\gamma = (\gamma^{\mathbf{L}}, \gamma^{\mathbf{A}}, \gamma^{\mathbf{U}}, \gamma_{\#})$ as a tuple consisting of the active entries in $\gamma^{\mathbf{L}}$ and the entries for exclusive borrows in $\gamma^{\mathbf{A}}$ and for shared borrows in $\gamma^{\mathbf{U}}$, and suspended entries in $\gamma_{\#} = \gamma^{\mathbf{A}}_{\#}, \gamma^{\mathbf{U}}_{\#}$ for affine and unrestricted entries. The suspended entries cannot be used directly, but they can be activated by appropriate borrowing on entry to a region.

$$\Delta \vdash \cdot, \cdot, \cdot, \cdot : \cdot$$

$$\frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma \qquad \Delta \vdash r : \sigma \qquad \mathrm{lin}_{\Delta}(r)}{\Delta \vdash \gamma^{\mathrm{L}}[x \mapsto r], \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma; (x : \sigma)} \qquad \frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma \qquad \Delta \vdash r : \sigma \qquad \mathrm{aff}_{\Delta}(r)}{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}[x \mapsto r], \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma; (x : \sigma)}$$

$$\frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma \qquad \Delta \vdash r : \sigma \qquad \mathrm{unr}_{\Delta}(r)}{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}[x \mapsto r], \gamma_{\#} : \Gamma; (x : \sigma)} \qquad \frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#}^{\mathrm{A}}, \gamma_{\#}^{\mathrm{U}} : \Gamma \qquad \Delta \vdash r : \sigma}{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#}^{\mathrm{A}}, \gamma_{\#}^{\mathrm{U}}[x \mapsto r] : \Gamma; [x : \sigma]_{\mathrm{U}}^{n}}$$

$$\frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#}^{\mathrm{A}}, \gamma_{\#}^{\mathrm{U}} : \Gamma \qquad \Delta \vdash r : \sigma}{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#}^{\mathrm{A}}[x \mapsto r], \gamma_{\#}^{\mathrm{U}} : \Gamma; [x : \sigma]_{\mathrm{A}}^{n}} \qquad \frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma \qquad \Delta \vdash \mathbf{U}\rho : \sigma}{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}[x \mapsto \mathbf{U}\rho], \gamma_{\#} : \Gamma; (x \div \sigma)_{\mathrm{U}}^{k}}$$

$$\frac{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}, \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma \qquad \Delta \vdash \mathbf{A}\rho : \sigma}{\Delta \vdash \gamma^{\mathrm{L}}, \gamma^{\mathrm{A}}[x \mapsto \mathbf{A}\rho], \gamma^{\mathrm{U}}, \gamma_{\#} : \Gamma; (x \div \sigma)_{\mathrm{A}}^{k}}$$

Extending environments and stores.

$$\Delta \leq \Delta \qquad \frac{\Delta \leq \Delta' \qquad \ell \notin \mathrm{dom}(\delta)}{\Delta \leq \Delta'(\ell : \sigma)}$$

$$\delta \leq \delta \qquad \frac{\delta \leq \delta' \qquad \ell \notin \mathrm{dom}(\delta)}{\delta \leq \delta'[\ell \mapsto w]}$$

Lemma G.1 (Store Weakening). *$\Delta \vdash \gamma : \Gamma$ and $\Delta \leq \Delta'$ implies $\Delta' \vdash \gamma : \Gamma$.*

Lemma G.2 (Store Extension).
- *If $\Delta_1 \leq \Delta_2$ and $\Delta_2 \leq \Delta_3$, then $\Delta_1 \leq \Delta_3$.*
- *If $\delta_1 \leq \delta_2$ and $\delta_2 \leq \delta_3$, then $\delta_1 \leq \delta_3$.*

We write getloc$(\cdot)$ for the function that extracts a multiset of *raw locations* from a result or from the range of the variable environment.

$$\begin{aligned} \mathrm{getloc}(\overline{\beta}\ell) &= \{\ell\} \\ \mathrm{getloc}(c) &= \{\} \\ \mathrm{getloc}(\cdot) &= \{\} \\ \mathrm{getloc}(\gamma(x \mapsto r)) &= \mathrm{getloc}(\gamma) \cup \mathrm{getloc}(r) \end{aligned}$$

We write $\mathrm{reach}_{\delta}(\gamma)$ for the multiset of all *addresses* reachable from getloc(γ) assuming that getloc$(\gamma) \subseteq \mathrm{dom}(\delta)$[5]. The function $\mathrm{reach}_{\delta}(\cdot)$ is defined in two steps. First a helper function for results, storables, and environments.

[5]In mixed comparisons between a multiset and a set, we tacitly convert a multiset M to its supporting set $\{x \mid M(x) \neq 0\}$.

$$\begin{aligned}
\mathrm{reach}_0(\cdot) &= \cdot \\
\mathrm{reach}_0(\gamma(x \mapsto r)) &= \mathrm{reach}_0(\gamma) \cup \mathrm{reach}_0(r) \\
\mathrm{reach}_0(\rho) &= \{\rho\} \\
\mathrm{reach}_0(c) &= \{\} \\
\mathrm{reach}_0(\mathrm{STPOLY}(\gamma, \overline{\kappa}, C, k, x, e)) &= \mathrm{reach}_0(\gamma) \\
\mathrm{reach}_0(\mathrm{STCLOS}(\gamma, k, x, e)) &= \mathrm{reach}_0(\gamma) \\
\mathrm{reach}_0(\mathrm{STPAIR}(k, r_1, r_2)) &= \mathrm{reach}_0(r_1) \cup \mathrm{reach}_0(r_2) \\
\mathrm{reach}_0(\mathrm{STRSRC}(r)) &= \mathrm{reach}_0(r) \\
\mathrm{reach}_0(\bullet) &= \{\}
\end{aligned}$$

This multiset is closed transitively by store lookup. We define $\mathrm{reach}_\delta(\gamma)$ as the smallest multiset Θ that fulfills the following inequations. We assume a nonstandard model of multisets such that an element ℓ may occur infinitely often as in $\Theta(\ell) = \infty$.

$$\begin{aligned}
&\Theta \supseteq \mathrm{reach}_0(\gamma) \\
&\Theta \supseteq \mathrm{reach}_0(w) && \text{if } \overline{\beta}\ell \in \Theta \wedge w = \delta(\ell)
\end{aligned}$$

Definition G.3 (Wellformed permission). A permission π is *wellformed* if it contains at most one address for each raw location.

Definition G.4 (Permission closure). The closure of a permission $\downarrow\pi$ is the set of addresses reachable from π by stripping an arbitrary number of borrows from it. It is the homomorphic extension of the closure $\downarrow\rho$ for a single address.

$$\downarrow\ell = \{\ell\} \qquad\qquad \downarrow(\beta\rho) = \{\beta\rho\} \cup \downarrow\rho$$

Lemma G.5 (Containment). *Suppose that* $\vdash \delta : \Delta$, $\Delta \vdash r : \tau$, $C \vdash_{\mathrm{e}} (\tau \le k) \wedge (k \le \mathbf{L}_{m-1})$. *Then* $\mathrm{reach}_\delta(r)$ *cannot contain addresses* ρ *such that* $\rho = b_n\rho'$ *with* $n \ge m$.

Proof. By inversion of result typing there are three cases.

Case $\Delta \vdash c : \mathrm{CType}(c)$. Immediate: reachable set is empty.

Case $\dfrac{\overline{\beta} \nwarrow b_n \qquad \Delta \vdash \ell : \tau}{\Delta \vdash \overline{\beta}\ell : \&^b(b_n, \tau)}$. The typing constraint enforces that $n < m$.

Case $\Delta \vdash \ell : \Delta(\ell)$. We need to continue by dereferencing ℓ and inverting store typing.

Case $\Delta \vdash \bullet : \tau$. Trivial.

Case $\dfrac{\Delta \vdash r : \mathrm{IType}(\mathrm{T}, \overline{\tau})}{\Delta \vdash [r] : \mathrm{T}\,\overline{\tau}}$. We assume the implementation type of a result to be unrestricted.

Case $\dfrac{\Delta \vdash r_1 : \tau_1 \qquad \Delta \vdash r_2 : \tau_2 \qquad \cdot \vdash_{\mathrm{e}} (\tau_1 \le k) \wedge (\tau_1 \le k)}{\Delta \vdash (r_1, r_2)^k : \tau_1 \times^k \tau_2}$.

The typing constraint yields that $k \le \mathbf{L}_{m-1}$. By induction and transitivity of $\le$, we find that $\mathrm{reach}_\delta(r_1)$ and $\mathrm{reach}_\delta(r_2)$ cannot contain offending addresses.

Case $\dfrac{(\exists \Gamma, C)\, \Delta \vdash \gamma : \Gamma \qquad \mathrm{dis}(\Gamma) \qquad C \mid \Gamma; (x : \tau_2) \vdash_{\mathrm{s}} e : \tau_1 \qquad C \vdash_{\mathrm{e}} (\Gamma \le k)}{\Delta \vdash (\gamma, \lambda^k x.e) : \tau_2 \xrightarrow{\mathrm{k}} \tau_1}$.

