

Avoiding Signature Avoidance in ML Modules with Zippers

CLÉMENT BLAUDEAU, Inria, France and Université de Paris Cité, France

DIDIER RÉMY, Inria, France

GABRIEL RADANNE, Inria, France, EnsL, France, UCBL, France, CNRS, France, and LIP, France

We present ZIPML, a new path-based type system for a fully fledged ML-module language that avoids the signature avoidance problem. This is achieved by introducing *floating fields*, which act as additional fields of a signature, invisible to the user but still accessible to the typechecker. In practice, they are handled as *zippers* on signatures, and can be seen as a lightweight extension of existing signatures. Floating fields allow to delay the resolution of instances of the signature avoidance problem as long as possible or desired. Since they do not exist at runtime, they can be simplified along type equivalence, and dropped once they became unreachable. We give rewriting rules for the simplification of floating fields without loss of type-sharing and present an algorithm that implements them. Remaining floating fields may fully disappear at signature ascription, especially in the presence of toplevel interfaces. Residual unavoidable floating fields can be shown to the user as a last resort, improving the quality of error messages. Besides, ZIPML implements early and lazy strengthening, as well as lazy inlining of definitions, preventing duplication of signatures inside the typechecker. The correctness of the type system is proved by elaboration into M^ω , which has itself been proved sound by translation to F^ω . ZIPML has been designed to be an improvement over OCAML that could be retrofitted into the existing implementation.

CCS Concepts: • **Software and its engineering** → *Functional languages; Semantics; Modules / packages; Theory of computation* → **Type theory; Type structures.**

Additional Key Words and Phrases: Modules, ML, Zippers, Signature avoidance, Strengthening, OCaml

ACM Reference Format:

Clément Blaudeau, Didier Rémy, and Gabriel Radanne. 2025. Avoiding Signature Avoidance in ML Modules with Zippers. *Proc. ACM Program. Lang.* 9, POPL, Article 66 (January 2025), 30 pages. <https://doi.org/10.1145/3704902>

1 Introduction

Modularity is essential to the design, development, and maintenance of complex systems. For software systems, language-level mechanisms are crucial to manage namespaces, enforce interfaces, provide encapsulation, and promote code factorization and reuse. A wide variety of modularity techniques appear in different programming languages: from simple functions to libraries, compilation units, objects, type-classes, packages, etc. In languages of the ML family (OCAML, SML, Moscow ML, etc.), modularity is provided by a *module system*, which forms a *separate* language layer built on top of the core language.¹ ML modules are renowned for their expressiveness and flexibility, making them adaptable to numerous contexts, from large-scale OS-libraries [19] to

¹1ML [26] is a proposal to unify the two layers in principle, although some stratification should persist in practice.

Authors' Contact Information: Clément Blaudeau, Inria, Paris, France and Université de Paris Cité, Paris, France, clement.blaudeau@inria.fr; Didier Rémy, Inria, Paris, France, didier.remy@inria.fr; Gabriel Radanne, Inria, Lyon, France and EnsL, Lyon, France and UCBL, Lyon, France and CNRS, Lyon, France and LIP, Lyon, France, gabriel.radanne@inria.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/1-ART66
<https://doi.org/10.1145/3704902>

complex parameterized interpreters [24]. At the most basic level, types and values are gathered in *modules*, which can be *parameterized*. The interface of a module (its type) is called a *signature*.

A key feature of ML modules is *encapsulation*, provided by *ascription*: a module can be forced to a certain signature. If the signature contains *abstract-type fields*, the actual definition of the type, hence the implementation details, are hidden. By using ascription, a developer can make sure that the data-structures are accessed only through the functions defined inside the module, before the ascription. In essence, this allows enforcing complex invariants: library designers only have to ensure that public functions preserve such invariants, thanks to the *language-wide guarantee* that users cannot break abstractions [4]. An overview of other advanced features (generative and applicative functors, transparent ascription, etc.) can be found in [2]. Providing a theoretical background for the seemingly simple mechanism of abstract-type fields turned out to be the crux of module systems. In their foundational paper of 1985, Mitchell and Plotkin [22] suggested to represent abstract types as *existential types*. This vision was opposed by MacQueen [18] who proposed using a restriction of dependent types (strong sums) to describe module systems, although further work on phase separation [10] supported the idea that dependent types may actually be too powerful for module systems. Later on, two approaches for the formalization and improvement of abstract types in SML were independently yet simultaneously described by Leroy [13] using manifest types, inspired by the dot notation [3], an extension of existential types with open scope, and Harper and Lillibridge [9] via an adaptation of F^ω , the higher-order polymorphic lambda calculus, with translucent sums. It is Russo [31] who reintroduced existential types for module typing and explained their open scope with an *extrusion* mechanism. Applicative functors were identified as having higher-order existential types by Biswas [1] and the extrusion mechanism was extended to support *skolemization* [2, 28, 31]. In this setting, syntactic, user-writable signatures are elaborated in the more expressive language of types of F^ω and type-sharing between modules is expressed differently: abstract type fields introduce existential type *variables*, possibly higher-order, quantified in front of the signature, so that two modules that share an abstract type simply refer to the same type expression. This has culminated in the successful, so-called, *F-ing* line of works [2, 26–28, 31, 32] where most significant ML module features are specified and proved sound by *elaboration* in F^ω , thus benefiting from the meta-theoretical properties of the target language. By contrast, purely syntactic systems for ML modules use somewhat nonstandard typing rules and their meta-theory, which has proved hard to formalized, has to be redone from scratch [7, 9, 10], requiring complex syntactic techniques or semantic objects.

Yet, ML modules system remained a case where advanced theoretical works had a limited impact on real-world implementations: neither OCAML nor SML compilers are based on an elaboration into F^ω . Notably, OCAML relies on a *path-based* system, initially described by Leroy [13, 14], which specification has not been extended to more complex constructs. The compiler has evolved to support new features and be more efficient, but maintaining an internal representation of signatures that is *more-or-less* syntactic. The SML compiler has an official specification [20, 21] and a subset of the language has been mechanically formalized using singleton types. The language has also evolved over the years, but not towards an elaboration in F^ω (see [17] for more details). Notably, Moscow ML, an implementation of the SML standard extended with applicative functors, recursive modules, and first-class modules [29–31], is an exception, as it uses F^ω -like types internally.

Overall, for ML modules, *explainability*, *usability*, and *soundness* seems to be misaligned goals. While the *F-ing* approach fulfills the last one, usability of industrial-grade *F-ing*-based compiler remained to be demonstrated.² Moreover, it creates a significant gap between the user writable

²In particular, the *lazy inlining* of definitions, which seems essential for efficiency, has not been formalized nor implemented in Moscow ML and could be problematic for applications that make a very intensive use of modules.

```

1  module type Comparable = sig type t val eq : t → t → bool end
2  module type Keys = sig type t type k val getKey : t → k val fast_eq : k → k → bool end

3  module Map (E: Comparable)
4    (K: Keys with type t := E.t) = struct
5    type map = (K.k * (E.t * int)) list
6    let insert x n (l : map) =
7      let k = K.getKey x in ...
8      if (K.fast_eq k k') then ...
9      if (E.eq x x') then ...
10   let get x m = ...
11   let get_from_key k m = ... (* ... *)
12 end

13 module Map (E: Comparable)
14   (K: Keys with type t := E.t) : sig
15   type map
16   val insert
17     : E.t → int → map → map
18   val get : E.t → map
19   val get_from_key
20     : K.k → map → (E.t * int) list
21   (* ... *)
22 end

```

Fig. 1. Example of OCaml code which could trigger signature avoidance. The left-hand side is the code file while the right-hand side is the interface (typically, with the code in a `.ml` file and the interface in a `.mli` file).

source signatures and their internal representations in F^ω , which can undermine explainability. Requiring users to think in terms of the internal language while still writing types in the surface language imposes mental gymnastics, as the elaboration between the two languages is quite involved. Another option would be to conceal the internal representation with a reverse translation, called *anchoring* in [2]. Unfortunately, it seems difficult to truly hide the internal representation: (1) Anchoring can fail because types of F^ω are more expressive than the source language signatures. Consequently, some inferred signatures cannot be expressed in the source signature syntax—and the program must be rejected. This issue, called *signature avoidance*, is a serious problem in all syntactic approaches. It would trigger a specific class of typechecking errors that might be tricky for users to understand, as it might require exposing the internal representation of types. (2) The printing of a signature for error messages might cause an unrelated signature avoidance error, which might be confusing for the user. (3) Even when the typechecking succeeds, the M^ω type-system [2] combined with anchoring is not *fully syntactic* in the sense of [32]: if a module expression admits a signature, not all sub-expressions necessarily do so. This might be counter-intuitive for the user, as simply exposing a submodule might make typechecking fail in non-trivial ways, which affects usability.

In the rest of this section, we discuss our design for a new syntactic module system, called ZIPML. We address explainability maintaining syntactic signatures as internal representation which should match the user’s intuitive understanding of syntax and usability by proposing a somewhat *conservative* extension of a path-based system that *should not* interfere with other features or affect performance trade-offs. We have not implemented ZIPML, so this claim is not backed by experiments yet. We maintain *soundness* by translation in M^ω [2], leveraging the *F-ing* line of works. Namely, (1) we delay, and resolve *signature avoidance* by zipping out-of-scope fields; (2) we handle applicative functors and module identities; (3) we ensure a proper sharing of types, via the so-called *strengthening* operation [13], but with a new *lazy and early* strategy; (4) we maintain intermediate module type definitions, only lazily inlining them, so as to print concise interfaces respecting the user’s intent—an aspect previously left behind in the *F-ing* line of works.

