

The implementation of the *Mezzo* type-checker

Jonathan Protzenko

INRIA

jonathan.protzenko@ens-lyon.org

Abstract

We present the implementation of *Mezzo*, a programming language based on the notion of *permission* that provides strong guarantees about aliasing, ownership and mutable state. The unique features of *Mezzo* make its implementation challenging, both from a formal and a technical perspective. In this paper, we formalize the core operations used for type-checking, and give corresponding algorithms. Two central operations, *subtraction* and *merging*, are detailed. Both operations are closely related to shape analysis and frame inference.

1. Introduction

Mezzo is a strict, functional programming language in the spirit of ML [18]. *Mezzo* is based on the concept of permission, which blends together the *type* of an object as well as the *state* the object is in.

The type system of *Mezzo* differs significantly from that of ML. Writing a type-checker for the language thus poses several new challenges. The purpose of this paper is to show how we managed to design a type-checker that could account for the many novel features of the *Mezzo* type system.

One could think of *Mezzo* as separation logic turned into a type system. Our main contribution is a presentation of type-checking that is capable of dealing with the *frame inference* problem (computing the portions of the heap that are left untouched by a call to a sub-routine) and the *join* problem (finding a description of the heap that subsumes two, more precise descriptions), within the context of ML. This implies dealing with arbitrary, higher-order quantifiers; with duplicable and non-duplicable portions of the heap; with subtyping relations. These issues, to the best of our knowledge, have not yet been dealt with in the literature. Before discussing the type-checking of *Mezzo*, we offer a quick introduction to help the reader become acquainted with the core concepts of the language. We use the example of the *map* function, which will be recurring throughout the present paper. We refer the reader to a more exhaustive document [18] for a more thorough introduction to the language¹.

¹The reader familiar with *Mezzo* already should know that in the present document, we omit the syntactic conventions of the surface syntax of *Mezzo*. Therefore, the *consumes* keyword is not mentioned, and neither is the name introduction construct. This means that, in the present paper, any permission found in a function’s argument will be consumed from

1.1 Why *Mezzo*?

Programming with mutable, heap-allocated data structures is hard. Various pitfalls await the unwary programmer: one may unwittingly share a mutable object between concurrent threads, leading to race conditions; etc. Such mistakes are typically not detected statically by mainstream programming languages. We assert that one core reason for these mistakes is a failure to track *ownership* and *aliasing*. *Mezzo* aims to fill this gap. Programs written in *Mezzo* enjoy strong properties, among which that all mutable portions of the heap have a unique owner. This entails data-race freedom.

Another common pattern, in imperative programs, is for objects to follow a protocol. A structure may start with its fields uninitialized, to be filled out later. In C++, one may use the `NULL` pointer; in ML, one may use the `option` type. Both solutions are unsatisfactory, as they fail to check that the only transition possible is from an uninitialized state to an initialized one. By reasoning about *state changes*, *Mezzo*, again, can check such properties statically. Moreover, the finer-grained reasoning also allows for new patterns that were previously unattainable in ML, such as memory re-use and type-changing updates.

We do not wish to rely on an external SMT-solver, or require manual, interactive intervention performed by a program proof expert in order for the program to be verified. In contrast to other tools, *Mezzo* is a type system, whose reasoning is performed automatically, and which requires no manual intervention apart from type annotations.

We have written more than 6000 lines of *Mezzo*, in the form of various libraries and data structures, such as lists (including destination-passing style functions not writable in ML [18]), tree maps (with memory re-use patterns), iterators [14], channels, which our implementation successfully type-checks, using a very low number of annotations.

1.2 Introduction to the type system of *Mezzo* or, permissions

A proof of soundness has been performed [2] for “Core *Mezzo*”, a lower-level version of the language. In this section, we present a higher-level language, which differs from “Core *Mezzo*” mostly in that it uses nominal data constructors such as `Nil` instead of integer tags such as 0.

Figure 1 presents the syntactic category of types (which contains permissions); the syntax of expressions is omitted, as it forms a fairly standard dialect of ML. Types have kinds: permissions (`empty`, `p * q`, `x @ t`) have kind `perm`, regular types (`arrow`, `constructor`, `tuple`) have kind `type`. The only types with kind `term` are program variables `x`, which may appear in types, as our only form of dependency. We informally refer to some of the type-checking rules; these are discussed more formally in §2.

the caller. We also do not mention the adoption/abandon mechanism, and discard mode constraints from the discussion.

$\kappa ::=$	type term perm $\kappa \rightarrow \kappa$	kind
$T, t, P ::=$		type or permission
	X	variable (a, x, \dots)
	$t \rightarrow t$	function type
	(\vec{t})	tuple type
	$A \{ \vec{f} : \vec{t} \}$	structural type
	$T \vec{T}$	n -ary type application
	$\forall (X : \kappa) T$	universal quantification
	$\exists (X : \kappa) T$	existential quantification
	$=x$	singleton type
	$(t \mid P)$	type/permission conjunction
	$x @ t$	atomic permission
	empty	empty permission
	$P * P$	permission conjunction
$D ::=$		algebraic data type definition
	data mutable? $d(\vec{X} : \vec{\kappa}) = \vec{b}$	
	adopts t	
$b ::=$	$A \{ \vec{f} : \vec{t} \}$	algebraic data type branch

Figure 1. Syntax of types and permissions

In order to illustrate our type system, we start by introducing the definition of *immutable* lists.

```
data list a =
| Nil
| Cons { head: a; tail: list a }
```

If e_1 is an expression that creates a list of elements with type a , writing “let $x = e_1$ in e_2 ” generates a new *permission* “ $x @ \text{list } a$ ”, which is available in e_2 . Permissions do not exist at runtime, and are purely an artefact of the type system.

At any program point, a *current permission* is available. When execution starts, the current permission is *empty*. As program execution advances, the current permission evolves, possibly becoming the *conjunction* of several permissions.

We thus replace the common concept of a typing judgement (“ x has type t ”) by that of a permission (“we have permission $x @ t$ ”).

Permissions have an ownership reading: the permission “ $x @ \text{list } a$ ” not only grants its owner the right to use x as a list, but it also embodies the *ownership* of the list. Whoever owns “ $x @ \text{list } a$ ” owns a *list spine*, made up of a succession of Cons cells, as well as the elements themselves. Here, a is a type variable; the ownership reading of the elements depends on the actual instantiation of a .

Permissions may disappear, for instance when calling a function. Here is the type of the map function. Square brackets denote universal quantification.

```
val map: [a, b] ((a -> b), list a) -> list b
```

The map function requires a pair of a function and a list. What happens whenever one wishes to call the map function is described in rule **APPLICATION** (Figure 2). Calling map with (f, l) requires exhibiting the permissions “ $f @ a \rightarrow b$ ” and “ $l @ \text{list } a$ ”. The signature of map indicates that the function does not hand back these permissions: they are *consumed*. The caller will, however, obtain a fresh permission “ $\text{ret} @ \text{list } b$ ”, where *ret* denotes the return value of the function.

In *Mezzo*, permissions are either *duplicable* or *exclusive*. A duplicable permission denotes shared, read-only knowledge, while an exclusive permission denotes uniquely-owned, read-write knowl-

edge. The unique-owner property allows the owner of a mutable object to change its type, which is safe as the system guarantees no one else knows about the object. A third mode, *affine*, exists, which is a superset of the two.

Duplicable permissions can be copied at any time. Function types, for instance, are duplicable. Back to the map example, this means that the caller will not lose the permission “ $f @ a \rightarrow b$ ”, since it can save a copy of it before calling map. What about “ $l @ \text{list } a$ ”?

Determining whether “ $l @ \text{list } a$ ” may be copied or not is non-trivial, and depends on what the type variable a happens to be. If l is a list of duplicable elements, e.g. if we have “ $l @ \text{list int}$ ”, then this permission denotes a read-only fragment of the heap. It is thus duplicable. If, however, l is a list of exclusive elements, e.g. if we have “ $l @ \text{list (ref int)}$ ”, then the permission is *affine*².

1.3 Type-checking by example

We illustrate our type system through the example of the map function. While the function can be written the exact same way in ML, the version in *Mezzo* provides stronger guarantees, via the ownership reading of the function, which we discuss below. The short example allows subsequent sections to pick specific steps of the type-checking to illustrate the discussion; explaining a bigger, more original example would require extra digressions that are outside the scope of the present document.

```
1  val rec map [a, b] (
2      f: (a -> b),
3      l: list a
4  ): list b =
5  match l with
6  | Nil ->
7      l
8  | Cons { head = h; tail = t } ->
9      let h' = f h
10     and t' = map f t in
11     Cons { head = h'; tail = t' }
12  end
```

Type-checking in *Mezzo* is performed in a forward manner similar to symbolic execution. One starts by assuming the pre-condition or, in *Mezzo* lingo, by adding the function’s arguments to the *current permission*, and then moves forward, updating the permission as type-checking progresses through the program. Thus, at line 5, permissions for f , l and map are available (**FUNCTION**):

```
f @ a -> b *
l @ list a *
map @ [a, b] (...) -> list b
```

Let us say a word about functions. We require that functions be annotated with the type of their arguments³ as well as their return type. The type-checker can thus easily detect which permissions are modified or left untouched by a function call. We believe these annotations represent a light burden, and that writing them is good practice. Moreover, functions can only close over duplicable permissions, meaning that a function type simply is always duplicable. A function that closes over a non-duplicable permission p can be written as the conjunction of a function that takes p as an argument and of p itself; the conjunction thus emulates “unique” function types.

² l itself does not denote a read-write portion of the heap, as the list cells are immutable; the permission “ $l @ \text{list } a$ ” is therefore neither duplicable nor exclusive: it is affine.

³ An experimental mechanism allows partly lifting these restrictions.