The typing constraint yields that $k \le \mathbf{L}_{m-1}$. By transitivity of $\le$ and $\Delta \vdash \gamma : \Gamma$, we find that the types of all addresses in γ have types bounded by $\mathbf{L}_{m-1}$ and, by induction, they cannot contain offending addresses. □

Theorem G.6 (Type Soundness). *Suppose that*

(A1) $C \mid \Gamma \vdash_s e : \tau$
(A2) $\Delta \vdash \gamma : \Gamma$
(A3) $\vdash \delta : \Delta$
(A4) π *is wellformed and* $\mathrm{getloc}(\pi) \subseteq \mathrm{dom}(\delta) \setminus \delta^{-1}(\bullet)$
(A5) $\mathrm{reach}_0(\gamma) \subseteq \pi$, $\mathrm{reach}_\delta(\gamma) \subseteq \downarrow\pi$.
(A6) $\mathrm{getloc}(\gamma^{\mathbf{L}})$, $\mathrm{getloc}(\gamma^{\mathbf{A}})$, $\mathrm{getloc}(\gamma^{\mathbf{U}})$, *and* $\mathrm{getloc}(\gamma_\#)$ *are all disjoint*
(A7) Incoming Resources:
  (a) $\forall \ell \in \mathrm{getloc}(\mathrm{reach}_\delta(\gamma)), \delta(\ell) \neq \bullet$.
  (b) $\forall \ell \in \Theta = \mathrm{getloc}(\mathrm{reach}_\delta(\gamma^{\mathbf{L}}, \gamma^{\mathbf{A}}, \gamma^{\mathbf{A}}_\#)), \Theta(\ell) = 1$.

For all $i \in \mathbb{N}$, *if* R' = eval $\delta\ \pi\ \gamma$ i e *and* $R' \neq \mathrm{TimeOut}$, *then* $\exists\ \delta', \pi', r', \Delta'$ *such that*

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$
(R2) $\Delta \leq \Delta'$, $\delta \leq \delta'$, $\vdash \delta' : \Delta'$
(R3) $\Delta' \vdash r' : \tau$
(R4) π' *is wellformed and* $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
(R6) Frame:
  For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ *it must be that*
  - $\delta'(\ell) = \delta(\ell)$ *and*
  - *for any* ρ *with* $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

(R7) Unrestricted values, resources, and borrows:
  For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{U}}, \gamma^{\mathbf{U}}_\#)$ *with* $\mathrm{getloc}(\rho) = \{\ell\}$, *it must be that* $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ *and* $\rho \in \pi'$.
(R8) Affine borrows and resources:
  For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}, \gamma^{\mathbf{A}}_\#)$ *with* $\mathrm{getloc}(\rho) = \{\ell\}$, *it must be that* $\ell \in \mathrm{dom}(\delta)$. *If* $\rho \neq \ell$, *then* $\delta'(\ell) \neq \bullet$. *If* $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}_\#)$, *then* $\rho \in \pi'$.
(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. *Let* $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathbf{L}})$.
  For all $\ell \in \Theta$ *it must be that* $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, *and if* $\delta(\ell)$ *is a resource, then* $\delta'(\ell) = \bullet$.
(R10) No thin air permission:
  $\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.

Some explanations are in order for the resource-related assumptions and statements.

Incoming resources are always active (i.e., not freed). Linear and affine resources as well as suspended affine borrows have exactly one pointer in the environment.

The Frame condition states that only store locations reachable from the current environment can change and that all permissions outside the reachable locations remain the same.

Unrestricted values, resources, and borrows do not change their underlying resource and do not spend their permission.

Affine borrows and resources may or may not spend their permission. Borrows are not freed, but resources may be freed. The permissions for suspended entries remain intact.

A linear resource is always freed.

Outgoing permissions are either inherited from the caller or they refer to newly created values.

Proof. By induction on the evaluation of eval $\delta\ \pi\ \gamma$ i e.

The base case is trivial as eval $\delta\ \pi\ \gamma\ 0$ e = TimeOut.

For $i > 0$ consider the different cases for expressions. For lack of spacetime, we only give details on some important cases.

Case e **of**
| Let (x, e_1, e_2, sp) →
  let (γ_1, γ_2) = vsplit γ sp **in**
  let* (δ_1, π_1, r_1) = eval $\delta\ \pi\ \gamma_1$ i' e_1 **in**
  let* (δ_2, π_2, r_2) = eval $\delta_1\ \pi_1\ \gamma_2$(x ↦ r_1) i' e_2 **in**
  Ok (δ_2, π_2, r_2)

We need to invert rule Let for monomorphic let:

$$\textsc{Let}\ \frac{sp : C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad C \mid \Gamma_1 \vdash_{\mathrm{s}} e_1 : \tau_1 \qquad C \mid \Gamma; (x : \tau_1) \vdash_{\mathrm{s}} e_2 : \tau_2}{C \mid \Gamma \vdash_{\mathrm{s}} \mathtt{let}\ x =_{sp}\ e_1\ \mathtt{in}\ e_2 : \tau_2}$$

As sp is the evidence for the splitting judgment and vsplit distributes values according to sp, we obtain

$$\Delta \vdash \gamma_1 : \Gamma_1 \tag{1}$$
$$\Delta \vdash \gamma_2 : \Gamma_2 \tag{2}$$

Moreover (using $\uplus$ for disjoint union),

- $\gamma^{\mathbf{L}} = {\gamma_1}^{\mathbf{L}} \uplus {\gamma_2}^{\mathbf{L}}$,
- $\gamma^{\mathbf{A}} = {\gamma_1}^{\mathbf{A}} \uplus {\gamma_2}^{\mathbf{A}}$,
- $\gamma^{\mathbf{U}} = {\gamma_1}^{\mathbf{U}} = {\gamma_2}^{\mathbf{U}}$,
- $\gamma_{\#} = \gamma_{1_{\#}} \uplus \gamma_{2_{\#}}$ (this splitting does not distinguish potentially unrestricted or affine bindings)

We establish the assumptions for the call eval $\delta\ \pi\ \gamma_1$ i' e_1.

(A1-1) From inversion: $C \wedge D \mid \Gamma_1 \vdash_{\mathrm{s}} e_1 : \tau_1$
(A1-2) From (1): $\Delta \vdash \gamma_1 : \Gamma_1$
(A1-3) From assumption
(A1-4) From assumption
(A1-5) From assumption because γ_1 is projected from γ.
(A1-6) From assumption because γ_1 is projected from γ.
(A1-7) From assumption because γ_1 is projected from γ.

Hence, we can apply the induction hypothesis and obtain

(R1-1) $R_1 = \mathrm{Ok}(\delta_1, \pi_1, r_1)$
(R1-2) $\Delta \leq \Delta_1, \delta \leq \delta_1, \vdash \delta_1 : \Delta_1$
(R1-3) $\Delta_1 \vdash r_1 : \tau_1$
(R1-4) π_1 is wellformed and $\mathrm{getloc}(\pi_1) \subseteq \mathrm{dom}(\delta_1) \setminus {\delta_1}^{-1}(\bullet)$.
(R1-5) $\mathrm{reach}_0(r_1) \subseteq \pi_1$, $\mathrm{reach}_{\delta_1}(r_1) \subseteq {\downarrow}\pi_1 \cap (\mathrm{reach}_{\delta_1}(\gamma) \setminus \mathrm{reach}_{\delta_1}(\gamma_{\#}) \cup \mathrm{dom}(\delta_1) \setminus \mathrm{dom}(\delta))$.
(R1-6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_1}(\gamma_1))$ it must be that
- $\delta_1(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_1$.

(R1-7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_1}({\gamma^{\mathrm{U}}}, {\gamma_{\#}^{\mathrm{U}}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta_1(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_1$.

(R1-8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{A}}, \gamma_{\#}^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_1(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_1}(\gamma_{\#}^{\mathbf{A}})$, then $\rho \in \pi_1$.

(R1-9) Resources: Let $\Theta = \mathrm{reach}_{\delta}(\gamma^{\mathbf{L}})$. Let $\Theta_1 = \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_1(\ell) = 1$, $\ell \notin \pi_1$, and if $\delta(\ell)$ is a resource, then $\delta_1(\ell) = \bullet$.

(R1-10) No thin air permission:
$\pi_1 \subseteq \pi \cup (\mathrm{dom}(\delta_1) \setminus \mathrm{dom}(\delta))$.

To establish the assumptions for the call
eval δ_1 π_1 γ_2(x↦r_1) i' e_2, we write $\gamma_2' = \gamma(x \mapsto r_1)$.

(A2-1) From inversion: $C \mid \Gamma; (x : \tau_1) \vdash_{\mathsf{s}} e_2 : \tau_2$

(A2-2) From (2) we have $\Delta \vdash \gamma_2 : \Gamma_2$. By store weakening (Lemma G.1) and using (R1-2), we have $\Delta_1 \vdash \gamma_2 : \Gamma_2$. With (R1-3), we obtain $\Delta_1 \vdash \gamma_2(x \mapsto r_1) : \Gamma_2; (x : \tau_1)$.

(A2-3) Immediate from (R1-2).

(A2-4) Immediate from (R1-4).

(A2-5) Show $\mathrm{reach}_0(\gamma_2') \subseteq \pi_1$, $\mathrm{reach}_{\delta_1}(\gamma_2') \subseteq \downarrow\pi_1$.
From (A1-5), we have $\mathrm{reach}_0(\gamma_2) \subseteq \pi_1$, $\mathrm{reach}_{\delta_1}(\gamma_2) \subseteq \downarrow\pi_1$. The extra binding $(x \mapsto r_1)$ goes into one of the compartments according to its type. We conclude by (R1-5).