1.1 Modules, Abstraction, and Type Sharing

1.1.1 Introductory example. Let us start with a concrete example of modular code given in Figure 1. We want to build a library that provides parametric integer maps. To that aim, we define the module-type `Comparable` (line 1) that represents the minimal signature that a module must satisfy for our library to build maps on: it must contain a type definition `type t` (left abstract for polymorphism), and an equality function `val eq : t → t → bool`. Then, we define `Map` (line 3) as a *functor* that

creates a structure given a module parameter E that satisfies the interface `Comparable`. However, for performance reasons, we want to limit the number of equality tests. To do so, we define a module-type `Keys` (line 2) for modules that provide a type k of keys, a key generator `getKey`, and a fast equality test `fast_eq` for keys. `Map` then takes a second module parameter $K : \text{Keys}$ as input (line 4), used to shortcut equality tests. Note that the generator `getKey` does not have to be injective: `Map` uses the default equality on elements when their keys are equal (lines 8 and 9). Internally, `Map` represents maps as lists of triples containing a key, an element, and its associated integer (line 5). Among the functions provided by `Map`, we have `get_from_key` (line 12) that returns the list of stored elements that have the same key. In the interface of `Map`, we force the type `map` to be abstract (line 16). It serves two purposes: first, it hides implementation details, especially the fact that we used lists to represent maps; second, it prevents users from inserting objects along with the wrong key, which would violate internal invariants of the module. A user can then instantiate `Map` with any comparable module along with a type of keys. Here, we create a datatype for keys:

```
module T = struct type t = Node of string * t * t | Leaf let eq = ... end
module K = struct type k = Root of string | Length of int ... end
module M = Map(T)(K)
```

However, this reveals the implementation details of K . One might hide them by writing, directly:

```
module M = Map(T)(struct type k = Root of string | Length of int ... end)
The parameter cannot be eliminated in the result type.
```

Unfortunately, this fails to typecheck, as the interface of `Map` mentions the type $K.k$ (line 21), while there is no name to refer to it outside of its enclosing structure. This is a case of the *signature avoidance* problem. Naming K solves this case. However, forcing users to name every module would impact usability and make some code patterns impractical.

1.1.2 Signature Avoidance. At a more abstract level, the signature avoidance problem can occur whenever type declarations become inaccessible while dependencies on those declarations remain. We illustrated it above with a functor call on an unnamed structure, but we now use module-level let-binding `let X = M in M'` or, equivalently³, projection out of an arbitrary module $M.X$. Neither constructs are present in OCAML which only allows for projection out of named modules. They can both encode functor calls on arbitrary module expressions⁴, but the converse is more verbose and less general, as it requires a parameter signature. Overall, the core problem remains the same: given a signature S that depends on a type t , can we extract S and keep it well-formed even when t has become inaccessible?

```
module M = (... : sig type t module X : S end).X
```

There might not exist a principal signature S' that is equivalent to S while *avoiding* the type t . The key issue here is to assume that hiding the field `type t` necessarily means *deleting* it. Indeed, one could argue that hiding a field makes it inaccessible *to the user* but not necessarily to the typechecker. This is the stance taken by Harper and Stone [11] and Dreyer et al. [7]: their approach relies on an elaboration with reserved names:

```
module M : sig module HIDDEN : sig type t end module VISIBLE : S end
```

³`let X = M in M'` can be defined as syntactic sugar for the projection `(struct module X = M module Y = M') . Y`. Conversely, `M.X` can itself be defined as `let Y = M in Y.X`, using only let-binding and projection on a named module.

⁴A functor call on an arbitrary module $F(M)$ can be defined as either the module let-binding `let X = M in F(X)` or the projection `(struct module Arg = M module Res = F(Arg) end) . Res`.

We share their intuition, but find the technical solution to achieve it too rigid: we would like to *focus* on S while keeping the rest of the fields *on the side*, if needed. A classical technique in functional programming, *zippers* [12], allows precisely to extend any tree-like type definition with the ability to focus on certain parts of the tree while keeping the rest on the side. As signatures are record trees, we can change the meaning of hiding to be the *zipping* of a signature onto the visible field, keeping the rest, if needed, as *floating fields* in the *zipper context*. Here, ZIPML would return:

```
module M : < type t > S
```

ZIPML uses zippers to preserve relevant type information and therefore delay the resolution of signature avoidance until a source signature is forced by ascription. Zippers also behave gracefully in case of errors, allowing to present users with the field names that were lost during typechecking. An orthogonal approach to the avoidance problem based on focusing can be found in [5].

1.1.3 Applicativity, abstraction safety, and aliasing. Applicative functors, introduced by Leroy [14], are one of the key features of OCAML. They initially relied on a simple *syntactic criterion*: two paths with functor applications are considered equal when they are syntactically equal:

```
module F (Y : sig end) = struct end
module G (Y : sig end) = struct type t end
module X = struct end
let compatible (x : G(F(X)).t) = (x : G(F(X)).t)      (* OCaml - typechecks *)
```

The strength of this criterion is that it preserves *abstraction safety* (see [2, 28]). However, it is also fragile, as naming a subexpression can break syntactic equality:

```
module FX = F(X)
let compatible (x : G(FX).t) = (x : G(F(X)).t)      (* OCaml - fails *)
```

Here, the type of module FX should manifest its module-level equality with $F(X)$. For this purpose, ZIPML reuses a syntactic construction, *transparent signatures* ($= P < S$), proposed by Blaudeau et al. [2], which generalizes module aliases [8] and allows expressing module level-sharing.⁵

1.1.4 Strengthening. In path-based systems, type sharing is expressed with manifest types [13] using explicit equalities between type fields. Consequently, to avoid a loss of type sharing when copying or aliasing a module, its signature must be rewritten, hence copied, to refer to the type fields of the original module. This operation, called *strengthening*, is central to *path-based* systems such as OCAML. However, rewriting a whole signature can be costly⁶ and hinder performance for very large libraries. To prevent useless rewriting while keeping type sharing, ZIPML implements three innovations: strengthening is made lazy and early. Laziness is achieved by using transparent signatures ($= P < S$) to mark the strengthening of S by P in the syntax tree, but delay the actual duplication of S . Besides, signatures are strengthened when entering the typing environment rather than when retrieving them from the typing environment, which becomes a standard typing rule.

1.1.5 Lazy expansion. Module languages allow for both type definitions (**type** $t = \text{int} * \text{float}$) and module type definitions (**module type** $T = \text{sig} \dots \text{end}$). However, while OCAML retains such definitions internally, as any real-world implementation, both M^ω and *F-ing* handle them by *eager inlining*. Inlining of module-type definitions in signatures might considerably increase their size and make typechecking of large-scale libraries with intensive use of modules quite expensive. It also loses the programmer's intent by inferring large signatures whose names or aliases have been

⁵In fact, transparent signatures were already expressible in *F-ing* [28]—see [2] for details.

⁶See the discussion at <https://github.com/ocaml-flambda/flambda-backend/pull/1337>.

lost. ZIPML keeps them internally and in inferred signatures, instead of systematically inlining them. This makes ZIPML closer to an actual implementation, such as OCAML, which could quickly and easily benefit from our solution to signature avoidance, as well as transparent ascriptions.

1.2 ZIPML

In this paper, we propose ZIPML, a fully-syntactic specification of an OCAML-like module system that supports both generative and applicative (higher-order) functors, opaque and transparent ascription, type and module-type definitions, extended with transparent signatures and the new concept of zippers with floating fields. Floating fields follow the way the user would resolve instances of signature avoidance manually, by adding extra fields in structures. However, while this pollutes the namespace with fields that were not meant to be visible, ZIPML floating fields are added automatically and can only be used internally: they are not accessible to the user and are absent at runtime. Besides, floating fields can be simplified and dropped after they became unreferenced. This mechanism is an internalized, improved counterpart of the anchoring of [2].

Our contributions are:

- The introduction of a syntactic type system for a fully-fledged OCAML-like module language, including both generative and applicative functors, and extended with transparent signatures.
- A new concept of *floating fields*, implemented as *zippers*, that enables to internalize and avoid (or delay) signature avoidance resolution.
- A criterion for signature avoidance resolution without loss of type sharing, along with the description of an algorithm to compute such resolution.
- A formal treatment of type and module type definitions that are kept in inferred signatures.
- A systematic, lazy and early treatment of strengthening that prevents useless inlining of module type definitions and increases sharing inside the typechecker.
- A soundness proof by elaboration of ZIPML into M^ω .

Plan. In §2, we start with a detailed overview of floating fields and zippers. In §3, we present the syntax and typing rules of ZIPML. In §4, we introduce narrow subtyping, a restriction of subtyping that allows for the *simplification* of floating fields without lost of type sharing. We present an algorithm that computes such simplification. In §5, we state formal properties of ZIPML, including type soundness, for which we give the structure of a proof by elaboration into M^ω . Finally, we discuss missing features and future works in §6.