APPLICATION
 $K; x_1 @ t_2 \rightarrow t_1 * x_2 @ t_2 \vdash x_1 x_2 : t_1$

FUNCTION

$$\frac{K, \vec{X} : \vec{\kappa}, x : \text{term}; P * x @ t_1 \vdash e : t_2 \quad P \text{ is duplicable}}{K; P \vdash \text{fun } [\vec{X} : \vec{\kappa}] (x : t_1) : t_2 = e : \forall(\vec{X} : \vec{\kappa}) t_1 \rightarrow t_2}$$

LET

$$\frac{K; P \vdash e_1 : t_1 \quad K, x : \text{term}; x @ t_1 \vdash e_2 : t_2}{K; P \vdash \text{let } x = e_1 \text{ in } e_2 : t_2}$$

FRAME

$$\frac{K; P_1 \vdash e : t}{K; P_1 * P_2 \vdash e : (t \mid P_2)}$$

Figure 2. Selected typing rules

We now type-check the function and omit, for brevity, the permissions for `f` and `map`. The reader can safely assume that they remain available throughout the whole body of `map`. We thus have:

```
1 @ list a
```

When the type-checker enters the `Nil` branch of the `match` construct (line 6), the permission is refined to:

```
1 @ Nil
```

`Nil` is a structural type that asserts that `1` is a block in the heap whose tag is `Nil`, that has zero fields.

Naming `ret` the return value of the function, we obtain after line 7:

```
PNil = 1 @ Nil * ret = 1
```

This is our first encounter of the *singleton type*. Having the permission “`x @ =y`” means that `x` and `y` are equal; in particular, if they are pointers, they point to the same object in the heap. We write “`x = y`”, as syntactic sugar.

Similarly, entering the `Cons` branch (line 8), the type-checker refines `1 @ list a` into the following permission:

```
1 @ Cons { head: a; tail: list a }
```

This permission is expanded (`DECOMPOSEBLOCK`).

```
1 @ Cons { head = h; tail = t } *
h @ a *
t @ list a
```

Just like the “`ret = 1`” example above, we write “`head = h`” to signify “`head = h`”⁴

The two function calls at lines 9 and 10 consume permissions from the environment (`APPLICATION`, `FRAME`): calling `f` consumes “`h @ a`”, which is *subtracted* from the environment, while a permission for `h'` is returned. Similarly, calling `map` consumes “`t @ list a`” which is traded for a new permission for `t'`. The two operations are performed in sequence, which gives in the end:

```
1 @ Cons { head = h; tail = t } *
h' @ b *
t' @ list b
```

At this stage, we can no longer do anything useful with `h`, since the permission is gone: any reference to it will be type-checked with the top type. The return value, which we call again `ret`, is then created:

```
PCons = 1 @ Cons { head = h; tail = t } *
h' @ b * t' @ list b *
ret @ Cons { head = h'; tail = t' }
```

The type-checker now needs to type-check the `match` construct. The type-checker needs to apply subsumption rules on both

⁴In practice, the type-checker introduces automatic names for the `head` and `tail` fields when expanding the permission, then performs pattern-matching, which results in additional equations that identify the user-provided names with the type-checker names.

`PNil` and `PCons`, until they match. This is called the merge operation, and constitutes the topic of §4. Here, the type-checker sees that `ret` is either `Nil` or `Cons`. It thus decides to reconcile the two permissions for `ret` as “`ret @ list b`”.

One may object that the type-checker is doing unnecessary work. Indeed, the return type of the `map` function constitutes a type annotation that tells us already what we should seek to obtain. We could skip the merge operation (reconciling the two permissions) altogether, and merely assert that `ret @ list b` is a permission that is available in both branches. Even though this is actually what our type-checker does in this particular case, type annotations are not always available. We thus discuss the general procedure.

The merge operation wants “`ret @ list b`”, which implies subtracting the permission “`ret @ list b`” from each branch. This operation consumes permissions. In the first branch, the constructor `Nil` is duplicable: we thus make a fresh copy of “`ret @ Nil`” to obtain “`ret @ list b`”, so that the following permission remains:

```
1 @ Nil * ret = 1
```

In the second branch, obtaining “`ret @ list b`” consumes permission “`h' @ b * t' @ list b`”. Indeed, `b` is a type variable which is, lacking any other information, considered non-duplicable, or affine. Conversely, the permission for `ret` is retained, as it denotes a `Cons` block in the heap which is immutable, hence duplicable. We obtain “`ret @ list b`” and are left with:

```
1 @ Cons { head = h; tail = t } *
ret @ Cons { head = h'; tail = t' }
```

Continuing the *reconciliation*, the type-checker infers that “`f @ a -> b * map @ ...`” is a valid conjunction for both branches, meaning that the final conjunction obtained after the `match` construct is:

```
ret @ list b * f @ a -> b * map @ ...
```

The final step consists in matching the final permission above with the expected return type for the function. This is, again, achieved by performing the *subtraction* of “`ret @ list b`”. The remaining permission is discarded; in this case, the subtraction merely serves as a way to assert that a particular permission is available.

1.4 Challenges for a type-checker

The example of the `map` function highlights several salient points.

- The representation of a permission is not unique; one may “unfold” a structural permission (tuple, constructor) by introducing singleton types; one may also float out quantifiers (`EXISTSSTAR`) or permissions (`MIXSTAR`). We start by introducing *prefixed* permissions, along with a normalized representation (§2) of a permission.
- The discussion mentions in several occasions that the type-checker needs to perform a *subtraction*, an operation that is required for implementing the typing rules from Figure 2. This

implies recombining permissions using rule [SUBSUMPTION](#). Which subsumption rules should be applied? In what order? There is a wide range of options for subsumption. We present an algorithm formulation of subtraction along with an implementation strategy (§3).

- Subtraction allows one to type-check function calls (extracting a permission, obtaining a remainder), function bodies (checking that the return type is satisfied), but we need an additional operation to type-check disjunctions, that is, `match` and `if-then-else` expressions. This is the `merge` operation. We present an algorithm that leverages subtraction (§4).

2. The representation of permissions

Subsumption rules and typing rules are shown in Figure 3 and Figure 2. Our typing judgements are of the form $K; P \vdash e : t$, meaning that under kinding environment K , using permission P , expression e is shown to have type t .

This formulation gives no clue as to how one can effectively type-check a *Mezzo* program. Type-checking a function call requires a combination of [APPLICATION](#) and [FRAME](#): how does one compute the part that is framed out? Similarly, the [LET](#) rule mentions $x @ t_1$, but t_1 will most certainly be of the form $(t'_1 \mid Q)$ where Q is a permission big enough to type-check e_2 with type t_2 . How does one know the value of t_2 or Q ?

2.1 Why is type-checking difficult?

We need to adopt a presentation where we carry a current permission P , and where a type-checking step takes P and returns P' , as this will form the basis of an algorithm. Finding the right representation for P is difficult, though. Indeed, a “good” representation shall satisfy quite a few constraints, which we review below.

- The existence of possibly several permissions for the same variable requires *exploration* of a solution space.
- Our conjunction operator $*$, unlike that of separation logic, works on non-exclusively owned portions of the heap. Indeed, some permissions can be duplicated, while some others cannot. When subtracting a permission, whether it will be consumed and disappear depends on whether it is duplicable or not.
- We mention in §1.4 the various, equivalent ways of representing a permission. Function types can also be represented in various equivalent ways: for instance,

$$x @ \forall a, \forall (y : \text{term}). (=y \mid y @ a) \rightarrow () \equiv x @ \forall a, a \rightarrow ()$$

Indeed, first-class functions are central in ML-like languages. We thus need to compare function types during the subtraction, something that other works do not need to do. We cannot use a syntactic criterion, as we often want powerful reasoning when comparing function types (for instance, implicit η -expansion).

- We manipulate quantifiers, both universal and existential. They may stem from user-provided quantifiers (polymorphic function types, one-shot functions), or be introduced by the translation of the name introduction construct found in the surface syntax [18]. They may appear in arbitrary order, which requires careful treatment to ensure they are properly introduced. Moreover, quantifiers may refer to variables that have kind term (e.g. x) type (e.g. t) or perm (e.g. $x @ t$). A quantifier at kind term refers to a program variable; this is akin to first-order logic: we’re quantifying over the program variable x that may appear in $x @ t$. A quantifier at kind type or perm is different; such a quantifier may refer to an entire permission $x @ t$, not just x . This would be the equivalent of second-order logic.

Per the first item, we know that we are bound to use a conjunction; a map from variables to types would not work. The second item encourages us to use an expanded representation. Indeed,

$$l @ \text{Cons} \{ \text{head} = h; \text{tail} = t \} * h @ a * t @ \text{list } a$$

is more precise than:

$$l @ \text{Cons} \{ \text{head} : a; \text{tail} : \text{list } a \}$$

The former conjunction tells us that the permission for l is duplicable, while the latter is more conservative and tells us that the permission for l is affine. The last two items highlight the difficulty of dealing with singleton types and quantifiers.

2.2 Prefixes

We quantify over the free variables of a permission using a *prefix* \mathcal{V} . A prefix is a succession of universal quantifiers $\forall(x : \kappa)$, existential quantifiers $\exists(x : \kappa)$ and instantiated existential quantifiers $\exists(x : \kappa = \tau)$. If κ is missing, the kind is assumed to be `type`.

A prefix \mathcal{V}' is more precise than \mathcal{V} if \mathcal{V}' can be obtained from \mathcal{V} by: *instantiating* existential quantifiers, or *inserting* existential quantifiers, or *appending* arbitrary quantifiers.

Type-checking initially starts with the permission empty, then threads it through various program points. Permissions are added (when entering function bodies, or when a function call returns) and removed (when exiting function bodies, or performing a function call). Thus, the type-checker carries a prefixed permission $\mathcal{V}.P$.