(A2-6) Disjointness holds by assumption for γ_2 and it remains to discuss r_1. But r_1 is either a fresh resource, a linear/affine resource from γ_1 (which is disjoint), or unrestricted. In each case, there is no overlap with another compartment of the environment.

(A2-7) We need to show Incoming Resources:
(a) $\forall \ell \in \mathrm{getloc}(\mathrm{reach}_{\delta_1}(\gamma_2'))$, $\delta_1(\ell) \neq \bullet$.
(b) $\forall \ell \in \Theta_1 = \mathrm{getloc}(\mathrm{reach}_{\delta_1}(\gamma_2'^{\mathbf{L}}, \gamma_2'^{\mathbf{A}}, \gamma_{2\#}'^{\mathbf{A}}))$, $\Theta_1(\ell) = 1$.
The first item holds by assumption, splitting, and (for r_1) by (R1-4) and (R1-5).
The second and third items hold by assumption (A1-7), splitting, and framing (R1-6).

Hence, we can apply the induction hypothesis and obtain

(R2-1) $R_2 = \mathrm{Ok}(\delta_2, \pi_2, r_2)$

(R2-2) $\Delta_1 \leq \Delta_2$, $\delta_1 \leq \delta_2$, $\vdash \delta_2 : \Delta_2$

(R2-3) $\Delta_2 \vdash r_2 : \tau_2$

(R2-4) π_2 is wellformed and $\mathrm{getloc}(\pi_2) \subseteq \mathrm{dom}(\delta_2) \setminus {\delta_2}^{-1}(\bullet)$.

(R2-5) $\mathrm{reach}_0(r_2) \subseteq \pi_2$, $\mathrm{reach}_{\delta_2}(r_2) \subseteq \downarrow\pi_2 \cap (\mathrm{reach}_{\delta_2}(\gamma_1) \setminus \mathrm{reach}_{\delta_2}(\gamma_{1\#}) \cup \mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta_1))$.

(R2-6) Frame:
For all $\ell \in \mathrm{dom}(\delta_1) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma_2'))$ it must be that
- $\delta_2(\ell) = \delta_1(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi_1 \Leftrightarrow \rho \in \pi_2$.

(R2-7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma_2'^{\mathbf{U}}, \gamma_{2\#}'^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta_1)$, $\delta_2(\ell) = \delta_1(\ell) \neq \bullet$ and $\rho \in \pi_2$.

(R2-8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma_2'^{\mathbf{A}}, \gamma_{2\#}'^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta_1)$. If $\rho \neq \ell$, then $\delta_2(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_2}(\gamma_{2\#}'^{\mathbf{A}})$, then $\rho \in \pi_2$.

(R2-9) Resources: Let $\Theta_1 = \mathrm{reach}_{\delta_1}(\gamma_2'^{\mathbf{L}})$. Let $\Theta_2 = \mathrm{reach}_{\delta_2}(\gamma_2'^{\mathbf{L}})$.
For all $\ell \in \Theta_1$ it must be that $\Theta_1(\ell) = \Theta_2(\ell) = 1$, $\ell \notin \pi_2$, and if $\delta_1(\ell)$ is a resource, then $\delta_2(\ell) = \bullet$.

(R2-10) No thin air permission:
$\pi_2 \subseteq \pi_1 \cup (\mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta_1))$.

It remains to establish the assertions for the let expression.

(R-1) $R_2 = \mathrm{Ok}(\delta_2, \pi_2, r_2)$
Immediate from (R2-1).

(R-2) $\Delta \leq \Delta_2$, $\delta \leq \delta_2$, $\vdash \delta_2 : \Delta_2$
Transitivity of store extension (Lemma G.2), (R2-2), and (R1-2).

(R-3) $\Delta_2 \vdash r_2 : \tau_2$
Immediate from (R2-3).

(R-4) π_2 is wellformed and $\mathrm{getloc}(\pi_2) \subseteq \mathrm{dom}(\delta_2) \setminus {\delta_2}^{-1}(\bullet)$.
Immediate from (R2-4).

(R-5) $\mathrm{reach}_0(r_2) \subseteq \pi_2$, $\mathrm{reach}_{\delta_2}(r_2) \subseteq {\downarrow}\pi_2 \cap (\mathrm{reach}_{\delta_2}(\gamma) \setminus \mathrm{reach}_{\delta_2}(\gamma_\#) \cup \mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta))$.
Immediate from (R2-5) because $\mathrm{reach}_{\delta_2}(\gamma_1) \subseteq \mathrm{reach}_{\delta_2}(\gamma)$ and $\mathrm{reach}_{\delta_2}(\gamma_{1\#}) \subseteq \mathrm{reach}_{\delta_2}(\gamma_\#)$. Moreover, $\mathrm{dom}(\delta) \subseteq \mathrm{dom}(\delta_1)$.

(R-6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma))$ it must be that
- $\delta_2(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_2$.

Suppose that $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma))$.
Then $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_1}(\gamma_1))$.
By (R1-6), $\delta_1(\ell) = \delta(\ell)$ and for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$: $\rho \in \pi \Leftrightarrow \rho \in \pi_1$.
But also $\ell \in \mathrm{dom}(\delta_1) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma_2'))$.
By (R2-6), $\delta_2(\ell) = \delta_1(\ell)$ for applicable ρ, $\rho \in \pi_1 \Leftrightarrow \rho \in \pi_2$.
Taken together, we obtain the claim.

(R-7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma^{\mathbf{U}}, \gamma_\#^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta_2(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_2$.
Follows from (R2-7) or (R1-7) because $\gamma^{\mathbf{U}} = {\gamma_1}^{\mathbf{U}} = {\gamma_2}^{\mathbf{U}}$.

(R-8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_2(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_2}(\gamma_\#^{\mathbf{A}})$, then $\rho \in \pi_2$.
Follows from (R2-8), (R1-8), and framing.

(R-9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. Let $\Theta_2 = \mathrm{reach}_{\delta_2}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_2(\ell) = 1$, $\ell \notin \pi_2$, and if $\delta(\ell)$ is a resource, then $\delta_2(\ell) = \bullet$.
Follows from disjoint splitting of $\gamma^{\mathbf{L}}$, (R2-9), (R1-9), and framing.

(R-10) No thin air permission:
$\pi_2 \subseteq \pi \cup (\mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta))$.
Immediate from (R2-10).

Case e **of**
| VApp (x_1, x_2) →
let* $r_1 = \gamma(x_1)$ **in**
let* ℓ_1 = getloc r_1 **in**
let*? () = $\ell_1 \in \pi$ **in**
let* w = $\delta(\ell_1)$ **in**
let* (γ', k', x', e') = getstclos w **in**
let π' = (**if** k' ≤ **U then** π **else** $\pi - \ell_1$) **in**
let* δ' = $\delta(\ell_1) \leftarrow$ (**if** k' ≤ **U then** w **else** •) **in**
let* $r_2 = \gamma(x_2)$ **in**
let* (δ_3, π_3, r_3) = eval δ' π' γ'(x'↦ r_2) i' e' **in**
Ok (δ_3, π_3, r_3)

We need to invert rule App:

$$\frac{\begin{array}{c}\textsc{App}\\ (x_1 : \tau_2 \xrightarrow{\mathrm{k}} \tau_1) \in \Gamma \\ (x_2 : \tau_2') \in \Gamma \\ C \vdash_{\mathrm{e}} (\tau_2' \le \tau_2) \\ C \vdash_{\mathrm{e}} (\Gamma \backslash \{x_1, x_2\} \le \mathbf{A}_\infty)\end{array}}{C \mid \Gamma \vdash_{\mathrm{s}} (x_1\ x_2) : \tau_1}$$

We need to establish the assumptions for the recursive call eval δ' π' γ'(x'↦r_2) i' e'. We write $\gamma_2' = \gamma'(x' \mapsto r_2)$.

(A1-1) $C' \mid \Gamma'; (x' : \tau_2) \vdash_{\mathrm{s}} e' : \tau_1$, for some C' and Γ'
Applying the first premise of App to $r_1 = \gamma(x_1)$, $\Delta \vdash \gamma : \Gamma$, and inversion of result typing yields that $r_1 = \ell_1$ with $\Delta(\ell_1) = \tau_2 \xrightarrow{\mathrm{k}} \tau_1$. By inversion of store typing and storable typing, we find that there exist Γ' and C' such that
(a) $\delta(\ell_1) = (\gamma', \lambda^k x'.e')$
(b) dis(Γ')
(c) $\Delta \vdash \gamma', \Gamma'$
(d) $C' \mid \Gamma'; (x' : \tau_2) \vdash_{\mathrm{s}} e' : \tau_1$
(e) $C' \vdash_{\mathrm{e}} (\Gamma' \le k)$

(A1-2) $\Delta \vdash \gamma'(x' \mapsto r_2) : \Gamma'; (x' : \tau_2)$
By (A1-1)c, assumption on γ, and the subtyping premise.