2 An introduction to floating fields

The main novelty of ZIPML is the introduction of floating fields as a way to delay and resolve instances of the signature avoidance problem. Floating fields provide additional expressiveness that allows to describe all inferred signatures. We use OCAML-like syntax [15] for examples. However, we use self-references for signatures⁷ to refer to other components of themselves. For instance, while we would write in OCAML:

```
sig type t type u = t → t val f : u end
```

we instead write in ZIPML:

```
sig (A) type t = A.t type u = A.t → A.t val f : A.u end
```

The first occurrence of A is a binder that refers to the whole signature, so that all internal references to a field of that signature go through this self-reference. This is also used for abstract types who are represented as aliases to themselves as **type** $t = A.t$ (in gray). Self-references are just a convenient

⁷ZIPML also uses self-references in structures, but we do not in examples to keep closer to OCAML syntax.

syntactic notation that does not bring additional expressiveness, nor any cyclic references, but avoids some forms of shadowing and simplifies the definition of strengthening.

We show examples with a module-level let-binding. We also use **type** *t* to introduce abstract types directly in a structure (a feature present in OCAML), but more realistic examples would use algebraic data-types⁸ or ascriptions.

2.1 Basic typing with zippers

We start with a simple example where OCAML fails to typecheck:

```
let module M = struct type t module X = struct let l : t list = [] end end in M.X
The value "l" has no valid type if "X" is hidden.
```

Instead, ZIPML will return⁹ the following *zipper* signature:

```
module M : < B: type t = B.t > sig(A) val l : B.t list end
```

The highlighted expression is the zipper-context. It contains a single *floating field* **type** *t* = *B.t*, bound to the self-reference *B*. Intuitively, the zipper signature is obtained as follows. Before projection, the signature of *M* is:

```
sig(B) type t = B.t module X : sig(A) val l : B.t list end end
```

When projecting on field *X*, we would like to return **sig(A) val l : B.t list end**, but the reference *B.t* would become dangling. The solution is to keep the type of *t* defined as a floating field *B.t*, which gives exactly the signature just above.

Another typical situation of signature avoidance is when a type is used in other type definitions:

```
(struct type t module X = struct type u = t list type v = t list end end).X
sig(A) type u type v end (* OCaml - over-abstraction *)
```

Here, instead of failing, OCAML silently turns *u* and *v* into abstract types, not only losing their list structure but also forgetting that these are actually equal types. We describe this as being erroneously resolved by *over-abstraction*. While abstraction is *safe*, as it could be achieved through an ascription, it is *incomplete* and therefore should be performed only when explicitly required by the user. In ZIPML, we would return the following signature:

```
< B: type t = B.t > sig(A) type u = B.t list type v = B.t list end
```

This is a correct answer, as no type information has been lost.

2.2 Deep projection and simplification

We now consider a module *M* with some nested submodules *X* and *Y*. We then project on a deeply nested module *M.X.Y*. Again, OCAML returns a non-principal signature due to over-abstraction.

```
let M = struct type t
  module X = struct type u = t list
    module Y = struct type v = t type w = u end end
end in M.X.Y
sig(C) type v type w end (* OCaml - over-abstraction *)
```

⁸From a module level point of view, a declaration of an ADT **type** *t* = *C*₁ **of** *int* | *C*₂ **of** *bool* can be thought of as an abstract type **type** *t* followed by value bindings **val** *C*₁ : *int* → *t* and **val** *C*₂ : *bool* → *t*.

⁹By default, input programs are colored in blue; errors in red; output signatures of both OCAML and ZIPML are in green, with highlights in yellow, orange or red. We may still temporarily use other colors to emphasize specific subexpressions.

In ZIPML, we first return the following signature S_0 with floating fields:

```
⟨ A : type t = A.t end ⟩ ⟨ B : type u = A.t list ⟩ (* S0 *)
  sig(C) type v = A.t type w = B.u end
```

which we can then simplify to its equivalent final form S_1 :

```
sig(C) type v type w = C.v list end (* S1 *)
```

We now explain both steps, projection and simplification, separately.

2.2.1 *Projection.* To understand S_0 , let us look at the signature of M , before the projection:

```
M : sig(A) type t = A.t
  module X : sig(B) type u = A.t list
    module Y : sig(C) type v = A.t type w = B.u end end
end
```

It is of the form $S_A [S_B [S_C]]$, where the signature S_C is placed in the context $S_A [S_B [\cdot]]$ and S_A and S_B are the outer and inner contexts. Unfortunately, the signature S_C is ill-formed outside of those contexts. For the first projection $M.X$, we turn the outer signature into a zipper context, which gives $\langle S_A \rangle S_B [S_C]$. For the second projection $(M.X).Y$, we need to project out of a zipper signature, which is not a structural signature. However, it actually composes well: any operation on a zipped signature correspond to pushing the zipper context in the typing environment, doing the operation and popping the zipper context back. For our projection, we therefore first push the zipper context in the environment, leaving us with $S_B [S_C]$. Projecting gives $\langle S_B \rangle S_C$, and popping the zipper context back gives: $\langle S_A \rangle (\langle S_B \rangle S_C)$. Finally, we merge the two zipper contexts into a single one and obtain $\langle S_A; S_B \rangle S_C$, which is exactly S_0 . Conceptually, we may sum up those steps as:

$$(S_A [S_B [S_C]].X).Y \rightsquigarrow ((\langle S_A \rangle S_B [S_C]).Y \rightsquigarrow \langle S_A \rangle (\langle S_B \rangle S_C).Y) \rightsquigarrow \langle S_A \rangle (\langle S_B \rangle S_C) \rightsquigarrow \langle S_A; S_B \rangle S_C$$

2.2.2 *Simplifying zippers.* We now explain the simplification process from S_0 to S_1 . Each component of the zipper contexts $\langle A : S_A \rangle$ or $\langle B : S_B \rangle$ can be directly accessed from S_C via its self-reference name A or B . Hence, S_C need not be renamed—provided we have chosen disjoint self-references while zipping. We may first inline the definition of $B.u$ in field w of S_C , leading to the signature:

```
⟨ A : type t = A.t ⟩ ⟨ B : type u = A.t list ⟩
  sig(C) type v = A.t type w = A.t list end
```

The floating component B is now unreferenced from the signature C and can be dropped. Since the type $C.v$ is equal to $A.t$, the field $C.w$ can be rewritten as $C.v \text{ list}$. We obtain the signature S_3 :

```
⟨ A : type t = A.t ⟩ sig(C) type v = A.t type w = C.v list end (* S3 *)
```

The signature could also be seen as the projection of the unzipped signature (using some reserved field Z for the lost¹⁰ projection path):

```
(sig(A) type t = A.t module Z : sig(C) type v = A.t type w = C.v list end end).Z
```

Currently $C.v$ is an alias to the floating abstract type $A.t$, which comes first. However, since the module X will become a floating field, absent at runtime, it may be moved after field Z , letting $C.v$ become the defining occurrence for the abstract type and $A.t$ be an alias to $C.v$. The key is that the two unzipped signatures, before their projection, are subtype of one another, hence equivalent.

¹⁰Our zippers are *partial*, since we dropped both the name of the field we projected on and the fields following the projection. We therefore cannot, but need not, recover the exact original signature.


```
(sig(A) module Z : sig(C) type v = C.v type w = C.v end type t = A.Z.v end).Z
```

We may now project back the signature on field Z . It no longer depends on the following field t , which can safely be dropped. This leads to the signature S_1 , thus equivalent to S_3 and to S_0 .

```
sig(C) type v type w = C.v list end (* S1 *)
```

In this case, we were able to eliminate all floating fields and therefore successfully and correctly resolve signature avoidance, while OCAML incorrectly removes the equality between types v and w .

In general, simplification may remove some, but not necessarily all, floating fields. This is acceptable, as we are able to pursue typechecking in presence of floating fields, which may be dropped later on. For example, while S_3 could be simplified by removing its floating field, this was not the case of the signature of $M.X$, since the floating field $A.t$ of $\langle S_A \rangle$ is referenced in $S_B[S_C]$ before being aliased. If we disallowed floating fields, we would have failed at that program point—or used over-abstraction as in OCAML in this case, likely causing a failure later as a consequence of over-abstraction.

Interestingly, when typechecking an ascription ($M : S$), the signature returned in case of success will be the elaboration of the *source signature* S , which never contains floating fields. Thus, an OCAML library defined by an implementation file M together with an interface file S never returns floating fields in ZIPML, even if internally some signatures will carry floating fields. In other cases, we may return an answer with floating fields, giving the user the possibility to remove them via a signature ascription. As a last resort, we may still keep them until link time—or fail, leaving both options to the language designer or the user.

3 Formal presentation

An overview of the type system is given in Figure 2. The core of the system is composed of three judgments: **module typing** (§3.6), **signature typing** (§3.5) and **subtyping** (§3.4). The main judgment, module typing $\Gamma \vdash M : S$, infers the signature S of a module expression M . Since module expressions may contain signatures, module typing relies on signature typing ①, which is used to check them¹¹. We use subtyping to check ascription in both module expressions ② and transparent signatures ③. Since paths appear in signatures, we extract **path typing** (§3.3.2) from module typing to avoid all typing judgments to be recursively defined. All judgments handle paths and therefore depend on path typing (dashed blue arrows). Path typing also depends on subtyping ④ since paths include functor applications to cope for applicative functors.¹² As ZIPML maintains user-written definitions of core-language types and module types, we need a notion of normalization (dotted gray arrows). Finally, we use helper judgments to handle lazy and shallow **strengthening** (§3.2) and early strengthening **environment extension**. After presenting the grammar and some technical choices in §3.1, we detail the judgments in reverse order of dependencies (from right to left).