We define $\text{restrict}(\mathcal{V}', \mathcal{V})$ to be the smallest prefix of \mathcal{V}' that contains all the variables quantified in \mathcal{V} . We see its usage in §3.3.

2.3 Normal form for a permission

We maintain the following invariants.

Existential quantifiers Existential quantifiers are all hoisted into the prefix.

Expanded form All structural permissions (tuples, constructors) are *expanded*, meaning that their fields are all singleton types. All function types have a domain that is a singleton type; structural permissions in the domain of function types are also expanded.

No hidden permissions All permissions nested inside structural types using $(t \mid P)$ are floated out into the outer conjunction.

No redundant permissions Some conjunctions, such as $x @ \text{Nil} * x @ \text{list } a$, are redundant. We simplify these.

Let us now explain how these invariants are enforced.

Existential quantifiers The existence of [EXISTSATOMIC](#) allows us to move existential quantifiers into the prefix. Along with rules [DECOMPOSEBLOCK](#) and [DECOMPOSETUPLE](#), any existentially-quantified type that was found “inside” a tuple or a constructor will be assigned to a fresh name, thus making [EXISTSATOMIC](#) applicable. The reason why existential quantifiers must be hoisted out is the topic of §3.5.

Expanded form Combining [EQUALITYREFLEXIVE](#), [CoSTAR](#), [MIXSTAR](#) and [EXISTSINTRO](#) allows one to derive the following subtyping rule, which we call [DECOMPOSE](#):

$$\text{DECOMPOSE} \\ t \equiv \exists(x : \text{term}). (=x \mid x @ t)$$

Intuitively, one can always give a name to the portion of the heap denoted by type t . Rules [DECOMPOSEBLOCK](#) and [DECOMPOSETUPLE](#) allow us to enforce the expanded form invariant on structural types, while [DECOMPOSE](#) allows the invariant to be respected in the domain of functions, in combination with [CoARROW](#).

This invariant is important for two reasons. First, having it simplifies the application of rules [TUPLE](#) and [CONSTRUCTOR](#), as t is

$\frac{\text{DUPLICATE}}{P \text{ is duplicable}} \frac{}{P \leq P * P}$	$\frac{\text{COSTAR}}{P_1 \leq P_2 \quad Q_1 \leq Q_2} \frac{}{P_1 * Q_1 \leq P_2 * Q_2}$	$\frac{\text{EQUALSFOR EQUALS}}{(y_1 = y_2) * [y_1/x]P} \equiv (y_1 = y_2) * [y_2/x]P$	$\frac{\text{COPYDUP}}{u \text{ is duplicable}} \frac{}{x @ t * y @ u \leq x @ [u/=y]t}$
$\frac{\text{DECOMPOSEBLOCK}}{\equiv \exists(x : \text{term}) (y @ A \{F[f : t]\})} \frac{}{\equiv \exists(x : \text{term}) (y @ A \{F[f = x]\}) * x @ t}$	$\frac{\text{DECOMPOSETUPLE}}{\equiv \exists(x : \text{term}) (y @ (\dots, t, \dots))} \frac{}{\equiv \exists(x : \text{term}) (y @ (\dots, =x, \dots)) * x @ t}$	$\frac{\text{EXISTSATOMIC}}{x @ \exists(X : \kappa) t \equiv \exists(X : \kappa) (x @ t)}$	
$\frac{\text{EXISTSTAR}}{P_1 * \exists(X : \kappa) P_2 \equiv \exists(X : \kappa) (P_1 * P_2)}$	$\frac{\text{EQUALITYREFLEXIVE}}{\text{empty} \leq (x = x)}$	$\frac{\text{MIXSTAR}}{x @ t * P \equiv x @ (t P)}$	$\frac{\text{FOLD}}{A \{\vec{f} : \vec{t}\} \text{ is an unfolding of } T \vec{T}} \frac{}{x @ A \{\vec{f} : \vec{t}\} \leq x @ T \vec{T}}$
$\frac{\text{COARROW}}{u_1 \leq t_1 \quad t_2 \leq u_2} \frac{}{x @ t_1 \rightarrow t_2 \leq x @ u_1 \rightarrow u_2}$	$\frac{\text{COMMUTEARROW}}{x @ (\exists(\alpha : \kappa).t) \rightarrow u \equiv x @ \forall(\alpha : \kappa).(t \rightarrow u)}$		$\frac{\text{SUB}}{\forall(x : \text{term}).x @ t - x @ u} \frac{}{t \leq u}$

Figure 3. Selected permission subsumption rules

already of the form $=x$, meaning that we don't have to introduce an extra variable y . Second, it reveals “implicit” existential quantifiers which are then moved into the prefix. Again, the technical discussion on why it is important to eagerly move existentials into the prefix takes place in §3.5.

No hidden permissions Decomposition, along with **MIXSTAR**, allows enforcing this invariant. Having this invariant simplifies the reasoning all throughout the present paper, as we don't have to consider the case where the permission we're looking for is “stashed” inside another one.

No redundant permissions The algorithm applies a last set of rules to further simplify P . The following conjunction is not inconsistent:

$$x @ \text{list } a * x @ \text{Cons } \{\text{head} = h; \text{tail} = t\}$$

It can be simplified into:

$$x @ \text{Cons } \{\text{head} = h; \text{tail} = t\} * h @ a * t @ \text{list } a$$

Another is example is

$$x @ (=y_1, =y_2) * x @ (=z_1, =z_2)$$

which is simplified into

$$x @ (=y_1, =y_2) * y_1 = z_1 * y_2 = z_2$$

A similar rule exists for constructors.

All the subsumption rules used for normalization are *reversible*, meaning that we do not lose information when applying them. Applying a non-reversible rule (e.g. **FOLD**) is a bad idea, as it will *lose* information about the current permission.

Once the normalization procedure is in place, one can implement the *addition* of a permission, which boils down to applying the normalization rules above. Additions take place when entering function bodies, and when returning from function calls.

2.4 Data structures of the type-checker

The type checker of *Mezzo* is written in OCaml. Representing a current permission P is done using the following data structures.

Persistent union-find We use a union-find structure to keep track of equalities between variables, that is, permissions of the form $x = y$. The union-find is persistent, as it facilitates backtracking and branching on disjunctions. Moreover, the union-find maps the representative of an equivalence class x to a list of types, which represent the permissions available for x .

Floating permissions Permissions that are not of the form $x @ t$ are not attached to a program variable; they are abstract. We dub them “floating permissions”, and store them in a separate list.

Flexible variables Flexible variables (introduced later on, in §3.3) are implemented using unification variables. A flexible variable is a globally unique atom. We keep a persistent map of these atoms to a structure containing the variable's level as well as its instantiation, if any. This structure allows inserting existential quantifiers at arbitrary position (**FLEX-TUPLE-L**).

Levels In order to efficiently check the non-occurrence premise of a rule such as **INSTANTIATION**, we use levels [19]. This prevents illegal instantiations, and allows for an efficient implementation of the “is local” check during merges (§4). The $\text{restrict}(\mathcal{V}', \mathcal{V})$ operation is implemented by restricting the flexible variables of \mathcal{V}' to the maximum level of \mathcal{V} .

3. The subtraction operation

The basic building block of a type-checking algorithm is the *subtraction* operation, as mentioned at numerous occasions in §1.2. A subtraction extracts Q from P , computing the remainder R . This operation is used when exiting function bodies, to check that the return type is satisfied (the remainder R is then discarded), and when calling functions, to compute the part of the current permission that is *consumed* by the function call, and the remainder (the frame). This problem is known as *frame inference* in separation logic.

3.1 Subtraction examples

Subtractions can be found in the map example. At line 9, for instance, a call to **f** takes place. Rule **APPLICATION** requires a permission for **f** that is a function type; we happen to possess one, namely “**f @ a -> b**”. Consistently with §1.2, we omit this permission, for brevity. The function is called with argument **h**. We thus need to obtain “**h @ a**” from the current permission; in other words, we need to perform the following subtraction:

$$\begin{array}{l} l @ \text{Cons } \{ \text{head} = h; \text{tail} = t \} * \\ h @ a * t @ \text{list } a \\ - h @ a \end{array}$$

Will “**h @ a**” disappear? As the type variable **a** is abstract, we have to be conservative: **a** could be anything, and denote uniquely-owned data. We thus take a conservative stance, and assume “**h @ a**” is not a duplicable permission. It is affine, and we cannot save a copy of it. Thus, the subtraction *consumes* the permission, leading to the following remainder:

$\frac{\text{SUBSUMPTION} \quad P \leq P' \quad Q' \leq Q \quad \mathcal{V}.P' - Q' = \mathcal{V}'.R}{\mathcal{V}.P - Q = \mathcal{V}'.R}$	$\frac{\text{EMPTY} \quad \mathcal{V}.P - \text{empty} = \mathcal{V}.P}{\mathcal{V}.P - \text{empty} = \mathcal{V}.P}$	$\frac{\text{STAR} \quad \mathcal{V}.P - Q = \mathcal{V}'.R \quad \mathcal{V}'.R - Q' = \mathcal{V}''.R'}{\mathcal{V}.P - Q * Q' = \mathcal{V}''.R'}$
$\frac{\text{ARROW} \quad \mathcal{V}, \forall(y : \text{term}). \bar{P} * y @ t'_1 - y @ t_1 = \mathcal{V}'.P' \quad \mathcal{V}', \forall(z : \text{term}). P' * z @ t_2 - z @ t'_2 = \mathcal{V}''.P''}{\mathcal{V}.P * x @ t_1 \rightarrow t_2 - x @ t'_1 \rightarrow t'_2 = \text{restrict}(\mathcal{V}'', \mathcal{V}).P}$	$\frac{\text{APP} \quad \mathcal{V}_i, \forall(y : \text{term}). \bar{P} * y @ u_i \stackrel{v_i}{\dashv} y @ u'_i = \mathcal{V}'.P' \quad \mathcal{V}_{i+1} = \text{restrict}(\mathcal{V}', \mathcal{V}_i)}{\mathcal{V}_0.P * x @ t \bar{u} - x @ t \bar{u}' = \mathcal{V}_n.P}$	$\frac{\text{FORALL-L} \quad \mathcal{V}, \exists(\alpha : \kappa). P - Q = \mathcal{V}'.R}{\mathcal{V}.(\forall(\alpha : \kappa). P) - Q = \mathcal{V}'.R}$
$\frac{\text{FORALL-R} \quad \mathcal{V}, \forall(\alpha : \kappa). P - Q = \mathcal{V}'.R}{\mathcal{V}.P - (\forall(\alpha : \kappa). Q) = \mathcal{V}'.R}$	$\frac{\text{EXISTS-L} \quad \mathcal{V}, \forall(\alpha : \kappa). P - Q = \mathcal{V}'.R}{\mathcal{V}.(\exists(\alpha : \kappa). P) - Q = \mathcal{V}'.R}$	$\frac{\text{EXISTS-R} \quad \mathcal{V}, \exists(\alpha : \kappa). P - Q = \mathcal{V}'.R}{\mathcal{V}.P - (\exists(\alpha : \kappa). Q) = \mathcal{V}'.R}$
$\frac{\text{VARIABLEPERM} \quad \mathcal{V}.P * p - p = \mathcal{V}.P}{\mathcal{V}.P * p - p = \mathcal{V}.P}$	$\frac{\text{VARIABLE} \quad \mathcal{V}.P * x @ \alpha - x @ \alpha = \mathcal{V}.P}{\mathcal{V}.P * x @ \alpha - x @ \alpha = \mathcal{V}.P}$	$\frac{\text{TUPLE} \quad \mathcal{V}_i, \forall(y : \text{term}). P_i * y @ t_i - y @ t'_i = \mathcal{V}_{i+1}.P_{i+1}}{\mathcal{V}_0.P_0 * x @ (\vec{t}) - x @ (\vec{t}') = \mathcal{V}_n.P_n}$
$\frac{\text{CONSTRUCTOR} \quad \mathcal{V}_i, \forall(y : \text{term}). P_i * y @ t_i - y @ t'_i = \mathcal{V}_{i+1}.P_{i+1}}{\mathcal{V}_0.P_0 * x @ \mathbf{A} \{\vec{f} : \vec{t}\} - x @ \mathbf{A} \{\vec{f}' : \vec{t}'\} = \mathcal{V}_n.P_n}$	$\frac{\text{SUBSTITUTEFLEXIBLE} \quad \mathcal{V}, \exists(\alpha : \kappa = \tau), \mathcal{V}'.[\tau/\alpha]P - [\tau/\alpha]Q = \mathcal{V}''.R}{\mathcal{V}, \exists(\alpha : \kappa = \tau), \mathcal{V}'.P - Q = \mathcal{V}''.R}$	$\frac{\text{UNKNOWN} \quad \mathcal{V}.P - x @ \text{unknown} = \mathcal{V}'.P}{\mathcal{V}.P - x @ \text{unknown} = \mathcal{V}'.P}$
$\frac{\text{INSTANTIATION} \quad \mathcal{V}, \exists(\alpha : \kappa = \tau), \mathcal{V}'.P - Q = \mathcal{V}''.R \quad \mathcal{V}' \# \tau}{\mathcal{V}, \exists(\alpha : \kappa), \mathcal{V}'.P - Q = \mathcal{V}''.R}$	$\frac{\text{EXISTSINTRO} \quad \mathcal{V}, \exists(\alpha : \kappa), \mathcal{V}'.P - Q = \mathcal{V}''.R}{\mathcal{V}, \mathcal{V}'.P - Q = \mathcal{V}''.R}$	

Figure 4. The rules of subtraction

$\frac{\text{VARIANCE-CO} \quad \mathcal{V}.P * x @ t - x @ t' = \mathcal{V}'.P'}{\mathcal{V}.P * x @ t \stackrel{\text{co}}{\dashv} x @ t' = \mathcal{V}'.P'}$	$\frac{\text{VARIANCE-CONTRA} \quad \mathcal{V}.P * x @ t' - x @ t = \mathcal{V}'.P'}{\mathcal{V}.P * x @ t \stackrel{\text{contra}}{\dashv} x @ t' = \mathcal{V}'.P'}$	$\frac{\text{VARIANCE-INV} \quad \mathcal{V}.P * x @ t - x @ t' = \mathcal{V}'.P' \quad \mathcal{V}'.P' * x @ t' - x @ t = \mathcal{V}''.P''}{\mathcal{V}.P * x @ t \stackrel{\text{inv}}{\dashv} x @ t' = \mathcal{V}''.P''}$
		$\frac{\text{VARIANCE-BI} \quad \mathcal{V}.P * x @ t \stackrel{\text{bi}}{\dashv} x @ t' = \mathcal{V}'.P'}{\mathcal{V}.P * x @ t \stackrel{\text{bi}}{\dashv} x @ t' = \mathcal{V}'.P'}$

Figure 5. Variance-dependent subtraction

```
l @ Cons { head = h; tail = t } * t @ list a
```

(This subtraction is followed by the *addition* of “h’ @ b” to the current permission, as per the return type of the function f. Addition is a much easier operation which we briefly mentioned in §2.3.)

A more sophisticated subtraction takes place when type-checking the whole match construct. The type-checker decides (and the reason for this is the topic of §4) to extract “ret @ list b” from both branches. Thus, the following subtraction takes place in the Cons branch:

```
l @ Cons { head = h; tail = t } *
ret @ Cons { head = h'; tail = t' }
h' @ b * t' @ list b *
- ret @ list b
```

The type-checker uses a syntactic criterion and figures out that if we could show that ret is a well-formed Cons cell, this would be enough to justify that it is also a list (FOLD). The type-checker hence tries to perform the following subtraction:

```
l @ Cons { head = h; tail = t } *
ret @ Cons { head = h'; tail = t' }
h' @ b * t' @ list b *
- ret @ Cons { head: b; tail: list b }
```

Applying **DECOMPOSEBLOCK** and floating up existential quantifiers through **EXISTSSTAR**, this amounts to showing that:

```
∃h'', ∃t'' (
l @ Cons { head = h; tail = t } *
```

```
ret @ Cons { head = h'; tail = t' }
h' @ b * t' @ list b *
- ret @ Cons { head = h''; tail = t'' } *
h'' @ b * t'' @ list b
)
```

One can instantiate h'' and t'' with h' and t' respectively. The subtraction then becomes straightforward. Again, the type variable b being abstract, the permission “h’ @ b * t’ @ list b” disappears. Conversely, “ret @ Cons { head = h’; tail = t’ }” is duplicable⁵: a copy of the permission can be saved through **DUPLICATE**. The final remainder of the subtraction is thus:

```
l @ Cons { head = h; tail = t } *
ret @ Cons { head = h'; tail = t' }
```

3.2 Formally defining subtraction

The result of a subtraction of Q from P under V is a remainder R along with a more precise prefix V'. We write:

$$\mathcal{V}.(P - Q) = \mathcal{V}'.R$$

where P and Q may themselves be prefixed. This is analogous to the frame inference problem, for which several algorithms have been proposed [4, 11, 17]. Calcagno *et al.* write this as $\delta \vdash H * ?\text{frame}$ [8] when assuming an existing procedure for frame inference.

⁵ The permission denotes an immutable block, whose fields are singleton types, which are themselves duplicable.

In the algorithmic presentation of the subtraction, we often wish to focus the subtraction on a particular variable x . We write:

$$\mathcal{V}.(P * x @ t - x @ t') = \mathcal{V}'.R$$

We sometimes need to keep only the duplicable parts of a permission. If $P = p_1 * \dots * p_n$, this merely consists in dropping all permissions p_i that do not satisfy the “is duplicable” predicate. We write \bar{P} .

3.3 Flexible variables vs. rigid variables

In a subtraction $\mathcal{V}.(P - Q)$, one can think of P as our current hypothesis, and of \mathcal{V} and Q as our goal. Thus, Q is, like V , in positive position, while P is in negative position. Consider:

$$\forall(x : \text{term}).(x @ \exists a.a - x @ \exists b.b)$$

In the example above, the quantifier $\exists a$ is in negative position, meaning that it turns into a universal when moved into the outer prefix. The $\exists b$ quantifier, being in positive position, moves unchanged into the outer prefix. We obtain:

$$\forall(x : \text{term}), \forall a, \exists b.(x @ a - x @ b)$$

We say that a is a *rigid* variable, while b is an *uninstantiated flexible variable*. This subtraction problem admits the following solution, where b is *instantiated*:

$$\forall(x : \text{term}), \forall a, \exists(b = a).\text{empty}$$

The treatment of flexible and rigid variables is a central issue and constitutes the topic of §3.5.

3.4 An algorithmic presentation of type-checking

The subtraction operation (Figure 4) provides a set of rules that an algorithm can choose to apply; the rules make it explicit which part of a permission is preserved, and which permission is returned. The subtraction rules work at a lower level than the type-checking rules; yet, they leave plenty of non-determinism, meaning that an actual algorithm needs an implementation strategy in order to know which rules should be applied (§3.5).

Rule **SUBSUMPTION** is central, as it allows powerful recombinations of a permission on both sides of the subtraction (§3.6, §2.3).

Rules **EMPTY** and **STAR** are standard. The $*$ conjunction being commutative and associative, rule **STAR** leaves the order of subtractions unspecified. The algorithm is free to apply an efficient strategy (§3.6).