(A1-3) $\vdash \delta' : \Delta'$
by assumption and the released rule of store typing (where we write $\Delta' = \Delta$ henceforth)

(A1-4) π' is wellformed and getloc(π') ⊆ dom(δ') \ δ'^{-1}(•)
the possible removal of a permission does not violate wellformedness; the permission is taken away exactly when the closure is destroyed

(A1-5) $\mathrm{reach}_0(\gamma_2') \subseteq \pi$, $\mathrm{reach}_\delta(\gamma_2') \subseteq \downarrow\pi$.
as the reach set is a subset of the incoming environment's reach

(A1-6) getloc($\gamma_2'^{\mathbf{L}}$), getloc($\gamma_2'^{\mathbf{A}}$), getloc($\gamma_2'^{\mathbf{U}}$), and getloc($\gamma_{2\#}'$) are all disjoint follows from dis(Γ') and since $r_2 = \Gamma(x_2)$ which is an entry disjoint from the closure $\Gamma(x_1)$.

(A1-7) Incoming Resources:
(a) $\forall \ell \in \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma_2'))$, $\delta'(\ell) \neq \bullet$.
(b) $\forall \ell \in \Theta' = \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma_2'^{\mathbf{L}}, \gamma_2'^{\mathbf{A}}, \gamma_{2\#}'^{\mathbf{A}}))$, $\Theta'(\ell) = 1$.
The first item holds because of assumption (A1-7).
The second item holds because $\Theta' \subseteq \Theta$ from assumption (A1-7).

The inductive hypothesis yields that ∃ δ_3, π_3, r_3, Δ_3 such that

(R1-1) $R_3 = \mathrm{Ok}(\delta_3, \pi_3, r_3)$
(R1-2) $\Delta' \leq \Delta_3$, $\delta' \leq \delta_3$, $\vdash \delta_3 : \Delta_3$
(R1-3) $\Delta_3 \vdash r_3 : \tau_1$
(R1-4) π_3 is wellformed and $\mathrm{getloc}(\pi_3) \subseteq \mathrm{dom}(\delta_3) \setminus {\delta_3}^{-1}(\bullet)$.
(R1-5) $\mathrm{reach}_0(r_3) \subseteq \pi_3$, $\mathrm{reach}_{\delta_3}(r_3) \subseteq \downarrow\pi_3 \cap (\mathrm{reach}_{\delta_3}(\gamma) \setminus \mathrm{reach}_{\delta_3}(\gamma_\#) \cup \mathrm{dom}(\delta_3) \setminus \mathrm{dom}(\delta))$.
(R1-6) Frame:
For all $\ell \in \mathrm{dom}(\delta') \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_3}(\gamma_2'))$ it must be that
- $\delta_3(\ell) = \delta'(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi' \Leftrightarrow \rho \in \pi_3$.

(R1-7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_3}(\gamma_2'^{\mathbf{U}}, \gamma_{2\#}'^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta')$, $\delta_3(\ell) = \delta'(\ell) \neq \bullet$ and $\rho \in \pi_3$.
(R1-8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_3}(\gamma_2'^{\mathbf{A}}, \gamma_{2\#}'^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta')$. If $\rho \neq \ell$, then $\delta_3(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_3}(\gamma_{2\#}'^{\mathbf{A}})$, then $\rho \in \pi_3$.
(R1-9) Resources: Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma_2'^{\mathbf{L}})$. Let $\Theta_3 = \mathrm{reach}_{\delta_3}(\gamma_2'^{\mathbf{L}})$.
For all $\ell \in \Theta'$ it must be that $\Theta'(\ell) = \Theta_3(\ell) = 1$, $\ell \notin \pi_3$, and if $\delta'(\ell)$ is a resource, then $\delta_3(\ell) = \bullet$.
(R1-10) No thin air permission:
$\pi_3 \subseteq \pi' \cup (\mathrm{dom}(\delta_3) \setminus \mathrm{dom}(\delta'))$.

The desired results are immediate because $\mathrm{dom}(\delta) = \mathrm{dom}(\delta')$.

Case e **of**
| Region (e, n, x, τ_x, b) →
let+ ρ = γ(x) **in**
let* ρ' = b.ρ **in**
let* π' = reach ρ τ_x δ **in**
let* π'' = b.π' **in**
let γ' = γ(x ↦ρ') **in**
let π = (π ∪ π'') \ π' **in**
let* (δ_1, π_1, r_1) = eval δ π γ' i' e **in**
let π_1 = (π_1 \ π'') ∪ π' **in**
Ok (δ_1, π_1, r_1)
We need to invert rule Region

$$\textsc{Region}\quad \frac{[x:\tau_x]^n_b \in \Gamma \qquad C \vdash_e \Gamma \rightsquigarrow^x_n \Gamma' \qquad C \mid \Gamma' \vdash_s e : \tau \qquad C \vdash_e (\tau \le \mathbf{L}_{n-1})}{C \mid \Gamma \vdash_s \{\!|e|\!\}^n_{\{x \mapsto b\}} : \tau}$$

We need to establish the assumptions for the recursive call eval δ π' γ' i' e' where $\gamma' = \gamma(x \mapsto \rho')$.

(A1-1) $C \mid \Gamma' \vdash_s e : \tau$
immediate from the inverted premise
(A1-2) $\Delta \vdash \gamma' : \Gamma'$
the only change of the environments is at x; adding the borrow modifier b succeeds due to the second premise; the address ρ' stored into x is compatible with its type by store typing
(A1-3) $\vdash \delta : \Delta$
Immediate by outer assumption
(A1-4) π is wellformed and getloc(π) ⊆ dom(δ) \ $\delta^{-1}(\bullet)$
Immediate by outer assumption; adding the modifier does not change the underlying raw location
(A1-5) reach$_0(\gamma') \subseteq \pi$, reach$_\delta(\gamma') \subseteq \downarrow\pi$.
locations were swapped simultaneously
(A1-6) getloc($\gamma'^{\mathbf{L}}$), getloc($\gamma'^{\mathbf{A}}$), getloc($\gamma'^{\mathbf{U}}$), and getloc($\gamma'_{\#}$) are all disjoint
Immediate by assumption
(A1-7) Incoming Resources:
(a) $\forall \ell \in$ getloc(reach$_\delta(\gamma')$), $\delta(\ell) \neq \bullet$.
(b) $\forall \ell \in \Theta =$ getloc(reach$_\delta(\gamma'^{\mathbf{L}}, \gamma'^{\mathbf{A}}, \gamma'^{\mathbf{A}}_{\#})$), $\Theta(\ell) = 1$.
Immediate by assumption.

The induction hypothesis yields the following statements. ∃ δ_1, π_1, r_1, Δ_1 such that

(R1-1) $R_1 = \text{Ok}(\delta_1, \pi_1, r_1)$
(R1-2) $\Delta \le \Delta_1$, $\delta \le \delta_1$, $\vdash \delta_1 : \Delta_1$
(R1-3) $\Delta_1 \vdash r_1 : \tau$
(R1-4) π_1 is wellformed and getloc(π_1) ⊆ dom(δ_1) \ ${\delta_1}^{-1}(\bullet)$.
(R1-5) reach$_0(r_1) \subseteq \pi_1$, reach$_{\delta_1}(r_1) \subseteq \downarrow\pi_1 \cap$(reach$_{\delta_1}(\gamma)$ \ reach$_{\delta_1}(\gamma_{\#})$ ∪ dom(δ_1) \ dom(δ)).
(R1-6) Frame:
For all $\ell \in$ dom(δ) \ getloc(reach$_{\delta_1}(\gamma')$) it must be that
- $\delta_1(\ell) = \delta(\ell)$ and
- for any ρ with getloc(ρ) = {ℓ}, $\rho \in \pi \Leftrightarrow \rho \in \pi_1$.

(R1-7) Unrestricted values, resources, and borrows:
For all $\rho \in \text{reach}_{\delta_1}(\gamma'^{\mathbf{U}}, \gamma'^{\mathbf{U}}_{\#})$ with $\text{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \text{dom}(\delta)$, $\delta_1(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_1$.

(R1-8) Affine borrows and resources:
For all $\rho \in \text{reach}_{\delta_1}(\gamma'^{\mathbf{A}}, \gamma'^{\mathbf{A}}_{\#})$ with $\text{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \text{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_1(\ell) \neq \bullet$. If $\rho \in \text{reach}_{\delta_1}(\gamma'^{\mathbf{A}}_{\#})$, then $\rho \in \pi_1$.

(R1-9) Resources: Let $\Theta = \text{reach}_{\delta}(\gamma'^{\mathbf{L}})$. Let $\Theta_1 = \text{reach}_{\delta_1}(\gamma'^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_1(\ell) = 1$, $\ell \notin \pi_1$, and if $\delta(\ell)$ is a resource, then $\delta_1(\ell) = \bullet$.

(R1-10) No thin air permission:
$\pi_1 \subseteq \pi \cup (\text{dom}(\delta_1) \setminus \text{dom}(\delta))$.

It remains to derive the induction hypothesis in the last line. The only additional action is the exchange of permissions which withdraws the borrow.

(R1) $R_1 = \text{Ok}(\delta_1, \pi_1, r_1)$
Immediate

(R2) $\Delta \leq \Delta_1, \delta \leq \delta_1, \vdash \delta_1 : \Delta_1$
Immediate

(R3) $\Delta_1 \vdash r_1 : \tau$
Immediate

(R4) π_1 is wellformed and $\text{getloc}(\pi_1) \subseteq \text{dom}(\delta_1) \setminus {\delta_1}^{-1}(\bullet)$.
The addresses ρ and ρ' (as well as the elements of π' and π'') have the same raw location, so exchanging them does not affect wellformedness. The underlying set of locations does not change.