3.1 Grammar

The syntax of ZIPML is given in Figure 3. It reuses the syntax of OCAML, but with a few differences in notation, the addition of transparent ascriptions, and our key contribution: signatures with *floating fields*, also called *zippers*. As meta-syntactic conventions, we use lowercase letters (x, t, \dots) for elements of the core language, and uppercase letters (X, T, \dots) for modules. We also use slanted

¹¹Signature typing is also used as a meta-theoretic well-formedness of signatures, maintained throughout the system.

¹²Technically, this comes from the fact that we use path typing both as a *path-well-formedness check* and as a *path-lookup*. We could have two separate judgments: the former would depend on subtyping and be quite costly, while the latter, only used on known well-formed paths, would not recheck subtyping and would be faster.




Fig. 2. Relationship between the main judgments of ZIPML

Path and Prefix

$P ::= Q.X$ (Module)
 $\quad | Y$ (Module Parameter)
 $\quad | P(P)$ (Applicative application)
 $\quad | P.A$ (Zipper access)
 $Q ::= A | P$

Module Expression

$M ::= \dot{P}$ (Path)
 $\quad | M.X$ (Anonymous projection)
 $\quad | (\dot{P} : \dot{S})$ (Ascription)
 $\quad | \dot{P}()$ (Generative application)
 $\quad | () \rightarrow M$ (Generative functor)
 $\quad | (Y : \dot{S}) \rightarrow M$ (Applicative functor)
 $\quad | \text{struct}_A \bar{B} \text{ end}$ (Structure)

Binding

$B ::= \text{let } x = e$ (Value)
 $\quad | \text{type } t = \ddot{u}$ (Type)
 $\quad | \text{module } X = M$ (Module)
 $\quad | \text{module type } T = \dot{S}$ (Module type)

Core language types and expression

$e ::= \dots | Q.x$ (Qualified value)

Identifier

$I ::= x | t | X | T$

Typing environment

$\Gamma ::= \emptyset | \Gamma, A.D | \Gamma, Y : S$

Zipper context

$\gamma ::= \emptyset | A : \bar{D} | \gamma ; \gamma$

Signature

$S ::= \langle \gamma \rangle S$ (Zipper)
 $\quad | \dot{S}$ (Plain signature)
 $\dot{S} ::= (= P < \dot{S})$ (Transparent signature)
 $\quad | Q.T$ (Module type)
 $\quad | () \rightarrow S$ (Generative functor)
 $\quad | (Y : \dot{S}) \rightarrow S$ (Applicative functor)
 $\quad | \text{sig}_A \bar{D} \text{ end}$ (Structural signature)

Declaration

$D ::= \text{val } x : u$ (Value)
 $\quad | \text{type } t = u$ (Type)
 $\quad | \text{module } X : S$ (Module)
 $\quad | \text{module type } T = \dot{S}$ (Module type)

$u ::= \dots | Q.t$ (Qualified type)

Fig. 3. Syntax of ZIPML

letters (I, A, Q, \dots) for identifiers and paths, and upright letters (M, e, D, \dots) for syntactic categories. Finally, we designate lists with an overbar: \bar{B} is a list of B 's.

Main module constructs. The two main categories are **Module Expressions** (M) and **Signatures** (S). We put the content of structures and structural signatures, (expressions, types, modules, module types) in the separate categories of **Bindings** (B) and **Declarations** (D). As in OCAML we use a special unit argument to syntactically distinguish generative functors from applicative functors. Both structures and signatures are annotated with a *self-reference* A , explained below. Signatures contain transparent signatures ($= P < \dot{S}$) to express that a module has the identity of P but the interface of \dot{S} . It partially subsumes type-level module aliases of OCAML.

Self-references. A special class of identifiers A range over self-references, which are variables used in both module structures $\text{struct}_A \bar{B} \text{ end}$ and structural signatures $\text{sig}_A \bar{D} \text{ end}$ to refer to the structure or the signature itself from fields \bar{B} or \bar{D} . They are **not** used to define **recursive structures**, but only to refer to previously defined fields in a telescope. The subscript annotation $_A$ on a structure or a signature is a binding occurrence whose scope extends to \bar{B} or \bar{D} and that can be freely renamed. Thanks to self-references, field names are no longer binders and behave rather as record fields. We write $\text{dom}(D)$ and $\text{dom}(B)$ the field name I of D and B . We disallow

field shadowing in signatures, which is standard in module systems. By contrast with common usage, we also disallow field shadowing in structures: all field names in the *same* structure must be disjoint. However, fields names do not shadow other fields of the same name in a *substructure* (or *subsignature*) as these are accessed through their self-references which can be freely renamed to avoid clashes. This restriction of shadowing is for the sake of simplicity and is just a matter of name resolution that has little interaction with typechecking.

Self-references are also used to represent abstract types: all type fields are of the form $\text{type } t = u$ where u may be a core language type or an alias $P.t$ including the case $A.t$ where A is the self-reference of the field under consideration, which in this case means that $A.t$ is an abstract type. This is merely a syntactic trick to simplify the treatment of telescopes and strengthening. We may omit the self-reference and just write $\text{sig } D \text{ end}$ for $\text{sig}_A D \text{ end}$ when A does not appear free in D . Conversely, when we write $\text{sig } D \text{ end}$, we should read $\text{sig}_A D \text{ end}$ where the anonymous self-reference A is chosen fresh for D .

Identifiers and Paths. Paths P are the mechanism to access the content of modules. Paths may access the environment directly, either through a functor parameter Y or a field $A.I$, where I spans over any identifier. They may also access module fields by projection $P.X$. In the absence of applicative functors and floating fields, this would be sufficient. With applicative functors, a path may also designate the result of an immediate module application $P(P)$. Floating fields are accessed with $P.A$. The letter Q designates a distant access via a path P or a local access via a self-reference A . Finally, we extend the core language with qualified values $Q.x$ and qualified types $Q.t$.

Note that we distinguish module names X from module parameters Y . The latter are variables, can be renamed when in binding position and substituted during typechecking. By contrast, names are never used in a binding position.

Typing environments. Typing environments Γ bind module fields to declarations and functor parameters to signatures. Module fields are always prefixed by a self-reference in typing environments. As we disallow shadowing, $A.D$ can only be added to Γ when Γ does not already contain some $A.D'$ where D and D' define the same field. For convenience, we may write $\Gamma, A.\bar{D}$ for the sequence $\Gamma, \overline{A.D}$, where $\Gamma, (\overline{A.D'}, \overline{A.D})$ means $(\Gamma, \overline{A.D'}), \overline{A.D}$, adding fields one by one.

Zippers. Finally, the novelty is the introduction of *zippers* $\langle \gamma \rangle S$ where γ is a *zipper context*, i.e., a sequence of floating fields $A : \bar{D}$. The self-reference A , now used as a label to access fields of \bar{D} from S , **cannot be** freely **renamed** any longer¹³: the introduction of zippers turns an α -convertible self-reference into a fixed label. The concatenation of zippers $\gamma_1 ; \gamma_2$ is only well-formed when the domains, i.e., the set of self-references, of γ_1 and γ_2 are disjoint. We may then see a zipper as a map from self-references to declarations and define $\gamma(A)$ accordingly. The concatenation of zippers “;” is associative and the empty zipper \emptyset is a neutral element. We identify $\langle \emptyset \rangle S$ with S and $\langle \gamma_1 \rangle \langle \gamma_2 \rangle S$ with $\langle \gamma_1 ; \gamma_2 \rangle S$ whenever γ_1 and γ_2 have disjoint domains. Therefore, a signature S can always be written as $\langle \gamma \rangle \dot{S}$ where \dot{S} is a plain signature and γ concatenates all consecutive zipper contexts or is \emptyset if there were none. The introduction of zippers requires a new form of path $P.A$, invalid in source programs, to access floating fields.¹⁴

Zipper-free signatures. We define plain signatures \dot{S} as those without an initial zipper (but where subterms may contain zippers) and source signatures \ddot{S} as those that do not contain any zipper at all, which are used in source types appearing in source expressions, i.e., in a transparent ascription or the domain of a functor. Similarly, we let \ddot{P} stand for *source* paths P that do not contain any

¹³In section §4, we will allow for consistent renaming as an additional typing rule, but not as part of type equivalence.

¹⁴In the absence of floating fields, self references could only be at the origin of a path Q .

access to some zipper context (of the form $P'.A$). Consistently, \ddot{Q} means \ddot{P} or A and *source* types \ddot{u} means types u where all paths occurring in u are source paths.