Rules **UNKNOWN** and **VARIABLE** are standard. There is no subtraction rule for singleton types as it can be derived from **EQUALITYREFLEXIVE** and **SUBSUMPTION**.

Rules **FORALL-L**, **FORALL-R**, **EXISTS-L** and **EXISTS-R** deal with quantifiers; they are discussed in §3.5. The right-hand side of a subtraction preserves the nature of the quantifier (it refers, like the prefix, to the *goal* that we’re trying to prove); the left-hand side, however, flips quantifiers (it is our current *hypothesis*).

Rules **TUPLE** and **CONSTRUCTOR** descend into structural permissions. The remainders are *chained*, ensuring that a permission consumed in the i -th subtraction is no longer available in subsequent subtractions.

Rule **APP** is somehow complex; let us review it. First of all, we compare the parameters of the type applications according to the standard notion of *variance*. A type can be either *covariant* or *contravariant* in its i -th parameter. A type that is neither is *invariant*. A type that is both is *bivariant*, meaning it makes no use of its parameter. The $\overset{v_i}{\rightleftharpoons}$ symbol allows the rule to pick the appropriate operation from Figure 5 according to v_i , the variance of the i -th parameter.

APP takes the duplicable restriction of the current permission P before comparing the arguments of the type application. This is

crucial, as performing:

$$\mathcal{V}.y @ \text{ref int} * x @ \text{list} (=y) - x @ \text{list} (\text{ref int})$$

should be forbidden. The list may have length two, meaning we would need *two* copies of $y @ \text{ref int}$, which is not duplicable!

As a first approximation, let us then write the premise as:

$$\mathcal{V}.\bar{P} * y @ u_i \overset{v_i}{\rightleftharpoons} y @ u'_i = \mathcal{V}'.P'$$

The resulting P' should be discarded. Indeed, performing

$$\mathcal{V}.x @ \text{list} (a | Q) - x @ \text{list} a$$

should *not* give Q as a remainder. Indeed, the type $(a | Q)$ may not be available, as x may be Nil; therefore, Q may be not available at all!

This is still not precise enough. If α is a flexible variable, our approximation allows the following subtraction to succeed.

$$\forall(x : \text{term}), \exists \alpha.x @ t \alpha \alpha - x @ t \text{int} ()$$

Indeed, the approximation does not carry the flexible variable instantiations from one premise to another. This is the reason why we need to number our prefixes, and use **restrict** to import the flexible variable instantiations from one prefix onto the next.

ARROW is perhaps the most powerful rule. It first starts by taking the restriction of P to its duplicable bits. A closure in *Mezzo* can only capture *duplicable* variables (function types are duplicable); the corresponding function comparison rule behaves the same way. It then subtracts the domains; this is a contravariant position, meaning the order of the subtraction is changed. Then, it uses the remainder to compare codomains. Any information contained in t'_1 is then carried over, which amounts to performing an η -expansion. The resulting environment is then discarded, and the flexible variable instantiations are retained.

3.5 Implementation: introducing quantifiers

The order in which we introduce quantifiers in a prefix \mathcal{V} is of particular importance. Consider, for instance, the subtraction from §3.3. We luckily chose the right order for introducing quantifiers: the $\forall a.\exists b$ prefix allows picking $b = a$ via **INSTANTIATE**, which solves the subtraction. If we were, however, to introduce b first, the prefix would be $\exists b.\forall a$, meaning that the instantiation could not succeed. Indeed, the two quantifiers cannot commute, meaning that the freshness premise from **INSTANTIATE** is not satisfied.

Quantifier introduction via normalization We thus need to *eagerly introduce universal (rigid) quantifiers*. This is the *raison d’être* of the normalization. By hoisting all possible existential quantifiers out of our current hypothesis P , we are able to introduce them as universal variables in the prefix \mathcal{V} (§2.3).

One may wonder why normalization introduces singleton types in the domain of arrows. Let us consider the following subtraction problem, which we already mentioned in §2.1. This subtraction should succeed (we omit the term kind annotation, for clarity); however, performing it naïvely leads to a failure.

$$\frac{\text{no proof}}{\forall f, \exists x, \forall y.(y = x * y @ a - y @ a)} \text{FUNCTION} \\ \frac{\forall f, \exists x.(f @ (=x | x @ a) \rightarrow () - f @ a \rightarrow ())}{\forall f.(f @ \forall x.(=x | x @ a) \rightarrow () - f @ a \rightarrow ())} \text{FORALL-L}$$

Indeed, finding an x such that for any y , we have $x = y$ is impossible, however powerful the type-checker may be. The y quantifier was introduced *too late*. We have to perform extra work in order to introduce y earlier; this can be achieved by introducing a singleton type in the arrow’s domain. The type-checker properly takes care of it, using a combination of **DECOMPOSE**, **COARROW** and

$$\begin{array}{c}
\text{FLEX-TUPLE-L} \\
\frac{\mathcal{V}, \exists(\vec{\beta} : \vec{\kappa}) \exists(\alpha : \kappa = (\vec{\beta})), \mathcal{V}'. P * x @ (\vec{t}) - x @ \alpha = \mathcal{V}'' . R}{\mathcal{V}, \exists(\alpha : \kappa), \mathcal{V}'. P * x @ (\vec{t}) - x @ \alpha = \mathcal{V}'' . R} \\
\\
\text{FLEX-CONSTRUCTOR-L} \\
\frac{\mathcal{V}, \exists(\vec{\beta} : \vec{\kappa}) \exists(\alpha : \kappa = \mathbf{A} \{ \vec{f} : \vec{\beta} \}), \mathcal{V}'. P * x @ \mathbf{A} \{ \vec{f} : \vec{t} \} - x @ \alpha = \mathcal{V}'' . R}{\mathcal{V}, \exists(\alpha : \kappa), \mathcal{V}'. P * x @ \mathbf{A} \{ \vec{f} : \vec{t} \} - x @ \alpha = \mathcal{V}'' . R} \\
\\
\text{FLEX-DEFAULT} \\
\frac{\mathcal{V}, \exists(\alpha : \kappa = t), \mathcal{V}'. P * x @ t - x @ \alpha = \mathcal{V}'' . R}{\mathcal{V}, \exists(\alpha : \kappa), \mathcal{V}'. P * x @ t - x @ \alpha = \mathcal{V}'' . R}
\end{array}$$

Figure 6. Instantiation of flexible variables

$$\begin{array}{c}
\text{SUBSUMPTION} \\
\frac{\frac{\vdots}{\forall f. (f @ \forall x. (=x | x @ a) \rightarrow ()) - f @ \forall y. (=y | y @ a) \rightarrow ())}{\forall f. (f @ \forall x. (=x | x @ a) \rightarrow ()) - f @ a \rightarrow ()}}{\frac{\frac{\text{DECOMPOSE}}{a \equiv \exists x. (=x | x @ a)} \quad \frac{\text{COARROW}}{a \rightarrow () \equiv (\exists x. (=x | x @ a)) \rightarrow ()} \quad \frac{\text{COMMUTEARROW}}{a \rightarrow () \equiv \forall x. (=x | x @ a) \rightarrow ()}}{\forall f. (f @ \forall x. (=x | x @ a) \rightarrow ()) - f @ a \rightarrow ()}}
\end{array}$$

Figure 7. Singleton types need to be introduced in the domain of arrows

COMMUTEARROW. The example above can thus be type-checked, as shown in Figure 7.

Quantifier introduction using rules Normalization is not enough, because it only extrudes *existential* quantifiers. Rule **EXISTS-R** (resp. **FORALL-L**) introduces a flexible variable. At this stage, we should introduce all possible rigid variables. Failing to do that would lose instantiation options for the flexible variable. The type-checker thus aggressively searches for rigid variables, that is, existential quantifiers on the left-hand side (resp. universal quantifiers on the right-hand side), before appending an existential variable to the prefix.

Exploration The type-checker sometimes need to explore different branches; one such case is when confronted with a choice between applying **FORALL-L** and **EXISTS-R**. Examples exist where only one branch succeeds; these examples are rare, however, meaning that the type-checker hardly ever backtracks (we only have one hand-crafted example in our test suite).

3.6 Implementation: picking rules to apply

When confronted with a subtraction $P - Q$, the type-checker breaks down Q into a conjunction of atomic permissions of the form $x @ t$ or p , where p is a type variable at kind perm. The latter can be easily subtracted (or fail to be) through **VARIABLEPERM**. The former pose more difficulties. First of all, the type-checker is unable to subtract $x @ t$ if x is a flexible variable, meaning that a situation where all permissions have a flexible left-hand side is a failure. Assuming a certain $x @ t$ exists, where x is rigid, the type-checker is faced with a choice. Indeed, there may exist several permissions for x : for instance, $x @ =x$ is always available, and it may happen that multiple function types are available for the same variable x .

The type-checker will thus explore all possible solutions, for each permission $x @ t'$ present in P . This may seem expensive, but there is usually one “useful” permission for a given variable, which we try first, using a syntactic criterion; situations where multiple “useful” permissions are available happen, but pertain to the realm of guru *Mezzo* code.

Once a permission is focused and the subtraction is of the form $P' * x @ t' - x @ t$, the type-checker relies on syntactic criteria to pick which rule to apply. Rules such as **TUPLE**, **CONSTRUCTOR**, **APP** are triggered according to the shape of t . Some situations require the use of **SUBSUMPTION**. For instance, when confronted with a nominal type (e.g. $\text{list } a$) and a concrete type (e.g. Cons), the rule

is used in conjunction with **FOLD**, as follows:

$$\text{FOLD} \frac{\frac{\text{Nil} \leq \text{list } a \quad \frac{\forall (y : \text{term}). y @ \text{Nil} - y @ \text{Nil}}{\forall (y : \text{term}). y @ \text{Nil} - y @ \text{list } a}}{\text{CONSTRUCTOR}} \quad \text{SUBSUMPTION}$$

The **SUBSUMPTION** rule also allows one to save a copy of a permission at any time, using **DUPLICATE**. In practice, and if possible, we save a copy of the permission when focusing on it.