(R5) $\text{reach}_0(r_1) \subseteq \pi_1$, $\text{reach}_{\delta_1}(r_1) \subseteq \downarrow\pi_1 \cap (\text{reach}_{\delta_1}(\gamma) \setminus \text{reach}_{\delta_1}(\gamma_{\#}) \cup \text{dom}(\delta_1) \setminus \text{dom}(\delta))$.
This case is critical for region encapsulation. Here we need to argue that ρ' (and hence π'') is not reachable from r_1 because its type τ is bounded by $\mathbf{L}_{n-1}$ according to the fourth premise. We conclude with Lemma G.5.

(R6) Frame:
For all $\ell \in \text{dom}(\delta) \setminus \text{getloc}(\text{reach}_{\delta_1}(\gamma'))$ it must be that
- $\delta_1(\ell) = \delta(\ell)$ and
- for any ρ with $\text{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_1$.

Immediate

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \text{reach}_{\delta_1}(\gamma'^{\mathbf{U}}, \gamma'^{\mathbf{U}}_{\#})$ with $\text{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \text{dom}(\delta)$, $\delta_1(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_1$.
Immediate

(R8) Affine borrows and resources:
For all $\rho \in \text{reach}_{\delta_1}(\gamma'^{\mathbf{A}}, \gamma'^{\mathbf{A}}_{\#})$ with $\text{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \text{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_1(\ell) \neq \bullet$. If $\rho \in \text{reach}_{\delta_1}(\gamma'^{\mathbf{A}}_{\#})$, then $\rho \in \pi_1$.
Immediate

(R9) Resources: Let $\Theta = \text{reach}_{\delta}(\gamma'^{\mathbf{L}})$. Let $\Theta_1 = \text{reach}_{\delta_1}(\gamma'^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_1(\ell) = 1$, $\ell \notin \pi_1$, and if $\delta(\ell)$ is a resource, then $\delta_1(\ell) = \bullet$.
Immediate

(R10) No thin air permission:
$\pi_1 \subseteq \pi \cup (\text{dom}(\delta_1) \setminus \text{dom}(\delta))$.
Immediate

Case e **of**
| Create (x) $\rightarrow$
let* r = γ(x) **in**
let w = STRSRC (r) **in**
let (ℓ_1, δ_1) = salloc δ w **in**
let $\pi_1 = \pi + \ell_1$ **in**
Ok $(\delta_1, \pi_1, \ell_1)$

We need to invert the corresponding rule

$$\textsc{Create}\ \frac{C \mid \Gamma \vdash_{\mathrm{s}} \tau : k \qquad C \vdash_{\mathrm{e}} (k \le \mathbf{U}_0) \wedge (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\mathrm{s}} \mathtt{create} : \tau \rightarrow \mathrm{R}\ \tau}$$

It is sufficient to show that there is some $\Delta_1 = \Delta(\ell_1 : \mathrm{R}\ \tau)$ such that δ_1, π_1, and $r_1 = \ell_1$ fulfill the following requirements.

(R1) $R_1 = \mathrm{Ok}(\delta_1, \pi_1, r_1)$

(R2) $\Delta \le \Delta_1, \delta \le \delta_1, \vdash \delta_1 : \Delta_1$
For the last item, we need to show that $\Delta(\ell_1) : \mathrm{R}\ \tau$, but this follows from the setting of w to a resource storable in the semantics.

(R3) $\Delta_1 \vdash r_1 : \mathrm{R}\ \tau$
Immediate from the discussion of the preceding case

(R4) π_1 is wellformed and $\mathrm{getloc}(\pi_1) \subseteq \mathrm{dom}(\delta_1) \setminus \delta_1^{-1}(\bullet)$.
Follows from the assumption on π and for ℓ_1 from the allocation of the resource.

(R5) $\mathrm{reach}_0(r_1) \subseteq \pi_1$, $\mathrm{reach}_{\delta_1}(r_1) \subseteq \downarrow\pi_1 \cap (\mathrm{reach}_{\delta_1}(\gamma) \setminus \mathrm{reach}_{\delta_1}(\gamma_\#) \cup \mathrm{dom}(\delta_1) \setminus \mathrm{dom}(\delta))$.
Immediate from the assignment to π_1.

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_1}(\gamma))$ it must be that
- $\delta_1(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_1$.

Obvious as no existing location is changed.

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{U}}, \gamma_\#^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta_1(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_1$.
Obvious as no existing location has changed and no permission is withdrawn.

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_1(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_1}(\gamma_\#^{\mathbf{A}})$, then $\rho \in \pi_1$.
Obvious as no existing location has changed and no permission is withdrawn.

(R9) Resources: Let $\Theta = \mathrm{reach}_{\delta}(\gamma^{\mathbf{L}})$. Let $\Theta_1 = \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_1(\ell) = 1$, $\ell \notin \pi_1$, and if $\delta(\ell)$ is a resource, then $\delta_1(\ell) = \bullet$.
By the constraint on Γ in the Create rule, $\gamma^{\mathbf{L}} = \emptyset$.

(R10) No thin air permission:
$\pi_1 \subseteq \pi \cup (\mathrm{dom}(\delta_1) \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| VDestroy (x) →
let* r = γ(x) **in**
let* ρ = getaddress r **in**
let* ℓ = getloc r **in**
let* w = $\delta(\ell)$ **in**
let* r = getstrsrc w **in**
let*? () = $\rho \in \pi$ **in**
let* $\delta_1 = \delta(\ell) \leftarrow \bullet$ **in**
let $\pi_1 = \pi - \ell$ **in**
Ok (δ_1, π_1, ())

We need to invert rule Destroy.

$$\textsc{Destroy}\ \frac{C \mid \Gamma \vdash_s \tau : k \qquad C \vdash_e (k \leq \mathbf{U}_0) \wedge (\Gamma \leq \mathbf{A}_\infty)}{C \mid \Gamma \vdash_s \texttt{destroy} : \mathrm{R}\ \tau \rightarrow \texttt{Unit}}$$

It is sufficient to show that $\Delta_1 = \Delta$, δ_1, π_1, and $r_1 = ()$ fulfill the following requirements.

(R1) $R_1 = \mathrm{Ok}(\delta_1, \pi_1, r_1)$

(R2) $\Delta \leq \Delta_1, \delta \leq \delta_1, \vdash \delta_1 : \Delta_1$
Immediate: ℓ was updated to void, which has any type.

(R3) $\Delta_1 \vdash () : \texttt{Unit}$

(R4) π_1 is wellformed and $\mathrm{getloc}(\pi_1) \subseteq \mathrm{dom}(\delta_1) \setminus {\delta_1}^{-1}(\bullet)$.
By assumption on π and because ℓ was removed.

(R5) $\mathrm{reach}_0(r_1) \subseteq \pi_1$, $\mathrm{reach}_{\delta_1}(r_1) \subseteq \downarrow\pi_1 \cap (\mathrm{reach}_{\delta_1}(\gamma) \setminus \mathrm{reach}_{\delta_1}(\gamma_\#) \cup \mathrm{dom}(\delta_1) \setminus \mathrm{dom}(\delta))$.
Immediate because the reach set is empty

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_1}(\gamma))$ it must be that
- $\delta_1(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_1$.

Only $\delta(\ell)$ was changed, which is not reachable from the frame.

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{U}}, \gamma^{\mathbf{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta_1(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_1$.
Immediate because we updated (destroyed) a resource (in $\gamma^{\mathbf{L}}$).

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{A}}, \gamma^{\mathbf{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_1(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{A}}_\#)$, then $\rho \in \pi_1$.
Immediate because we updated (destroyed) a resource (in $\gamma^{\mathbf{L}}$).

(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. Let $\Theta_1 = \mathrm{reach}_{\delta_1}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_1(\ell) = 1$, $\ell \notin \pi_1$, and if $\delta(\ell)$ is a resource, then $\delta_1(\ell) = \bullet$.
By the constraint on Γ, ℓ was the only resource passed to this invocation of eval. The claimed condition holds as ℓ was removed from π_1 and the location's contents cleared.

(R10) No thin air permission:
$\pi_1 \subseteq \pi \cup (\mathrm{dom}(\delta_1) \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| Var (x) →
let* r = γ(x) **in**
Ok (δ, π, r)
We need to invert rule VAR.

$$\textsc{Var}\ \frac{(x:\tau)\in\Gamma \qquad C\vdash_{\mathrm{e}}(\Gamma\backslash\{x\}\le \mathbf{A}_\infty)}{C\mid\Gamma\vdash_{\mathrm{s}} x:\tau}$$

We establish that the claims hold for $\delta' = \delta$, $\pi' = \pi$, $r = \gamma(x)$, and $\Delta' = \Delta$.

(R1) $R = \mathrm{Ok}(\delta, \pi, r)$

(R2) $\Delta \le \Delta, \delta \le \delta, \vdash \delta : \Delta$
Immediate by reflexivity and assumption.

(R3) $\Delta \vdash r : \tau$
Immediate by assumption (A1-1)c.

(R4) π is wellformed and $\mathrm{getloc}(\pi) \subseteq \mathrm{dom}(\delta) \setminus \delta^{-1}(\bullet)$.
Immediate by assumption (A1-4)

(R5) $\mathrm{reach}_0(r) \subseteq \pi$, $\mathrm{reach}_\delta(r) \subseteq \downarrow\pi \cap (\mathrm{reach}_\delta(\gamma) \setminus \mathrm{reach}_\delta(\gamma_\#) \cup \mathrm{dom}(\delta) \setminus \mathrm{dom}(\delta))$.
Immediate

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_\delta(\gamma))$ it must be that
- $\delta(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi$.