Invariants. We also define several syntactic subcategories of signatures to capture some invariants.¹⁵ The head value form S^v of a signature gives the actual shape of a signature, which is either a structural signature or a functor. The head normal form S^n is similar, but still contains the *identity* of a signature, if it has one, via a transparent ascription.

$$S^v ::= \text{sig}_A \bar{D} \text{ end} \mid (Y : \ddot{S}) \rightarrow S \mid () \rightarrow S \quad S^n ::= S^v \mid (= P < S^v)$$

Notice that head normal forms and value forms are superficial and a signature appearing under a value form may itself be any signature and therefore contain inner zippers. A *Transparent* signature \mathbb{S} is a generalization of the syntactic form $(= P < \dot{S})$ that also allows zippers provided their signatures eventually start with transparent signatures. The definition is the following:

$$\begin{aligned} \mathbb{S} &::= \langle \gamma \rangle \mathbb{S} \mid \dot{\mathbb{S}} & \dot{\mathbb{S}} &::= (= P < \dot{S}) & \gamma &::= A : \bar{D} \mid \gamma ; \gamma \mid \emptyset \\ \mathbb{D} &::= \text{module } X : \mathbb{S} \mid \text{val } x : u \mid \text{module type } T = \dot{\mathbb{S}} \mid \text{type } t = u \end{aligned}$$

Syntactic choices and syntactic sugar. Our grammar has some superficial syntactic restrictions that simplify the presentation without reducing expressiveness. In particular, we only allow applications of paths to paths. The more general application $M_1(M_2)$ may be encoded as syntactic sugar for $(\text{struct}_A \text{ module } X_1 = M_1 \text{ module } X_2 = M_2 \text{ module } X = A.X_1(A.X_2) \text{ end}).X$. Indeed, our projection $M.X$ is unrestricted. By contrast, OCAML restricts M to be a path P and must encode general projection with an application. Our choice is more general, as it does not require any additional type annotation. Similarly, we restricted ascriptions to source paths, $(\ddot{P} : \ddot{S})$, but the general case can be encoded as $(\text{struct}_A \text{ module } X_1 = M \text{ module } X = (A.X_1 : \ddot{S}) \text{ end}).X$. Finally, note some grammatical ambiguity: $P.X$ may be a projection from a path P or from a module expression M which may itself be a path. This is not an issue as their typing will be the same.

3.2 Strengthening

Strengthening is a key operation in path-based module systems: intuitively, it is used to give module signatures and abstract types an identity that will then be preserved by aliasing. ZIPML reuses transparent ascription $(= P < S)$ to express strengthening of the signature S by the path P , which effectively gives an identity, i.e., the path P , to an existing signature S —under a subtyping condition between the signature of P and S . In other words, transparent ascription allows strengthening to be directly represented in the syntax of signatures. This is advantageously used to implement strengthening lazily, by contrast with OCAML's eager version.¹⁶

Given a signature S and a path P , we consider two forms of strengthening: delayed strengthening $S // P$ and shallow strengthening S / P , defined in Figure 4. Shallow strengthening is only defined on signatures in head normal form and is called during resolution to push strengthening just one level down. It then delegates the work to delayed strengthening $S // P$, which will insert a transparent ascription in a signature, if there is not one already, and push it under zippers if any. In particular, the shallow strengthening of a structural signature is a structural signature that does not use its self-reference any longer. Since the very purpose of strengthening is to make signatures transparent, both forms of strengthening indeed return a transparent signature \mathbb{S} . The rules for

¹⁵For sake of readability, we do not always use the most precise syntactic categories and sometimes just write S when S may actually be of a more specific shape. Conversely, we may use subcategories to restrict the application of a rule that only applies for signatures of a specific shape.

¹⁶There is actually a proposal to add lazy strengthening in OCAML for efficiency purposes, see <https://github.com/ocaml-flambda/flambda-backend/pull/1337>

Delayed strengthening (signatures)

$$\begin{aligned}
(= P' < S) // P &\triangleq (= P' < S) \\
\langle \gamma ; \gamma' \rangle S // P &\triangleq (\langle \gamma \rangle (\langle \gamma' \rangle S)) // P \\
\langle A : \bar{D} \rangle S // P &\triangleq \langle A : \bar{D}[A \leftarrow P.A] // P \rangle (S[A \leftarrow P.A] // P) \\
S // P &\triangleq (= P < S)
\end{aligned}$$

Shallow strengthening

$$\begin{aligned}
\text{sig}_A \bar{D} \text{ end} / P &\triangleq \text{sig } \overline{D[A \leftarrow P]} // P \text{ end} \\
(Y : \bar{S}) \rightarrow S / P &\triangleq (Y : \bar{S}) \rightarrow (S // P(Y)) \\
() \rightarrow S / P &\triangleq () \rightarrow S
\end{aligned}$$

Delayed strengthening (declarations)

$$\begin{aligned}
(\text{val } x : u) // Q &\triangleq \text{val } x : u & (\text{module } X : S) // Q &\triangleq \text{module } X : (S // Q.X) \\
(\text{type } t = u) // Q &\triangleq \text{type } t = u & (\text{module type } T = S) // Q &\triangleq \text{module type } T = S
\end{aligned}$$

Environment strengthening

$$\begin{aligned}
\Gamma \uplus (Y : S) &\triangleq \Gamma, Y : (S // Y) & \Gamma \uplus (A : \bar{D}; \gamma) &\triangleq (\Gamma \uplus A.\bar{D}) \uplus \gamma \\
\Gamma \uplus A.D &\triangleq \Gamma, A.(D // A) & \Gamma \uplus A.(D, \bar{D}) &\triangleq (\Gamma \uplus A.D) \uplus A.\bar{D}
\end{aligned}$$

Fig. 4. Strengthening (delayed by default) – S / P and $S // P$ and $D // Q$

delayed signature strengthening should be read in order of appearance, as they pattern match on the head of the signature:

- Delayed strengthening stops at a transparent ascription, since it is already transparent.
- It strengthens (zipped) signatures step by step. We decompose compound zippers as two successive zippers. For a simple zipper, we strengthen both the zipper context and the underlying signature in which we replaced the self-reference A by the strengthened path.
- Otherwise, delayed strengthening just inserts a transparent ascription, which is actually the materialization of the delaying.

We also defined delayed strengthening on declarations $D // Q$, which is called by strengthening on zipper contexts, shallow strengthening, and strengthening of the typing context: delayed strengthening is pushed inside module declarations and dropped on other fields. (Module type definitions never have an identity, so they cannot be strengthened.) Declaration strengthening $D // A$ may be called on a zipper self-variable A , but it will then immediately be dropped or expand to strengthening by a path $S // A.X$. We also define a binary operator \uplus that strengthens bindings as they enter the typing environment. We use it to maintain the invariant that all signatures entering the typing environment are transparent signatures \bar{S} , so that they can be duplicated without loss of sharing.

3.3 Path typing and normalization

Since paths include projections and applications, which require recursive lookups and substitutions, their types are not immediate to deduce. Moreover, signatures in the typing environment may themselves be module type definitions that must be inlined to be analyzed. We use path resolution, path typing, and normalization for this purpose.

Path typing and path resolution. Intuitively, typing a path $\Gamma \vdash P : \bar{S}$ returns *all* the environment information about the module at path P , including a potential zipper, an aliasing information with another path (or itself), and the plain signature. To factor out a common pattern-match on the result of path typing, we also introduce *path resolution* $\Gamma \vdash P \triangleright S$ to extract the module *content*, by dropping the zipper and aliasing information and forcing a shallow strengthening, and $\Gamma \vdash P \triangleright P'$ to extract the module *identity*, by dropping the zipper and signature.

3.3.1 $\boxed{\Gamma \vdash P \triangleright P'}$ and $\boxed{\Gamma \vdash P \triangleright S}$ – *Path resolution.* The two judgments are defined in Figure 5 by a single rule each. Rule **RES-P-ID** simply pattern-matches on the result of signature typing to return

$$\frac{\text{RES-P-ID} \quad \Gamma \vdash P : \langle \gamma \rangle (= P' < \dot{S})}{\Gamma \vdash P \triangleright P'} \quad \frac{\text{RES-P-VAL} \quad \Gamma \vdash P : \langle \gamma \rangle (= P' < \dot{S}) \quad \dot{S} / P' = S'}{\Gamma \vdash P \triangleright S'}$$

Fig. 5. Path Resolution – $\Gamma \vdash P \triangleright P'$ and $\Gamma \vdash P \triangleright \dot{S}$

$$\frac{\text{TYP-P-ARG} \quad Y : \mathbb{S} \in \Gamma}{\Gamma \vdash Y : \mathbb{S}} \quad \frac{\text{TYP-P-MODULE} \quad A. \text{module } X : \mathbb{S} \in \Gamma}{\Gamma \vdash A.X : \mathbb{S}} \quad \frac{\text{TYP-P-ZIP} \quad \Gamma \vdash P : \langle \gamma \rangle \mathbb{S} \quad \gamma(A) = \overline{\mathbb{D}}}{\Gamma \vdash P.A : \text{sig } \overline{\mathbb{D}} \text{ end}} \quad \frac{\text{TYP-P-NORM} \quad \Gamma \vdash P : \mathbb{S} \quad \Gamma \vdash \mathbb{S} \downarrow \mathbb{S}'}{\Gamma \vdash P : \mathbb{S}'}$$

$$\frac{\text{TYP-P-PROJ} \quad \Gamma \vdash P \triangleright \text{sig } \overline{\mathbb{D}} \text{ end} \quad \text{module } X : \mathbb{S} \in \overline{\mathbb{D}}}{\Gamma \vdash P.X : \mathbb{S}} \quad \frac{\text{TYP-P-APPA} \quad \Gamma \vdash P \triangleright (Y : \dot{S}_0) \rightarrow \mathbb{S} \quad \Gamma \vdash P' : \mathbb{S}' \quad \Gamma \vdash \mathbb{S}' \leq \dot{S}_0}{\Gamma \vdash P(P') : \mathbb{S}[Y \leftarrow P']}$$