3.7 Implementation: instantiation choices

The way **INSTANTIATE** is formulated leaves plenty of choice for the instantiation of a flexible variable. An implementation should thus have a strategy for instantiating flexible variables. The solution space is very large, and our algorithm explores only a subset of it: it is in no way complete. The behavior of our algorithm is governed by the rules found in Figure 6.

Rules **FLEX-TUPLE-L** and **FLEX-CONSTRUCTOR-L** express the fact that a flexible variable should never directly instantiate to a tuple type or a constructor type. This ensures that a subtraction such as:

$$\exists \alpha. x @ \text{int} * y @ \text{int} * (=x, =y) - \alpha$$

always contains $\alpha = (\text{int}, \text{int})$ as a solution. Lacking these two rules, the only solution considered by the algorithm would be $\alpha = (=x, =y)$ which is often problematic (§4). These rules are prioritized over **FLEX-DEFAULT**.

There are a variety of other situations where the algorithm applies a specific strategy for instantiating flexible variables. We do not detail these, as the discussion is quite technical, and the strategy likely to change in the future.

4. The merge operation

When introducing the language informally (§1.3), we mentioned that in order to type-check the `match` construct, the type-checker is confronted with the following permission, from the `Nil` branch:

```
1 @ Nil * ret = 1
```

and the following permission, from the `Cons` branch:

```
1 @ Cons { head = h; tail = t } *
h' @ b * t' @ list b *
ret @ Cons { head = h'; tail = t' }
```

We went on explaining that the type-checker “magically” knew that it should seek to obtain “`ret @ list b`”, a permission that

is valid in both branches. Solving this problem in general is called the *merge operation*.

The merge operation is what allows one to effectively type-check a *match* or an *if-then-else* expression. As *Mezzo* positions itself as a type-checker, not a static analyzer or a program prover, it is compulsory that disjunctions in control-flow be solved on the spot. Indeed, a type system shall be predictable; having errors that mention a particular path in the control-flow is no intuitive behavior, and postponing errors until the end of a function body is not acceptable either.

4.1 Why is merging difficult?

The merge problem does not admit, in general, a principal solution. Let us see a few examples.

First example We consider a first example where both branches of an *if-then-else* expression return a tuple. Please note that τ is exclusive, meaning only one copy of “ $x @ T$ ” is available.

```
data mutable t = T

val z =
  if ... then begin
    let x = T in
      (x, x)
  end else
    (T, T)
```

We obtain “ $\text{ret } @ (=x, =x) * x @ T$ ” in the first branch and “ $\text{ret } @ (T, T)$ ” in the second branch. This merge problem can be solved by either “ $z @ (T, \text{unknown})$ ” or “ $z @ (\text{unknown}, T)$ ”, where *unknown* denotes the \top type, i.e. the supertype of all types. Lacking any other annotation, the type-checker has no way of deciding which solution should be preferred over the other, as they are incomparable. It is to be remarked that if the type τ were to be duplicable, “ $z @ (T, T)$ ” would be a principal solution.

Second example Here is a slightly more complex example, where sharing occurs:

```
val z =
  let x = T in
    if ... then
      (x, x)
    else begin
      let y = T in
        (y, y)
      end
```

This time, the type-checker can pick several solutions (x has gone out of scope and becomes existentially-quantified):

- picking “ $x @ T * z @ (=u, =u)$ ” will preserve the permission for x while still retaining the fact the two components of the tuple are pointers to the same block in the heap (u is existentially quantified as well);
- picking “ $x @ \text{unknown} * z @ (=u, =u) * u @ T$ ” loses the permission for x in favor of a “better” one for z ;
- other, inferior solutions exist, such as “ $z @ (T, \text{unknown})$ ” or “ $z @ (\text{unknown}, T)$ ”, which would fail to preserve sharing information.

The first two solutions are incomparable, as none of them implies the other. The last two solutions, however, are derivable from the second one. Again, if τ were to be duplicable, this situation would admit a principal solution, namely “ $x @ T * z @ (=u, =u) * u @ T$ ”.

We conjecture that a principal solution always exists (and is always found by our algorithm) if the merge problem only contains

duplicable data, and does not contain any data type with more than one branch.

Third example Merging is even more difficult when nominal types come into play.

```
val z =
  let y = ... in
    if ... then
      Some { contents = y }
    else
      None
```

Let us consider the case where y refers to duplicable data, and assume for the sake of example that we have “ $y @ \text{int}$ ”. This merge problem admits a principal solution, namely “ $z @ \text{option } (=y) * y @ \text{int}$ ”. One can derive “ $z @ \text{option } \text{int}$ ” using [COPYDUP](#).

In the case that y refers to non-duplicable data, e.g. if we have “ $y @ \text{ref } \text{int}$ ”, then again, this problem does not admit a principal solution, as “ $z @ \text{option } (=y) * y @ \text{ref } \text{int}$ ” and “ $z @ \text{option } (\text{ref } \text{int})$ ” are not comparable⁶.

As a side-note, if y were bound in the *then* branch, the solutions would be similar, except y would be existentially-quantified in the resulting permission.

4.1.1 Three sub-problems

The merge operation poses several difficulties.

- We need to reconcile conflicting heap shapes into a common description. This implies traversing two graphs in parallel, which is similar to what has been done in the field of shape analysis [21]. This problem is predictably and efficiently treated by our algorithm.
- Because of mutable data, the algorithm has to decide, as in the first example, where to “assign” a piece of mutable data. We call this the *exclusive resource allocation problem*. Our algorithm can detect and warn the user about such cases, but lacking any backward analysis, is unable to decide which solution it should pick. Our first two examples would trigger such a warning.
- Folding inductive predicates is also a recurring problem in the literature; it is highlighted by the third and fourth examples. In *Mezzo*, the use of singleton types poses additional difficulties, as it offers multiple solutions for the parameter α when folding two types into, say, *list* α . Our algorithm is not complete and only considers a limited, yet predictable, subset of choices for the parameter.

4.2 Formal definition of the merge operation

We define a *root* to be a program variable defined before the merge operation, or the return value of the *match* expression. In the third example, y is a root, as it is bound before the disjunction; z is also a root, as it represents the value of the *if-then-else* expression.

We talk about the *original* permission (before the disjunction in control-flow), the *left* and *right* permissions, and the *destination* permission (the result of the merge operation).

We define a *point* to be any location in the heap accessible from a root. There are left points, right points, and destination points.

We assume a *bijection* ψ where $\psi(x) = x_l, x_r$ means that the left point x_l and the right point x_r map onto the destination point x . Initially, $\psi(\rho) = \rho_l, \rho_r$, meaning that the return values from the left and right expressions map onto the return value for the whole *match* expression. We also have $\psi(x) = x, x$ for all other roots, that is, variables that were already defined in the *original* environment.

⁶The system would need a rule for an η -expansion of data types. We currently do not have this, and do not plan to add it.

We define the merge of prefixed, *normalized* environments P_l and P_r to be:

$$(\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) = \mathcal{V}.R$$

A more useful presentation is one that, just like subtraction, focuses on a particular variable, and represents the working state of the algorithm using R , which stands for permissions that have been merged already.

$$\frac{\psi(x) = x_l, x_r}{(\mathcal{V}_l.P_l * x_l @ t \vee \mathcal{V}_r.P_r * x_r @ t) * \mathcal{V}.R = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.R * x @ t}$$

The difficulties lie, first, in finding the correct ψ function and second, in applying the right subsumption rules on P_l and P_r so as to obtain the same type on both sides of the disjunction.

4.3 Computing the mapping

Computing the ψ function is done by traversing the left and right permission graphs in parallel. One basically follows the structural permissions (tuples and constructors) starting from the roots, and associates a pair of matching points onto a destination point. We ensure that for a pair of matching x_l and x_r , only one destination x corresponds: ψ is really a bijection, meaning that it *preserves sharing*.

The ψ function, once completed, maps every pair of left and right points, reachable through the same path, onto a destination point.

The rules that implement the building of the ψ function are presented in Figure 8. They assume the environments to be normalized, since building the ψ function is only possible if structural types are *expanded* using singleton types.

4.4 Algorithmic presentation of merge

We assume the ψ function has been computed, and initially start with:

$$(\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.empty$$

If \mathcal{V}_0 is the prefix of the *original* permission, and if $\text{dom}(\psi)$ is the domain of the ψ function, we define \mathcal{V} to be:

$$\mathcal{V} = \mathcal{V}_0, \exists(\overline{\text{dom}(\psi)} : \text{term})$$

This is restrictive: a variable that is *local* to the left or right permission, and that has not been mapped by the ψ function, cannot appear in the destination permission. For instance, no meaningful permission will come out of the following merge $\{\{t\} \ t\}$ t is the concrete syntax for $\exists t.t$):

```
(* f @ () -> {t} t *)
val z = if ... then f () else f ()
```

Function f returns an existentially-quantified type. Per the restriction above, the algorithm will not know how to merge this. The user can solve the merge by providing a type annotation for z .

The reason for this restriction is that merging local types requires packing them under an existential; the results are not predictable, and we risk merging too many types. Indeed, with this strategy, “ $x @ \text{int} \vee x @ ()$ ” would be solved as “ $x @ \exists t.t$ ”!