Immediate as permissions and store stay the same.

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_\delta(\gamma^{\mathbf{U}}, \gamma^{\mathbf{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi$.
Immediate

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_\delta(\gamma^{\mathbf{A}}, \gamma^{\mathbf{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_\delta(\gamma^{\mathbf{A}}_\#)$, then $\rho \in \pi$.
Immediate

(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta(\ell) = 1$, $\ell \notin \pi$, and if $\delta(\ell)$ is a resource, then $\delta(\ell) = \bullet$.
As π remains the same, a linear resource in x is returned untouched.

(R10) No thin air permission:
$\pi \subseteq \pi \cup (\mathrm{dom}(\delta) \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| Const (c) →
Ok (δ, π, c)
We need to invert rule Const.

$$\frac{\textsc{Const}\quad C \vdash_{\mathrm{e}} (\Gamma \leq \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\mathrm{s}} c : \mathrm{CType}(c)}$$

We need to establish the claims for $\delta' = \delta$, $\pi' = \pi$, $r' = c$, and $\Delta' = \Delta$:

(R1) $R = \mathrm{Ok}(\delta, \pi, r)$
(R2) $\Delta \leq \Delta, \delta \leq \delta, \vdash \delta : \Delta$
By assumption (A3).
(R3) $\Delta \vdash c : \mathrm{CType}(c)$
by result typing.
(R4) π is wellformed and $\mathrm{getloc}(\pi) \subseteq \mathrm{dom}(\delta) \setminus \delta^{-1}(\bullet)$.
By assumption (A1-4).
(R5) $\mathrm{reach}_0(r) \subseteq \pi$, $\mathrm{reach}_\delta(r) \subseteq \downarrow\pi \cap (\mathrm{reach}_\delta(\gamma) \setminus \mathrm{reach}_\delta(\gamma_\#) \cup \mathrm{dom}(\delta) \setminus \mathrm{dom}(\delta))$.
As $\mathrm{reach}_\delta(c) = \emptyset$.
(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_\delta(\gamma))$ it must be that
- $\delta(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi$.

Immediate
(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_\delta(\gamma^{\mathrm{U}}, \gamma^{\mathrm{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi$.
Immediate
(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_\delta(\gamma^{\mathrm{A}}, \gamma^{\mathrm{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_\delta(\gamma^{\mathrm{A}}_\#)$, then $\rho \in \pi$.
(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathrm{L}})$. Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathrm{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta(\ell) = 1$, $\ell \notin \pi$, and if $\delta(\ell)$ is a resource, then $\delta(\ell) = \bullet$.
Immediate as $\gamma^{\mathrm{L}} = \emptyset$.
(R10) No thin air permission:
$\pi \subseteq \pi \cup (\mathrm{dom}(\delta) \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| VPair (k, x_1, x_2) →
let* $r_1 = \gamma(x_1)$ **in**
let* $r_2 = \gamma(x_2)$ **in**
let w = STPAIR (k, r_1, r_2) **in**
let (ℓ', δ') = salloc δ w **in**
let $\pi' = \pi + \ell'$ **in**
Ok (δ', π', ℓ')

We need to invert rule Pair.

$$\textsc{Pair}\ \frac{\begin{array}{c}(x_1 : \tau_1) \in \Gamma \\ (x_2 : \tau_2) \in \Gamma \\ C \vdash_{\mathrm{e}} (\Gamma \backslash \{x_1, x_2\} \leq \mathbf{A}_\infty)\end{array}}{C \mid \Gamma \vdash_{\mathrm{s}} (x_1, x_2)^k : \tau_1 \times \tau_2}$$

Show that $\delta', \pi', r' = \ell', \Delta' = \Delta(\ell' : \tau_1 \times^k \tau_2)$ such that

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$

(R2) $\Delta \leq \Delta', \delta \leq \delta', \vdash \delta' : \Delta'$

(R3) $\Delta' \vdash \ell' : \tau_1 \times \tau_2$

(R4) π' is wellformed and $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
By assumption (A1-4) and because ℓ' is properly initialized.

(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
By assumption (A1-5), $\mathrm{reach}_{\delta'}(\ell') = \mathrm{reach}_\delta(r_1, r_2) \cup \{\ell'\} \subseteq \mathrm{reach}_\delta(\gamma) \cup \{\ell'\}$ and $\{\ell'\} = \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta)$.

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ it must be that
- $\delta'(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

Immediate

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{U}}, \gamma_\#^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi'$.
Immediate

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta'(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta'}(\gamma_\#^{\mathbf{A}})$, then $\rho \in \pi'$.
Immediate

(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, and if $\delta(\ell)$ is a resource, then $\delta'(\ell) = \bullet$.
Every such ℓ must be reachable either from r_1 or r_2. So they become reachable from ℓ', as required.

(R10) No thin air permission:
$\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| Lam (k, x, e) →
let w = STCLOS (γ, k, x, e) **in**
let (ℓ', δ') = salloc δ w **in**
let π' = π + ℓ' **in**
Ok (δ', π', ℓ')
We need to invert rule Abs

$$\frac{\textsc{Abs}\quad C \mid \Gamma; (x : \tau_2) \vdash_{\mathrm{s}} e : \tau_1 \quad C \vdash_{\mathrm{e}} (\Gamma \leq k)}{C \mid \Gamma \vdash_{\mathrm{s}} \lambda x.e : \tau_2 \xrightarrow{\mathrm{k}} \tau_1}$$

Show that δ', π', $r' = \ell'$, and $\Delta' = \Delta(\ell' : \tau_2 \xrightarrow{\mathrm{k}} \tau_1)$ fulfill

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$

(R2) $\Delta \leq \Delta'$, $\delta \leq \delta'$, $\vdash \delta' : \Delta'$
Immediate by definition and store typing

(R3) $\Delta' \vdash r' : \tau_2 \xrightarrow{\mathrm{k}} \tau_1$
Immediate by store typing

(R4) π' is wellformed and $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
Wellformedness holds by assumption on π and because ℓ' is a new location. The domain constraint is assumed for π and ℓ' is initialized to a closure.

(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.

$$\begin{aligned}\mathrm{reach}_{\delta'}(r') &= \{\ell'\} \cup \mathrm{reach}_{\delta'}(\gamma)\\ &= \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta) \cup \mathrm{reach}_{\delta'}(\gamma)\end{aligned}$$

Moreover, the constraint $(\Gamma \leq k)$ implies that $\gamma_\# = \emptyset$.

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ it must be that
- $\delta'(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

Immediate

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathrm{U}}, \gamma^{\mathrm{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi'$.
Immediate

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathrm{A}}, \gamma^{\mathrm{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta'(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathrm{A}}_\#)$, then $\rho \in \pi'$.
Immediate

(R9) Resources: Let $\Theta = \mathrm{reach}_{\delta}(\gamma^{\mathrm{L}})$. Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathrm{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, and if $\delta(\ell)$ is a resource, then $\delta'(\ell) = \bullet$.
The second case is immediately applicable.

(R10) No thin air permission:
$\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| Borrow (b, x) →
let+ $\rho = \gamma(\mathrm{x})$ **in**
let*? () = ρ ? b **&&** $\rho \in \pi$ **in**
Ok (δ, π, ρ)

We have to invert rule Borrow

$$\frac{\begin{array}{c}(x \div \sigma)^k_b \in \Gamma \qquad C_x, \tau_x = \mathrm{Inst}(\Gamma, \sigma) \\ C \vdash_{\mathrm{e}} C_x \wedge (\Gamma \backslash \{x\} \leq \mathbf{A}_\infty)\end{array}}{C \mid \Gamma \vdash_{\mathrm{s}} \&^b x : \&^b(k, \tau_x)} \textsc{Borrow}$$

Show that $\delta' = \delta$, $\pi' = \pi$, $r' = \rho$, $\Delta' = \Delta$ such that

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$

(R2) $\Delta \leq \Delta'$, $\delta \leq \delta'$, $\vdash \delta' : \Delta'$
Immediate, no changes.

(R3) $\Delta' \vdash r' : \&^b(k, \tau)$
Immediate by result typing and because the interpreter checks that the permissions of the borrow are very restricted.

(R4) π' is wellformed and $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
Immediate (no change).

(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
By typing, ρ is not in $\gamma_\#$. Hence, the condition is immediate.

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ it must be that
- $\delta'(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

Immediate as no change.

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{U}}, \gamma^{\mathbf{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi'$.
Immediate as no change

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}, \gamma^{\mathbf{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta'(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}_\#)$, then $\rho \in \pi'$.
Immediate

(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, and if $\delta(\ell)$ is a resource, then $\delta'(\ell) = \bullet$.
Immediate because Θ, Θ' must be empty

(R10) No thin air permission:
$\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
Immediate.