Fig. 6. Path typing – $\Gamma \vdash P : \mathbb{S}$

$$\frac{\text{NORM-S-ZIP} \quad \Gamma \uplus \gamma \vdash S \downarrow S'}{\Gamma \vdash \langle \gamma \rangle S \downarrow \langle \gamma \rangle S'} \quad \frac{\text{NORM-S-TRANS-SOME} \quad \Gamma \vdash S \downarrow (= P' < S')}{\Gamma \vdash (= P < S) \downarrow (= P' < S')}$$

$$\frac{\text{NORM-S-TRANS-NONE} \quad \Gamma \vdash P \triangleright P' \quad \Gamma \vdash S \downarrow S'}{\Gamma \vdash (= P < S) \downarrow (= P' < S')} \quad \frac{\text{NORM-TYP-RES} \quad \Gamma \vdash P \triangleright P'}{\Gamma \vdash P.t \downarrow P'.t}$$

$$\frac{\text{NORM-S-LOCALMODTYPE} \quad \text{module type } A.T = S \in \Gamma}{\Gamma \vdash A.T \downarrow S} \quad \frac{\text{NORM-S-PATHMODTYPE} \quad \Gamma \vdash P \triangleright \text{sig } \overline{\mathbb{D}} \text{ end} \quad \text{module type } T = S \in \overline{\mathbb{D}}}{\Gamma \vdash P.T \downarrow S}$$

$$\frac{\text{NORM-TYP-LOCAL} \quad \text{type } A.t = u \in \Gamma}{\Gamma \vdash A.t \downarrow u} \quad \frac{\text{NORM-TYP-PATH} \quad \Gamma \vdash P \triangleright \text{sig } \overline{\mathbb{D}} \text{ end} \quad \text{type } t = u \in \overline{\mathbb{D}}}{\Gamma \vdash P.t \downarrow u}$$

Fig. 7. Signature and type normalization – $\Gamma \vdash S \downarrow S'$

the aliasing information. Rule **RES-P-VAL** pattern-matches on the result of signature typing both to extract \dot{S} and to force it to be in head-normal (so that \dot{S} / P' is well-defined).

3.3.2 $\boxed{\Gamma \vdash P : \mathbb{S}}$ – *Path typing*. The judgment is defined in Figure 6. Rules **TYP-P-ARG** and **TYP-P-MODULE** are straightforward lookups. Rule **TYP-P-ZIP** accesses a zipper context through its self-reference. Notice that the signature $\text{sig}_A \overline{\mathbb{D}} \text{ end}$ of $P.A$ need not be put back inside the zipper context γ , since the signature $\langle \gamma \rangle \mathbb{S}$, hence the declarations $\overline{\mathbb{D}}$, are transparent and no longer depend on the zipper context γ , but only on the environment Γ . Rule **TYP-P-NORM** means that path typing is defined up to signature normalization, which allows the inlining of module type definitions. This makes path typing non-deterministic, purposefully. The two remaining rules use path resolution to pattern-match on the content of a signature. In Rule **TYP-P-PROJ**, we omitted the self-reference of the signature of P since we know P is transparent, hence \mathbb{S} does not use its self-reference and is well-formed in Γ . Typing of functors (Rule **TYP-P-APPA**) requires the domain signature to be a super type of the argument signature.¹⁷ We then return the codomain signature after substitution of the argument P for the parameter Y .

The dependency between path typing and subtyping (the premise $\Gamma \vdash \mathbb{S}' \leq \dot{S}_0$ in **TYP-P-APPA**) is somewhat artificial. It comes from the fact that a single judgment is used for both *path lookups* and *path checking*. We could have (and an efficient implementation *would*) used two separate judgments, where path-lookup, only used on valid paths, would not recheck subtyping.

3.3.3 $\boxed{\Gamma \vdash S \downarrow S'}$ – *Normalization*. The judgment $\Gamma \vdash S \downarrow S'$ allows the (head) normalization of a signature S into S' , which inlines module type definitions, but only one step at a time. Hence, to achieve the head normal form, we may call normalization repeatedly. Rule **NORM-S-ZIP** normalizes the signature part of a zipper. We never normalize the zipper context itself, as we will first access the zipper context and normalize the result afterwards. Rules **NORM-S-TRANS-SOME** and **NORM-S-TRANS-NONE** allow normalization under a transparent ascription. If the head of S is itself a transparent

¹⁷The two first premises ensure that signatures \mathbb{S}' and S are well-typed as required when using the subtyping judgment.

$\frac{\text{SUB-S-SIG} \quad \bar{D}_0 \sqsubset \bar{D}_1 \quad \Gamma \uplus A : \bar{D}_1 \vdash \bar{D}_0 // A \leq \bar{D}_2}{\Gamma \vdash \text{sig}_A \bar{D}_1 \text{ end} \leq \text{sig}_A \bar{D}_2 \text{ end}}$		$\frac{\text{SUB-S-NORM} \quad \Gamma \vdash S_1 \downarrow S'_1 \quad \Gamma \vdash S_2 \downarrow S'_2 \quad \Gamma \vdash S'_1 \leq S'_2}{\Gamma \vdash S_1 \leq S_2}$	
$\frac{\text{SUB-S-ZIPPERL} \quad \Gamma \uplus \gamma \vdash S_1 \leq S_2}{\Gamma \vdash \langle \gamma \rangle S_1 \leq S_2}$	$\frac{\text{SUB-S-TRASCR} \quad \Gamma \vdash S_1 / P \leq S_2 / P}{\Gamma \vdash (= P < S_1) \leq (= P < S_2)}$	$\frac{\text{SUB-S-LOOSEALIAS} \quad \Gamma \vdash S_1 / P \leq S_2}{\Gamma \vdash (= P < S_1) \leq S_2}$	$\frac{\text{SUB-S-FCTG} \quad \Gamma \vdash S_1 \leq S_2}{\Gamma \vdash () \rightarrow S_1 \leq () \rightarrow S_2}$
$\frac{\text{SUB-S-FCTA} \quad \Gamma \vdash \check{S}_2 \leq \check{S}_1 \quad \Gamma \uplus Y : \check{S}_2 \vdash S'_1 \leq S'_2}{\Gamma \vdash (Y : \check{S}_1) \rightarrow S'_1 \leq (Y : \check{S}_2) \rightarrow S'_2}$		$\frac{\text{SUB-T-NORM} \quad \Gamma \vdash u_1 \downarrow u \quad \Gamma \vdash u_2 \downarrow u}{\Gamma \vdash u_1 \leq u_2}$	
$\frac{\text{SUB-D-MOD} \quad \Gamma \vdash S_1 \leq S_2}{\Gamma \vdash \text{module } X : S_1 \leq \text{module } X : S_2}$		$\frac{\text{SUB-D-MODTYPE} \quad \Gamma \vdash \check{S}_1 \leq \check{S}_2 \quad \Gamma \vdash \check{S}_2 \leq \check{S}_1}{\Gamma \vdash \text{module type } T = \check{S}_1 \leq \text{module type } T = \check{S}_2}$	
$\frac{\text{SUB-D-TYPE} \quad \Gamma \vdash u_1 \approx u_2}{\Gamma \vdash \text{type } t = u_1 \leq \text{type } t = u_2}$		$\frac{\text{SUB-S-DYNEQ} \quad \Gamma \vdash S_1 \leq S_2}{\Gamma \vdash S_1 \approx S_2}$	

Fig. 8. Subtyping rules $\mathcal{R}_{\leq}^{\sqsubset, \check{\cdot}}$ for code-free subtyping $\Gamma \vdash S_1 \leq S_2$ and coercion subtyping $\Gamma \vdash S_1 \leq S_2$

ascription, we return it as is; otherwise, we return its strengthened version by the resolved path P' . Finally, rules **NORM-S-LOCALMODTYPE** and **NORM-S-PATHMODTYPE** expand a module type definition. We also define $\Gamma \vdash P.t \downarrow u$, the normalization of types. Rules **NORM-TYP-RES** allows the resolution of the path P while rules **NORM-TYP-LOCAL** and **NORM-TYP-PATH** inlined the type definition $Q.t$.

3.4 Subtyping judgments

$$\boxed{\Gamma \vdash S_1 \leq S_2}$$

We define two subtyping judgments, *code-free* subtyping $\Gamma \vdash S \leq \check{S}$ and *coercion* subtyping $\Gamma \vdash S \leq \check{S}$. The latter can only be used at type ascription and functor applications, as it requires changing the representation of the underlying values.