The rules are presented in Figure 9. We write $t \in \mathcal{V}_0$ to signify that a type makes sense in the left, right, and destination permissions (thus rendering its merge trivial). Tuple and constructor types are translated using the ψ function (**MERGE-TUPLE**, **MERGE-CONSTRUCTOR**). Function types, type variables are merged using the default criterion, that is, syntactic equality (**MERGE-DEFAULT**). **MERGE-APP** deserves an explanation. We can only merge type applications of the same type. In that case, we recursively merge the parameters of the type application.

- If parameter i is bivariant, the type application will not use its parameter, meaning we can pick $\perp = \forall \alpha.\alpha$ for the resulting u_i .
- If parameter i is covariant, we can recursively merge the type parameters; we bind a fresh rigid variable y and re-use the merge algorithm to compute parameter u_i .
- If parameter i is invariant, we rely on a syntactic equality criterion to compute parameter u_i .
- If parameter i is contravariant, we would need to perform the *intersection* of two types, an operation that is not supported by our type-checker. We approximate, and use syntactic equality again.

4.5 Implementation strategy

Again, these rules provide very little in terms of implementation details. With the contents of Figure 9, we are still at loss for resolving the merge from the third example. Indeed, one needs to apply *subsumption rules* before the rule **MERGE-APP** can apply.

Difficulties related to merging only appear when merging distinct constructors (e.g. Nil vs. Cons), or a constructor with a nominal type (e.g. Cons vs. list t). Let us start with the latter, easier case.

Cons. vs. nominal We are faced with the following disjunction, where A is a data constructor that belongs to type t (the problem has no solution otherwise).

$$\mathcal{V}_l.P_l * x @ A \{\vec{f} : = \vec{y}\} \vee \mathcal{V}_r.P_r * x @ t \vec{u}$$

We could be subtle here, but for the sake of clarity, we require $\vec{u} \in \mathcal{V}_0$, which implies that all the type application’s parameters also make sense in the *left* environment.

We can thus perform, in the left environment, the following subtraction:

$$\mathcal{V}_l.P_l * x @ A \{\vec{f} : = \vec{y}\} - x @ t \vec{u} = \mathcal{V}'_l.P'_l$$

The left permission is weakened, but we gain in exchange a new merge problem, which is solvable using rule **MERGE-APP**.

$$\mathcal{V}'_l.P'_l * x @ t \vec{u} \vee \mathcal{V}_r.P_r * x @ t \vec{u}$$

Please note that we *keep* the resulting environment $\mathcal{V}'_l.P'_l$: the subtraction may have consumed permissions and instantiated flexible variables in the left permission, and we need to keep this information for future merges.

Cons. vs. cons We are now faced with the following disjunction, where both A and B are data constructors belonging to a common type t .

$$\mathcal{V}_l.P_l * x @ A \{\vec{f} : = \vec{y}_l\} \vee \mathcal{V}_r.P_r * x @ A \{\vec{f} : = \vec{y}_r\}$$

Again, subtraction is going to help us. We know that both these types should fold into “ $t \vec{u}$ ”. The difficulty is we don’t know the actual values of \vec{u} . We can perform the following subtraction:

$$\mathcal{V}_l, \exists(\vec{\alpha}_l : \vec{\kappa}).P_l * x @ A \{\vec{f} : = \vec{y}_l\} - x @ t \vec{\alpha}_l = \mathcal{V}'_l.P'_l$$

We introduce flexible variables along with the right kinds to stand for the parameters of the type application. The same operation is performed on the right-hand side, leading to the following subtraction, which is solvable using **MERGE-APP**:

$$\mathcal{V}'_l.P'_l * x @ t \vec{\alpha}_l \vee \mathcal{V}'_r.P'_r * x @ t \vec{\alpha}_r$$

If instantiated, the type-checker automatically substitutes a flexible variable α for its instantiation t . In the case where either one of α_l and α_r is uninstantiated, the type-checker will perform the intuitive instantiation. In the case where both remain flexible, a type annotation is mandatory.

PSI-TUPLE $\frac{\psi(x) = x_l, x_r \quad x_l @ (\dots, =y_l, \dots)}{\exists y'. \psi(y') = y_l, y_r \quad x_r @ (\dots, =y_r, \dots) \quad y \text{ fresh}} \quad \psi(y) = y_l, y_r$	PSI-CONSTRUCTOR $\frac{\psi(x) = x_l, x_r \quad x_l @ \mathbf{A} \{ \dots, f := y_l, \dots \}}{\exists y'. \psi(y') = y_l, y_r \quad x_r @ \mathbf{A} \{ \dots, f := y_r, \dots \} \quad y \text{ fresh}} \quad \psi(y) = y_l, y_r$
---	---

Figure 8. Computing the ψ function

MERGE-TUPLE $\frac{\psi(x) = x_l, x_r \quad \psi(\vec{y}) = \vec{y}_l, \vec{y}_r}{(\mathcal{V}_l.P_l * x_l @ (= \vec{y}_l) \vee \mathcal{V}_r.P_r * x_r @ (= \vec{y}_r)) * \mathcal{V}.R = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.R * x @ (= \vec{y})}$	MERGE-CONSTRUCTOR $\frac{\psi(x) = x_l, x_r \quad \psi(\vec{y}) = \vec{y}_l, \vec{y}_r}{(\mathcal{V}_l.P_l * x_l @ \mathbf{A} \{ \vec{f} := \vec{y}_l \} \vee \mathcal{V}_r.P_r * x_r @ \mathbf{A} \{ \vec{f} := \vec{y}_r \}) * \mathcal{V}.R = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.R * x @ \mathbf{A} \{ \vec{f} := \vec{y} \}}$
MERGE-APP $\frac{(\mathcal{V}_l, \exists(y : \text{term}). \overline{P}_l * y @ u_{l,l} \vee \mathcal{V}_r, \exists(y : \text{term}). \overline{P}_r * y @ u_{l,r}) = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.y @ u \quad \text{if } t \text{ is covariant in parameter } i}{\begin{array}{l} u = \forall \alpha. \alpha \quad \text{if } t \text{ is bivariate in parameter } i \\ u_l \in \mathcal{V}_0 \quad u_r \in \mathcal{V}_0 \quad u_l = u_r \quad \text{otherwise} \end{array}} \quad \psi(x) = x_l, x_r$ $\frac{}{(\mathcal{V}_l.P_l * t @ \vec{u}_l \vee \mathcal{V}_r.P_r * t @ \vec{u}_r) * \mathcal{V}.R = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.R * t @ \vec{u}}$	
MERGE-PERM-VAR $\frac{p \in \mathcal{V}_0}{(\mathcal{V}_l.P_l * p \vee \mathcal{V}_r.P_r * p) * \mathcal{V}.R = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.R * p}$	MERGE-DEFAULT $\frac{t \in \mathcal{V}_0 \quad \psi(x) = x_l, x_r}{(\mathcal{V}_l.P_l * x_l @ t \vee \mathcal{V}_r.P_r * x_r @ t) * \mathcal{V}.R = (\mathcal{V}_l.P_l \vee \mathcal{V}_r.P_r) * \mathcal{V}.R * x @ t}$

Figure 9. The merge operation

Instantiation choices The quality of the frequent *Cons* vs. *cons* case depends on the choice made by the type-checker for the flexible variables $\vec{\alpha}$. A heuristic is used, which the user can always bypass using a type annotation. It mainly consists in never instantiating a flexible variable with a type that contains singletons, for fear the singletons would refer to a local variable, thus rendering the instantiation choice useless for a merge.

4.6 Detecting exclusive resource allocation conflicts

The exclusive resource allocation conflicts from the first and second examples can be detected easily. When one merges $y_l @ t \vee y_r @ t$ onto $\psi(y) = y_l, y_r$, if t is exclusive, and there exists y' such that $\psi(y') = y_l, y'_r$ (resp. $\psi(y') = y'_l, y_r$), this means that there will be a conflict as to which of y and y' “gets” the exclusive permission from y_l (resp. y_r).

5. Related work

The entire design of *Mezzo* owns a great deal to separation logic [20]; our type system can be seen as an adaptation of separation logic into a type system for an ML-like language. This is achieved through the use of permissions: a permission *is* analogous to a separation logic assertion. This design choice has numerous consequences.

- Type annotations are mandatory for functions, meaning that the types of the arguments *are* the pre-condition, while the return type of the function is the post-condition. We do not need a separate language for pre- and post-conditions.
- We thus do not distinguish between spatial predicates and logical formulae: a permission embodies both concepts. The permission “ $x @ \text{list } a$ ” has a spatial reading, while a permission “ $x @ =y$ ” has a logical reading.
- Algebraic (inductive) data types are central in ML-like languages, and *Mezzo* makes no exception. Instead of defining separate hard-wired predicates for lists, trees, doubly-linked-lists,

Mezzo naturally relies on user-defined data types (such as `list`) for the spatial assertions (such as “ $x @ \text{list } a$ ”).

- Not having pointers, we do not have to prove that memory accesses are within bounds, which is a pervasive concern in other contexts [13]. We still share the same concerns about uniqueness, though.

Permissions The Java community has investigated better type-checking using permissions [6]. In this context, programs are first type-checked using the standard Java type system; then, permissions are elaborated so as to refine the usage protocols and possibly rule out mistakes. Various access control modifiers are proposed, to reflect the variety of usages in the Java world. We fuse the concepts of a type and a permission together, and keep a relatively simple access mechanism (duplicable / exclusive): so far, we haven’t felt the need for more elaborate mechanisms such as fractional permissions [7].

The Racket community, retrofitting a type system on top of Scheme, has had to deal with flow-sensitive typing [22]. Typing in *Mezzo* is flow-sensitive as well, and our constructor types, which encode the fact that we may be in the Nil or Cons case, are reminiscent of this. We do not employ the sophisticated implementation techniques like those they mention in §7.4.