Case e **of**
| VObserve (x) →
let* r = γ(x) **in**
let* ρ = getaddress r **in**
let*? () = $\rho \in \pi$ **in**
let* (b, _, ℓ) = getborrowed_loc r **in**
let*? () = (b = **U**) **in**
let* w = $\delta(\ell)$ **in**
let* r' = getstrsrc w **in**
Ok (δ, π, r')

We have to invert the rule Observe

$$\textsc{Observe}\ \frac{C \mid \Gamma \vdash_{\mathrm{s}} \tau : k \qquad C \vdash_{\mathrm{e}} (k \leq \mathbf{U}_0) \wedge (\Gamma \leq \mathbf{A}_\infty)}{C \mid \Gamma \vdash_{\mathrm{s}} \mathtt{observe} : \&^{\mathrm{U}}(k', \mathrm{R}\ \tau) \to \tau}$$

Show that $\delta' = \delta$, $\pi' = \pi$, r', and $\Delta' = \Delta$ fulfill

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$

(R2) $\Delta \leq \Delta'$, $\delta \leq \delta'$, $\vdash \delta' : \Delta'$
By reflexivity and assumption.

(R3) $\Delta' \vdash r' : \tau$
Immediate by store typing

(R4) π' is wellformed and $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
Immediate: no changes.

(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
Immediate

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ it must be that
- $\delta'(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

Immediate: no changes.

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathrm{U}}, \gamma^{\mathrm{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi'$.
Immediate: no changes to immutables.

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathrm{A}}, \gamma^{\mathrm{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta'(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathrm{A}}_\#)$, then $\rho \in \pi'$.
Immediate: one particular ρ is overwritten, but not freed.

(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathrm{L}})$. Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathrm{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, and if $\delta(\ell)$ is a resource, then $\delta'(\ell) = \bullet$.
Immediate because $\Theta = \emptyset$

(R10) No thin air permission:
$\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
Immediate

Case e **of**
| VUpdate (x_1, x_2) →
let* $r_1 = \gamma(x_1)$ **in**
let* ρ = getaddress r_1 **in**
let* (b, _, ℓ) = getborrowed_loc r_1 **in**
let*? () = (b = **A**) **in**
let* $r_2 = \gamma(x_2)$ **in**
let* w = $\delta(\ell)$ **in**
let* r = getstrsrc w **in**
let*? () = $\rho \in \pi$ **in**
let* δ' = $\delta(\ell) \leftarrow$ STRSRC (r_2) **in**
let π' = $\pi - \rho$ **in**
Ok (δ', π', ())

We need to invert rule Update

$$\frac{C \mid \Gamma \vdash_s \tau : k \qquad C \vdash_e (k \le \mathbf{U}_0) \wedge (\Gamma \le \mathbf{A}_\infty)}{C \mid \Gamma \vdash_s \mathtt{update} : \&^{\mathbf{A}}(k', \mathrm{R}\ \tau) \to \tau \xrightarrow{\mathbf{A}} \mathtt{Unit}} \textsc{Update}$$

We need to show that $\delta', \pi', r' = (), \Delta' = \Delta$ fulfill

(R1) $R' = \mathrm{Ok}(\delta', \pi', r')$

(R2) $\Delta \le \Delta', \delta \le \delta', \vdash \delta' : \Delta'$
Immediate by store typing for ℓ

(R3) $\Delta' \vdash r' : \mathtt{Unit}$
Immediate

(R4) π' is wellformed and $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$.
Immediate, as we remove a permission from π

(R5) $\mathrm{reach}_0(r') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(r') \subseteq \downarrow\pi' \cap (\mathrm{reach}_{\delta'}(\gamma) \setminus \mathrm{reach}_{\delta'}(\gamma_\#) \cup \mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.
Immediate, as we only update a reachable ℓ

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma))$ it must be that
- $\delta'(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi'$.

Immediate

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{U}}, \gamma^{\mathbf{U}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta'(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi'$.
Immediate

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}, \gamma^{\mathbf{A}}_\#)$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta'(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}_\#)$, then $\rho \in \pi'$.
Immediate; for ℓ, we observe that it is overwritten, but not freed.

(R9) Resources: Let $\Theta = \mathrm{reach}_\delta(\gamma^{\mathbf{L}})$. Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, and if $\delta(\ell)$ is a resource, then $\delta'(\ell) = \bullet$.
Immediate because $\gamma^{\mathbf{L}} = \emptyset$ and hence $\Theta = \emptyset$.

(R10) No thin air permission:
$\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.

Case e **of**
| VMatch (x, x', z, e_2, sp) $\rightarrow$
let (γ_1, γ_2) = vsplit γ sp **in**
let* r = γ_1(z) **in**
let* ℓ = getloc r **in**
let* w = $\delta(\ell)$ **in**
let* (k, r_1, r_1') = getstpair w **in**
let π' = **if** k $\leq$ U **then** π **else** $\pi - \ell$ **in**
let* δ' = $\delta(\ell) \leftarrow$ (**if** k $\leq$ U **then** w **else** $\bullet$) **in**
let γ_2' = γ_2(x $\mapsto$ r_1)(x' $\mapsto$ r_1') **in**
let* (δ_2, π_2, r_2) = eval δ' π' γ_2' i' e_2 **in**
Ok (δ_2, π_2, r_2)

We need to invert rule MatchPair

$$\textsc{MatchPair}\quad \frac{sp : C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad \Gamma_1 = (z : \phi(\tau_1 \times \tau_1')) \qquad C \mid \Gamma_2; (x : \phi(\tau_1)); (x' : \phi(\tau_1')) \vdash_s e_2 : \tau_2}{C \mid \Gamma \vdash_s \mathtt{match}_\phi\ x, x' = z\ \mathtt{in}\ e_2 : \tau_2}$$

The case VMatch corresponds to the match specification $\phi = \mathrm{id}$.

Establish the assumptions for the recursive call with $\gamma_2' = \gamma_2(x \mapsto r_1)(x' \mapsto r_1')$ and $\Delta' = \Delta$:

(A1-1) $C \mid \Gamma_2(x : \tau_1)(x' : \tau_1') \vdash_s e_2 : \tau_2$ by inversion

(A1-2) $\Delta \vdash \gamma_2 : \Gamma_2$ by assumption; moreover, $\Delta \vdash r_1 : \tau_1$ and $\Delta \vdash r_1' : \tau_1'$ by inversion of the store typing for ℓ. As $\Delta' = \Delta$, we have $\Delta' \vdash \gamma_2' : \Gamma_2(x : \tau_1)(x' : \tau_1')$.

(A1-3) $\vdash \delta' : \Delta'$: the only change from assumption is in ℓ which potentially maps to $\bullet$.

(A1-4) π' is wellformed and $\mathrm{getloc}(\pi') \subseteq \mathrm{dom}(\delta') \setminus \delta'^{-1}(\bullet)$: permission to ℓ is removed iff ℓ is mapped to $\bullet$.

(A1-5) $\mathrm{reach}_0(\gamma_2') \subseteq \pi'$, $\mathrm{reach}_{\delta'}(\gamma_2') \subseteq \downarrow\pi'$.
by assumption

(A1-6) $\mathrm{getloc}(\gamma_2'^{\mathbf{L}})$, $\mathrm{getloc}(\gamma_2'^{\mathbf{A}})$, $\mathrm{getloc}(\gamma_2'^{\mathbf{U}})$, and $\mathrm{getloc}(\gamma_{2\#}')$ are all disjoint: by assumption and splitting

(A1-7) Incoming Resources:
(a) $\forall \ell \in \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma_2')), \delta'(\ell) \neq \bullet$.
(b) $\forall \ell \in \Theta' = \mathrm{getloc}(\mathrm{reach}_{\delta'}(\gamma_2'^{\mathbf{L}}, \gamma_2'^{\mathbf{A}}, \gamma_{2\#}'^{\mathbf{A}})), \Theta'(\ell) = 1$.

Hence the call to eval yields $\exists\ \delta_2, \pi_2, r_2, \Delta_2$ such that

(R1-1) $R_2 = \mathrm{Ok}(\delta_2, \pi_2, r_2)$

(R1-2) $\Delta' \leq \Delta_2$, $\delta' \leq \delta_2$, $\vdash \delta_2 : \Delta_2$

(R1-3) $\Delta_2 \vdash r_2 : \tau_2$

(R1-4) π_2 is wellformed and $\mathrm{getloc}(\pi_2) \subseteq \mathrm{dom}(\delta_2) \setminus {\delta_2}^{-1}(\bullet)$.

(R1-5) $\mathrm{reach}_0(r_2) \subseteq \pi_2$, $\mathrm{reach}_{\delta_2}(r_2) \subseteq \downarrow\pi_2 \cap (\mathrm{reach}_{\delta_2}(\gamma) \setminus \mathrm{reach}_{\delta_2}(\gamma_\#) \cup \mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta))$.

(R1-6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma))$ it must be that
- $\delta_2(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_2$.

(R1-7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{U}}, \gamma_\#^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta), \delta'(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi'$.

(R1-8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta'}(\gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta'(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta'}(\gamma_\#^{\mathbf{A}})$, then $\rho \in \pi'$.

(R1-9) Resources: Let $\Theta = \mathrm{reach}_{\delta}(\gamma^{\mathbf{L}})$. Let $\Theta' = \mathrm{reach}_{\delta'}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta'(\ell) = 1$, $\ell \notin \pi'$, and if $\delta(\ell)$ is a resource, then $\delta'(\ell) = \bullet$.

(R1-10) No thin air permission:
$\pi' \subseteq \pi \cup (\mathrm{dom}(\delta') \setminus \mathrm{dom}(\delta))$.

As R_2 is also returned from the match, these results carry over.