To factor both definitions, we defined a set of subtyping rules $\mathcal{R}_{\leq}^{\sqsubset, \check{\cdot}}$ parameterized by three relations in Figure 8. The key rule, which is also the main difference between \leq and \leq is **SUB-S-SIG**. It uses a binary relation \sqsubset to tell which components can be dropped or reordered when subtyping between two structural signatures $\text{sig}_A \bar{D}_1 \text{ end}$ and $\text{sig}_A \bar{D}_2 \text{ end}$. Namely, we must find a sequence \bar{D}_0 related to \bar{D}_1 by \sqsubset that is, after strengthening by A (to keep all sharing), in pointwise \leq -subtyping relation with \bar{D}_2 in an environment extended with $A.\bar{D}_1$. We use two versions of \sqsubset : for coercion subtyping, we take \sqsubset , i.e., \bar{D}_0 can be any subset of \bar{D}_1 where fields may appear in a different order; for code-free subtyping, we use the relation \sqsubset that is the subrelation of \sqsubset that preserves dynamic fields (modules and values) and their order. Formally, $\bar{D}_0 \sqsubset \bar{D}_1$ means both $\bar{D}_0 \sqsubset \bar{D}_1$ and $\text{dyn}(\bar{D}_0) = \text{dyn}(\bar{D}_1)$ where $\text{dyn}(\bar{D})$ returns the subsequence of \bar{D} composed of dynamic fields only.

The subtyping relations are defined in two steps: we first define code-free subtyping \leq as the smallest relation that satisfies the rules $\mathcal{R}_{\leq}^{\sqsubset, \check{\cdot}}$. Once \leq is defined and fixed, we then define coercion subtyping \leq as the smallest relation that satisfies the rules $\mathcal{R}_{\leq}^{\sqsubset, \check{\cdot}}$. Both subtyping judgments are of the form $\Gamma \vdash S_1 \leq S_2$ and, for now, only defined when S_2 is a zipper-free signature \check{S}_2 —we relax this condition in §4. The judgment should (and will) only be used when both signatures are well-typed. Typically, the right-hand side signature is a source signature while the left-hand side signature S results from the typing of either a source signature or a module expression—and may contain zippers. The same invariants extend to auxiliary subtyping judgments for declarations and paths.

When reading the subtyping rules, it may help to think of coercion subtyping first, i.e., reading of \leq as \leq and \leq as \leq . The subtyping rules for signatures can be read by case analysis on the left-hand-side signature, except for Rule **SUB-S-NORM**, which is not syntax directed and can be applied anywhere and repeatedly to perform normalization before subtyping. For example, there

is no rule matching a signature $Q.T$ on the left-hand side. But normalization allows inlining the definition before checking for subtyping. Since a zipper may only occur on the left-hand side, Rule `SUB-S-ZIPPERL` just pushes the zipper context in the typing environment and pursues with subtyping. For other cases, we require the left-hand side to be in head normal form, which can be achieved by Rule `SUB-S-NORM`. When the left-hand side is a transparent ascription the right-hand side may also be a transparent ascription, in which case, we check subtyping between the respective signatures, but after pushing strengthening one level-down, lazily (`SUB-S-TRASCR`). Otherwise, we drop the transparent ascription from the left-hand side, which amounts to a loss of transparency, hence increase abstraction, as allowed by subtyping (`SUB-S-LOOSEALIAS`). In the remaining cases, the left-hand side is a head value form and the right-hand side must have the same shape. Functor types are contravariant (rules `SUB-S-FCTG` and `SUB-S-FCTA`). Finally, Rule `SUB-S-DYNEQ` is just used to define \approx on signatures as the kernel of \lesssim .

Subtyping uses two other helper judgments, for type and declaration subtyping. There is a single rule `SUB-T-NORM` for type subtyping that injects head type normalization into the subtyping relation, which is the analogue of Rule `SUB-S-NORM` for the normalization of core-language types. In fact, this rule should also be made available in the subtyping relation of the core language, which should be a congruent preorder. For module declarations (Rule `SUB-D-MOD`), we just require subtyping covariantly. Rule `SUB-D-VAL` for core language values is similar, requiring (code-free) subtyping in the core language. In OCAML, this would reduce to core-language type-scheme specialization, which we have not formalized in ZIPML. Since module types may be used in both covariant and contravariant positions, the rule `SUB-D-MODTYPE` requests code-free subtyping in both directions, i.e., type equivalence, which is well-defined since module type definitions are source signatures.

Subtyping and well-typedness. Subtyping is only meant to be well-behaved on well-typed signatures: it is the caller's responsibility to ensure that both signatures are well-formed.

Optimization. Judgments $\Gamma \vdash S_1 / P \leq S_2$ and $\Gamma \vdash S_1 / P \leq S_2 / P$ are in fact equivalent. Intuitively, a derivation of the former may abstract some types appearing in S_1 / P , but never has to, i.e., the same derivation could be reproduced without any abstraction. Therefore, Rule `SUB-S-LOOSEALIAS` could be replaced by the optimized rule `SUB-S-LOOSEALIASO` below.

$$\frac{\text{SUB-S-LOOSEALIASO} \quad \Gamma \vdash S_1 / P \leq S_2 / P}{\Gamma \vdash (= P < S_1) \leq S_2} \qquad \frac{\text{SUB-S-SIGO} \quad \bar{D}_0 \subseteq \bar{D}_1 \quad \Gamma \vdash \bar{D}_0 \leq \bar{D}_2}{\Gamma \vdash \text{sig } \bar{D}_1 \text{ end} \leq \text{sig } \bar{D}_2 \text{ end}}$$

We may then use Rule `SUB-S-SIGO`, an instance of `SUB-S-SIG`, when neither side uses its self-reference, so as to avoid adding useless information to the typing environment.

3.5 Signature typing

 $\boxed{\Gamma \vdash S}$

User-provided source signatures \check{S} are not necessarily well-formed and are checked using the judgment $\Gamma \vdash S$, defined on Figure 9. We still define this judgment on signatures with zippers, as we also use it for well-formedness at the meta-theoretical level.

Rule `TYP-S-MODTYPE` uses path typing to check the well-formedness of paths. Rule `TYP-S-ASCR` must also check that the signature S' of path P is a subtype of the source signature S .¹⁸ Rule `TYP-S-STR` for structural signatures delays most of the work to the elaboration judgment $\Gamma \vdash_A \bar{D}$ for declarations, which carries the self-reference variable A that should be chosen fresh for Γ , as it now appears free in declarations \bar{D} . Rule `TYP-D-SEQ` for sequence of declarations pushes $A.D_0$ in the context while typing the remaining sequence \bar{D} . Rule `TYP-S-ZIPPER` is a combination of `TYP-S-STR` and `TYP-D-SEQ`. All the other rules are straightforward. In practice, typing of signatures could simplify

¹⁸The two first premises ensure that signatures S and S' are well-typed, as required when using the subtyping judgment.

$\frac{\text{TYP-S-MODTYPE} \quad \Gamma \vdash Q.T : S}{\Gamma \vdash Q.T}$	$\frac{\text{TYP-S-GENFCT} \quad \Gamma \vdash S}{\Gamma \vdash () \rightarrow S}$	$\frac{\text{TYP-S-APPFCT} \quad \Gamma \vdash \check{S} \quad \Gamma \uplus (Y : \check{S}) \vdash S'}{\Gamma \vdash (Y : \check{S}) \rightarrow S'}$	$\frac{\text{TYP-S-ASCR} \quad \Gamma \vdash S \quad \Gamma \vdash P : S' \quad \Gamma \vdash S' \leq S}{\Gamma \vdash (= P < S)}$	
$\frac{\text{TYP-S-STR} \quad \Gamma \vdash_A \bar{D} \quad A \notin \Gamma}{\Gamma \vdash \text{sig}_A \bar{D} \text{ end}}$	$\frac{\text{TYP-S-ZIPPER} \quad \Gamma \vdash_A \bar{D} \quad \Gamma \uplus A : \bar{D} \vdash S}{\Gamma \vdash \langle A : \bar{D} \rangle S}$	$\frac{\text{TYP-D-VAL} \quad \Gamma \vdash u}{\Gamma \vdash_A (\text{val } x : u)}$	$\frac{\text{TYP-D-TYPE} \quad \Gamma \vdash u}{\Gamma \vdash_A (\text{type } t = u)}$	$\frac{\text{TYP-D-NIL}}{\Gamma \vdash_A \emptyset}$
$\frac{\text{TYP-D-TYPEABS}}{\Gamma \vdash_A (\text{type } t = A.t)}$	$\frac{\text{TYP-D-MOD} \quad \Gamma \vdash S}{\Gamma \vdash_A (\text{module } X : S)}$	$\frac{\text{TYP-D-MODTYPE} \quad \Gamma \vdash S}{\Gamma \vdash_A (\text{module type } T = S)}$	$\frac{\text{TYP-D-SEQ} \quad \Gamma \vdash_A D_0 \quad \Gamma \uplus A.D_0 \vdash_A \bar{D}}{\Gamma \vdash_A (D_0, \bar{D})}$	