Program proof Strong properties on programs can already be obtained using existing tools. One may fully prove a program using a proof assistant, or rely on program logics embedded in Coq such as Ynot [15] or CFML [10]. These tools’ underlying models are based on separation logic, and the notion of a type is replaced by a separation logic assertion. This is much more powerful, but this comes at the price of automation. Other tools include Why3 [12], which is not based on separation logic and uses weakest-precondition inference, sending off proof obligations to a battery of SMT-solvers. The Java community has been know to use KeY [3], which uses its own dynamic logic.

We believe these tools are targeted at experts in program proof, which is not our target audience. *Mezzo*, by contrast, is a type system, whose reasoning is performed automatically, and which requires no manual intervention apart from type annotations.

Frame inference Calcagno *et al.* perform a symbolic execution of pointer programs with Smallfoot [4, 5]. They introduce a decidable algorithm for inferring *frame axioms*, that is, portions of the current separation logic assertion that are left untouched by a call to a subroutine. They assume unquantified assertions, and only consider a subset of hard-wired inductive predicates.

Nguyen *et al.* [17] allow user-defined inductive predicates. They present an algorithm that, even if not complete, terminates. They allow a limited form of quantification.

Pérez and Rybalchenko [16] have a correct and complete algorithm for a fragment of separation logic with list segments. They deal with linear formulas that contain exactly one spatial predicate, and treat disjunction.

In the setting of *Mezzo*, function types are fully annotated. The problem of performing frame inference thus seems easy, and would merely consist in “taking out” the parts that are consumed by a function call. We do not have disequality permissions, which makes our task seemingly easier than other works [5] which, dealing with pointers, often need to prove that two variables must not alias. We do, however, need to deal with extra difficulties: arbitrary quantifiers and proper instantiation, as well as handling duplicable and non-duplicable portions are the main ones. It is thus unlikely that our algorithm would prove useful in traditional pointer languages with separation logic, as it is tightly coupled with an ML-like setting.

Computing joins What we call the merge problem received comparatively little attention from the program proof and static analysis communities [21, 23, 9]. There are several reasons for this. Indeed, computing the merge, or “join” of two branches amounts to performing an approximation. There are valid reasons not to do that, such as preserving the quality of the solutions, or avoiding the complexity of computing an approximation. In this case, each branch is treated separately, which may indeed lead to an exponential number of branches when one reaches the end of control-flow.

Our algorithm could be seen as performing the same task as a “join” in a shape domain for abstract interpretation. Our algorithm is different from a “regular” algorithm for shape domains, though, in the sense that it is targeted specifically at an ML-like language: thanks to ML’s tagged sums, we do not need to explore arbitrary unfoldings of inductive predicates; we also handle the difficulties that arise from duplicity *vs.* non-duplicity of permissions.

Adding an anonymous union (sum) type to our system would certainly simplify merge operations, but we would then lose the aforementioned benefits of the tagged sums of ML. Indeed, in *Mezzo*, disjunction is somehow easier to treat. For a pointer language, reasoning of the form $x = \text{null} \vee x \mapsto \text{list}$ takes place, which requires exploring disjunctions and manipulating predicates for *list segments*. Algebraic data types are *tagged* sums, meaning that disjunctions are destructed explicitly via `match` expressions, which make the job of the type-checker easier.

The need for a “join” operation arises in a different context for pointer languages. In pointer languages, loops are central. Computing the shape of a loop requires “joining” shapes that denote successive iterations of the loop. The procedure ultimately reaches a fixed point which then stands for the loop invariant. We do not have this issue: `while` loops are replaced with *recursive* functions, that express different pre- and post-conditions. Combined with recursive reasoning [18], this alleviates the need for list segments.

Proving our type-checker A version of the Smallfoot program analyzer has been proven sound using the Coq proof assistant [1].

We currently do not have any soundness or completeness result for our type-checking algorithms.

6. Conclusion

We presented the implementation of a type-checker for the *Mezzo* language. *Mezzo* is based on permissions and features singleton types, constructor types, first-class functions, which make implementing a type-checker difficult. Type-checking a *Mezzo* program leverages two central operations, *subtraction* and *merge*. We have shown how to perform both operations in the context of *Mezzo*, extending previous work in the area.

The *Mezzo* type-checker performs well: the merge problem, with the additional heuristic that resource allocation conflicts should always be solved in favor of the return value of the `match` expression, require less than five annotations over our entire code-base. Subtraction, once fine-tuned with the heuristic that a flexible variable should never instantiate onto a singleton type, performs well most of the time. Type annotations are required for polymorphic function calls as soon as a type variable α should instantiate onto $(t \mid P)$, because the type-checker never explore such a solution (there would be too many choices for P).

The implementation is slow on a couple of examples, namely those from our iterator library [14]: the 450-line file takes up to two minutes to type-check on a modern computer. The reason is the two files use a specific programming pattern, where a variable f has several function permissions available. The type-checker thus explores several branches when asked for a specific function type, which slows down operations. Another unsatisfactory aspect of the type-checker is the lack of predictability for type annotations on polymorphic function calls. As of now, we are unable to provide the user with a simple criterion for deciding when to annotate such a call.

We have strong confidence in the correctness of our type-checker. In the future, we would like to obtain a completeness result for our algorithm, even if on a limited subset of *Mezzo* programs. We conjecture that our type-checker is complete when the only quantifiers are at kind term. We would also like to improve inference, as our algorithm only explores a subset of the solution space.

References

- [1] Andrew W. Appel. VeriSmall: Verified Smallfoot shape analysis. In *Certified Programs and Proofs*, volume 7086 of *Lecture Notes in Computer Science*, pages 231–246. Springer, 2011.
- [2] Thibaut Balabonski and François Pottier. A Coq formalization of *Mezzo*, take 2, December 2013. <http://gallium.inria.fr/~fpottier/mezzo/mezzo-coq.tar.gz>.
- [3] Bernhard Beckert, Reiner Hähnle, and Peter H. Schmitt, editors. *Verification of Object-Oriented Software: The KeY Approach*. LNCS 4334. Springer-Verlag, 2007.
- [4] Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *Formal Methods for Components and Objects*, volume 4111 of *Lecture Notes in Computer Science*, pages 115–137. Springer, 2005.
- [5] Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. Symbolic execution with separation logic. In *Asian Symposium on Programming Languages and Systems (APLAS)*, volume 3780 of *Lecture Notes in Computer Science*, pages 52–68. Springer, 2005.
- [6] Kevin Bierhoff and Jonathan Aldrich. Modular typestate checking of aliased objects. In *Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 301–320, 2007.
- [7] John Boyland. Checking interference with fractional permissions. In *Static Analysis Symposium (SAS)*, volume 2694 of *Lecture Notes in Computer Science*, pages 55–72. Springer, 2003.

- [8] Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. Compositional shape analysis by means of bi-abduction. In *Principles of Programming Languages (POPL)*, pages 289–300, 2009.
- [9] Bor-Yuh Evan Chang and Xavier Rival. Relational inductive shape analysis. In *Principles of Programming Languages (POPL)*, pages 247–260, 2008.
- [10] Arthur Charguéraud. Program verification through characteristic formulae. In *International Conference on Functional Programming (ICFP)*, 2010.
- [11] Dino Distefano and Matthew J. Parkinson. jStar: towards practical verification for Java. In *Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 213–226, 2008.
- [12] Jean-Christophe Filliâtre and Claude Marché. Multi-prover verification of C programs. In *International Conference on Formal Engineering Methods (ICFEM)*, volume 3308 of *Lecture Notes in Computer Science*, pages 15–29. Springer, 2004.
- [13] Dan Grossman, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in Cyclone. In *Programming Language Design and Implementation (PLDI)*, pages 282–293, 2002.
- [14] Armaël Guéneau, François Pottier, and Jonathan Protzenko. The ins and outs of iteration in *Mezzo*. Workshop on Higher-Order Programming and Effects (HOPE), September 2013. <http://goo.gl/NrgKc4>.
- [15] Aleksandar Nanevski, Greg Morrisett, Avraham Shinnar, Paul Govereau, and Lars Birkedal. Ynot: dependent types for imperative programs. In *International Conference on Functional Programming (ICFP)*, pages 229–240, 2008.
- [16] Juan Antonio Navarro Pérez and Andrey Rybalchenko. Separation logic + superposition calculus = heap theorem prover. In *Programming Language Design and Implementation (PLDI)*, pages 556–566, 2011.
- [17] Huu Hai Nguyen, Cristina David, Shengchao Qin, and Wei-Ngan Chin. Automated verification of shape and size properties via separation logic. In *Verification, Model Checking and Abstract Interpretation (VMCAI)*, volume 4349 of *Lecture Notes in Computer Science*, pages 251–266. Springer, 2007.
- [18] François Pottier and Jonathan Protzenko. Programming with permissions in Mezzo. In *International Conference on Functional Programming (ICFP)*, pages 173–184, 2013.
- [19] François Pottier and Didier Rémy. The essence of ML type inference. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 10, pages 389–489. MIT Press, 2005.
- [20] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science (LICS)*, pages 55–74, 2002.
- [21] Xavier Rival. *Abstract Domains for the Static Analysis of Programs Manipulating Complex Data Structures*. Habilitation à diriger des recherches, École Normale Supérieure, 2011.
- [22] Sam Tobin-Hochstadt and Matthias Felleisen. Logical types for untyped languages. *ACM Sigplan Notices*, 45(9):117–128, 2010.
- [23] Hongseok Yang, Oukseh Lee, Josh Berdine, Cristiano Calcagno, Byron Cook, Dino Distefano, and Peter W. O’Hearn. Scalable shape analysis for systems code. In *Computer Aided Verification*, volume 5123 of *Lecture Notes in Computer Science*, pages 385–398. Springer, 2008.