Case e **of**
| VMatchborrow (x, x', z, e_2, sp) $\rightarrow$
let (γ_1, γ_2) = vsplit γ sp **in**
let* $r_1 = \gamma_1$(z) **in**
let* (b, _, ℓ) = getborrowed_loc r_1 **in**
let* w = $\delta(\ell)$ **in**
let* (k', r_1', r_2') = getstpair w **in**
let* ρ = getaddress r_1 **in**
let π'' = (**if** k' $\leq$ **U then** π **else** $\pi - \rho$) **in**
let δ'' = δ **in**
let* ρ_1 = getaddress r_1' **in**
let* ρ_2 = getaddress r_2' **in**
let* ρ_1' = b.ρ_1 **in**
let* ρ_2' = b.ρ_2 **in**
let r_1'' = ρ_1' **in**
let r_2'' = ρ_2' **in**
let γ_2'' = γ_2(x $\mapsto$ r_1'')(x' $\mapsto$ r_2'') **in**
let* (δ_2, π_2, r_2) = eval δ'' π'' γ_2'' i' e_2 **in**
Ok (δ_2, π_2, r_2)

We need to invert rule MatchPair

$$\frac{sp : C \vdash_e \Gamma = \Gamma_1 \ltimes \Gamma_2 \qquad \Gamma_1 = (z : \phi(\tau_1 \times \tau_1')) \qquad C \mid \Gamma_2; (x : \phi(\tau_1)); (x' : \phi(\tau_1')) \vdash_s e_2 : \tau_2}{C \mid \Gamma \vdash_s \mathtt{match}_\phi\ x, x' = z\ \mathtt{in}\ e_2 : \tau_2}$$

The case VMatchborrow corresponds to the match specification $\phi = \&^b$. In contrast to the non-borrowing match, the borrowed pair is never deallocated.

Establish the assumptions for the recursive call with $\Delta'' = \Delta$:

(A1-1) $C \mid \Gamma_2(x : \&^b\tau_1)(x' : \&^b\tau_1') \vdash_s e_2 : \tau_2$ by inversion

(A1-2) $\Delta \vdash \gamma_2 : \Gamma_2$ by assumption; moreover, $\Delta \vdash r_1'' : \&^b\tau_1$ and $\Delta \vdash r_2'' : \&^b\tau_1'$ by inversion of the store typing for ρ. As $\Delta'' = \Delta$, we have $\Delta'' \vdash \gamma_2'' : \Gamma_2(x : \&^b\tau_1)(x' : \&^b\tau_1')$.

(A1-3) $\vdash \delta' : \Delta''$: the only change from assumption is in ℓ which potentially maps to $\bullet$.

(A1-4) π'' is wellformed and $\mathrm{getloc}(\pi'') \subseteq \mathrm{dom}(\delta'') \setminus \delta''^{-1}(\bullet)$: permission to ℓ is removed iff ℓ is mapped to $\bullet$.

(A1-5) $\mathrm{reach}_0(\gamma_2'') \subseteq \pi''$, $\mathrm{reach}_{\delta''}(\gamma_2'') \subseteq \downarrow\pi''$.
by assumption

(A1-6) $\mathrm{getloc}(\gamma_2''^{\mathbf{L}})$, $\mathrm{getloc}(\gamma_2''^{\mathbf{A}})$, $\mathrm{getloc}(\gamma_2''^{\mathbf{U}})$, and $\mathrm{getloc}(\gamma_{2\#}'')$ are all disjoint: by assumption and splitting

(A1-7) Incoming Resources:
(a) $\forall \ell \in \mathrm{getloc}(\mathrm{reach}_{\delta''}(\gamma_2'))$, $\delta''(\ell) \neq \bullet$.
(b) $\forall \ell \in \Theta'' = \mathrm{getloc}(\mathrm{reach}_{\delta''}(\gamma_2'^{\mathbf{L}}, \gamma_2'^{\mathbf{A}}, \gamma_{2\#}'^{\mathbf{A}}))$, $\Theta''(\ell) = 1$.

Hence the call to eval yields δ_2, π_2, r_2, Δ_2 such that

(R1-1) $R_2 = \mathrm{Ok}(\delta_2, \pi_2, r_2)$

(R1-2) $\Delta'' \leq \Delta_2$, $\delta'' \leq \delta_2$, $\vdash \delta_2 : \Delta_2$

(R1-3) $\Delta_2 \vdash r_2 : \tau_2$

(R1-4) π_2 is wellformed and $\mathrm{getloc}(\pi_2) \subseteq \mathrm{dom}(\delta_2) \setminus {\delta_2}^{-1}(\bullet)$.

(R1-5) $\mathrm{reach}_0(r_2) \subseteq \pi_2$, $\mathrm{reach}_{\delta_2}(r_2) \subseteq \downarrow\pi_2 \cap (\mathrm{reach}_{\delta_2}(\gamma_2'') \setminus \mathrm{reach}_{\delta_2}(\gamma_{2\#}'') \cup \mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta_2''))$.

(R1-6) Frame:
For all $\ell \in \mathrm{dom}(\delta'') \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma_2''))$ it must be that
- $\delta_2(\ell) = \delta''(\ell)$ and

- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi'' \Leftrightarrow \rho \in \pi_2$.

(R1-7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma_2''^{\mathbf{U}}, \gamma_{2\,\#}''^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta'')$, $\delta_2(\ell) = \delta''(\ell) \neq \bullet$ and $\rho \in \pi_2$.

(R1-8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma_2''^{\mathbf{A}}, \gamma_{2\,\#}''^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta'')$. If $\rho \neq \ell$, then $\delta_2(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_2}(\gamma_{2\,\#}''^{\mathbf{A}})$, then $\rho \in \pi_2$.

(R1-9) Resources: Let $\Theta'' = \mathrm{reach}_{\delta''}(\gamma_2''^{\mathbf{L}})$. Let $\Theta_2 = \mathrm{reach}_{\delta_2}(\gamma_2''^{\mathbf{L}})$.
For all $\ell \in \Theta''$ it must be that $\Theta''(\ell) = \Theta_2(\ell) = 1$, $\ell \notin \pi_2$, and if $\delta''(\ell)$ is a resource, then $\delta_2(\ell) = \bullet$.

(R1-10) No thin air permission:
$\pi_2 \subseteq \pi'' \cup (\mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta''))$.

It remains to relate to result with the original call to eval.

(R1) $R_2 = \mathrm{Ok}(\delta_2, \pi_2, r_2)$

(R2) $\Delta \leq \Delta_2$, $\delta \leq \delta_2$, $\vdash \delta_2 : \Delta_2$ because $\Delta'' = \Delta$ and (R1-2).

(R3) $\Delta_2 \vdash r_2 : \tau_2$ by (R1-3)

(R4) π_2 is wellformed and $\mathrm{getloc}(\pi_2) \subseteq \mathrm{dom}(\delta_2) \setminus {\delta_2}^{-1}(\bullet)$. Immediate from (R1-4).

(R5) $\mathrm{reach}_0(r_2) \subseteq \pi_2$, $\mathrm{reach}_{\delta_2}(r_2) \subseteq {\downarrow}\pi_2 \cap (\mathrm{reach}_{\delta_2}(\gamma) \setminus \mathrm{reach}_{\delta_2}(\gamma_\#) \cup \mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta))$. By (R1-5) and because $\delta = \delta''$.

(R6) Frame:
For all $\ell \in \mathrm{dom}(\delta) \setminus \mathrm{getloc}(\mathrm{reach}_{\delta_2}(\gamma))$ it must be that
- $\delta_2(\ell) = \delta(\ell)$ and
- for any ρ with $\mathrm{getloc}(\rho) = \{\ell\}$, $\rho \in \pi \Leftrightarrow \rho \in \pi_2$.

Immediate from (R1-6) because $\delta = \delta''$

(R7) Unrestricted values, resources, and borrows:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma^{\mathbf{U}}, \gamma_\#^{\mathbf{U}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$, $\delta_2(\ell) = \delta(\ell) \neq \bullet$ and $\rho \in \pi_2$.

(R8) Affine borrows and resources:
For all $\rho \in \mathrm{reach}_{\delta_2}(\gamma^{\mathbf{A}}, \gamma_\#^{\mathbf{A}})$ with $\mathrm{getloc}(\rho) = \{\ell\}$, it must be that $\ell \in \mathrm{dom}(\delta)$. If $\rho \neq \ell$, then $\delta_2(\ell) \neq \bullet$. If $\rho \in \mathrm{reach}_{\delta_2}(\gamma_\#^{\mathbf{A}})$, then $\rho \in \pi_2$.

(R9) Resources: Let $\Theta = \mathrm{reach}_{\delta}(\gamma^{\mathbf{L}})$. Let $\Theta_2 = \mathrm{reach}_{\delta_2}(\gamma^{\mathbf{L}})$.
For all $\ell \in \Theta$ it must be that $\Theta(\ell) = \Theta_2(\ell) = 1$, $\ell \notin \pi_2$, and if $\delta(\ell)$ is a resource, then $\delta_2(\ell) = \bullet$.
Immediate by (R1-9) because the borrowing match does not deallocate.

(R10) No thin air permission:
$\pi_2 \subseteq \pi \cup (\mathrm{dom}(\delta_2) \setminus \mathrm{dom}(\delta))$.

□

Non-anonymous material for the article "Kindly Bent to Free Us"

An online playground for the type-checker is available at:

`https://drup.github.io/pl-experiments/affe/`

The implementation is available at

`https://github.com/Drup/pl-experiments`