Fig. 9. Signature typing – $\Gamma \vdash S$

$\frac{\text{TYP-M-NORM} \quad \Gamma \vdash^\diamond M : S \quad \Gamma \vdash S \downarrow S'}{\Gamma \vdash^\diamond M : S'}$	$\frac{\text{TYP-M-MODE} \quad \Gamma \vdash^\nabla M : S}{\Gamma \vdash^\nabla M : S}$	$\frac{\text{TYP-M-PATH} \quad \Gamma \vdash \check{P} : \langle \gamma \rangle (= P' < \check{S})}{\Gamma \vdash^\diamond \check{P} : (= \check{P} < \check{S})}$	$\frac{\text{TYP-M-ASCR} \quad \Gamma \vdash \check{S} \quad \Gamma \vdash \check{P} : \check{S} \quad \Gamma \vdash \check{S} \leq \check{S}}{\Gamma \vdash^\diamond (\check{P} : \check{S}) : \check{S}}$
$\frac{\text{TYP-M-FCTA} \quad \Gamma \vdash \check{S} \quad \Gamma \uplus Y : \check{S} \vdash^\nabla M : S'}{\Gamma \vdash^\diamond (Y : \check{S}) \rightarrow M : (Y : \check{S}) \rightarrow S'}$	$\frac{\text{TYP-M-FCTG} \quad \Gamma \vdash^\nabla M : S}{\Gamma \vdash^\diamond () \rightarrow M : () \rightarrow S}$	$\frac{\text{TYP-M-APPG} \quad \Gamma \vdash \check{P} \triangleright () \rightarrow S}{\Gamma \vdash^\nabla \check{P}() : S}$	$\frac{\text{TYP-M-STR} \quad \Gamma \vdash_A \bar{B} : \bar{D} \quad A \notin \Gamma}{\Gamma \vdash \text{struct}_A \bar{B} \text{ end} : \text{sig}_A \bar{D} \text{ end}}$
$\frac{\text{TYP-M-PROJ}\Gamma \quad \Gamma \vdash^\diamond M : \langle \gamma \rangle (= P < \text{sig}_A \bar{D}, \text{module } X : S, \bar{D}' \text{ end})}{\Gamma \vdash^\diamond M.X : \langle \gamma \rangle (= P.X < S[A \leftarrow P])}$		$\frac{\text{TYP-M-PROJ}A \quad \Gamma \vdash^\diamond M : \langle \gamma \rangle \text{sig}_A \bar{D}, \text{module } X : S, \bar{D}' \text{ end}}{\Gamma \vdash^\diamond M.X : \langle \gamma ; A : \bar{D} \rangle S}$	
$\frac{\text{TYP-B-SEQ} \quad \Gamma \vdash_A^\diamond B_0 : D_0 \quad \Gamma \uplus A.D_0 \vdash_A^\diamond \bar{B} : \bar{D}}{\Gamma \vdash_A^\diamond B_0, \bar{B} : D_0, \bar{D}}$	$\frac{\text{TYP-B-TYPE-BIND} \quad \Gamma \vdash \check{u}}{\Gamma \vdash_A^\diamond (\text{type } t = \check{u}) : (\text{type } t = \check{u})}$	$\frac{\text{TYP-B-EMPTY}}{\Gamma \vdash_A^\diamond \emptyset : \emptyset}$	
$\frac{\text{TYP-B-ABSType} \quad \Gamma \vdash_A^\diamond (\text{type } t = A.t) : (\text{type } t = A.t)}$	$\frac{\text{TYP-B-LET} \quad \Gamma \vdash^\diamond e : u}{\Gamma \vdash_A^\diamond (\text{let } x = e) : (\text{val } x : u)}$	$\frac{\text{TYP-B-MOD} \quad \Gamma \vdash^\diamond M : S}{\Gamma \vdash_A^\diamond (\text{module } X = M) : (\text{module } X : S)}$	$\frac{\text{TYP-B-MODTYPE} \quad \Gamma \vdash \check{S}}{\Gamma \vdash_A^\diamond (\text{module type } T = \check{S}) : (\text{module type } T = \check{S})}$

Fig. 10. Typing rules for modules – $\Gamma \vdash^\diamond M : S$

them on the fly, typically removing chains of transparent ascriptions if any. This would then require replacing the typing judgment $\Gamma \vdash S$ by an elaboration judgment¹⁹.

3.6 Module typing

 $\Gamma \vdash^\diamond M : S$

The typing judgment $\Gamma \vdash^\diamond M : S$ for module expressions is given in Figure 10. The \diamond symbol is a metavariable for modes that ranges over the applicative (or *transparent*) mode ∇ and the generative (or *opaque*) mode ∇ . The floating rule **TYP-M-MODE** allows us to consider an applicative judgment as a generative one at any time. Judgments for pure module expressions can be treated either as applicative or generative, hence they use the \diamond metavariable. Many rules use the same metavariable \diamond in premises and conclusion, which then stands for the same mode. This implies that if the premise can only be proved in generative mode, it will also be the case for the conclusion.

When considering a source path \check{P} as a module expression (Rule **TYP-M-PATH**), path typing returns a transparent signature $\langle \gamma \rangle (= P' < \check{S})$. However, we return $(= \check{P} < \check{S})$ with the most recent identity \check{P} , as the older identity P' can always be recovered by normalization and \check{P} is probably the one the user wants to see. Besides, we drop the transparent zipper-context γ , which is useless as

¹⁹This would also be necessary if we allowed declarations `open S` and `include S` that should always be elaborated. We have not included them, but the type system has been designed to allow them.

the information is already stored in Γ . We could even go as far as removing all the zippers in the *accessible part* of the signature, but it might require pushing the strengthening and cause inlining.

A signature ascription ($\ddot{P} : \ddot{S}$) has the signature \ddot{S} provided it is indeed a supertype²⁰ of the signature \mathbb{S} of \ddot{P} (Rule **TYP-M-ASCR**). This ascription is opaque²¹ since it returns the signature \ddot{S} which is not strengthened by \ddot{P} . Since subtyping is not code free, it is not present as a floating subtyping rule but only allowed here and at functor applications.

A generative functor **TYP-M-FCTG** is just an evaluation barrier: the functor itself is applicative while the body is generative. Correspondingly, applying a generative functor, which amounts to evaluating its body, is then generative. An applicative functor is typed in the obvious way.

Rule **TYP-M-STR** for structures delays the work to the typing rules for bindings, which carry the self-variable A of the structure as an annotation that should be chosen fresh for the environment Γ . Typing rules for bindings are as expected. In particular, Rule **TYP-B-SEQ** pushes the declaration $A.D_0$ into the context while typing declarations \bar{D} , much as **TYP-D-SEQ** for signatures.

Finally, we have two rules for projection. Rule **TYP-M-PROJT** projects on a module that is a known alias of P . Therefore, we do not need to introduce a new zipper, as we can *strengthen* the resulting signature to mention P instead of A , removing dependencies with \bar{D} which are then dropped. By contrast, Rule **TYP-M-PROJA** is the key rule that leverages zippers. Intuitively, it just returns the signature S' of the field X of the signature S of M zipped around the initial fields \bar{D} of S appearing before the field X . Still, we have to consider that the signature of M may itself be in a zipper context γ , which is then composed with the zipper context formed of the initial fields \bar{D} , resulting in $\gamma ; A : \bar{D}$. As we pattern-match on the signature, this might require normalization. Importantly, the name A becomes fixed during projection, and is no longer freely α -convertible afterwards.

4 Resolving signature avoidance by zipper simplification

So far, zippers allow us to *delay* signature avoidance, but are sometimes polluting the inferred signatures. Indeed, zippers may be removed by ascription, but they are not simplified otherwise, even if they became useless. Yet, floating fields are not present at runtime and should only maintain the minimal amount of type information needed for a signature.

In this section, we address this issue by extending the definition of code-free subtyping \lesssim to enable simplification of zippers. In particular, as we are interested in *removing* floating fields whenever possible, we define *signature simplification* $\Gamma \vdash S_1 \rightsquigarrow S_2$ as a restricted form of subtyping that preserves all *visible* type equalities. This simplification is our actual *solution* to the signature avoidance problem, which we inject into module typing with Rule **TYP-M-SIMP**. Simplifying along subtyping ensures that it is always sound to simplify. We also show that we never do simplifications by over-abstracting or other transformations that could prevent further typings, by contrast with OCAML. For that purpose, we introduce *narrow subtyping* \lesssim_n , a restriction of subtyping that has this property and show that it contains the simplification relation \rightsquigarrow . While the primary goal of zipper simplification is to eliminate (or reduce) floating fields in inferred types, it also helps print simpler error messages. Early simplification, especially immediately after Rule **TYP-M-PROJA**, may also speed up typechecking.

$$\frac{\text{TYP-M-SIMP} \quad \Gamma \vdash M : S \quad \Gamma \vdash S \rightsquigarrow S'}{\Gamma \vdash M : S'}$$

Overview. We start with the extension of subtyping in §4.1, where we allow subtyping with zipper signatures on the right-hand side and define narrow subtyping. In §4.2, we present a set of simplification rules, included in narrow subtyping. In §4.3, we give an algorithm that computes an

²⁰The two first premises ensure that signatures \mathbb{S} and \ddot{S} are both well-typed, as required when using subtyping.

²¹We can still use this construct to implement transparent ascription ($\ddot{P} < \ddot{S}$) as syntactic sugar for $(\ddot{P} : (= \ddot{P} < \ddot{S}))$, which then returns the view \ddot{S} but with the identity of \ddot{P} .

iteration of these simplification rules without revisiting the signature multiple times. In §4.4, we show that narrow subtyping preserves typability.

4.1 Subtyping and narrow subtyping with zippers

$$\Gamma \vdash S \lesssim S'$$

We first extend subtyping to support zippers on the right-hand side, i.e., allow judgments of the form $\Gamma \vdash \langle A : \bar{D}_1 \rangle S_1 \leq \langle A : \bar{D}_2 \rangle S_2$. Even though we are mainly interested in code-free subtyping \lesssim , we simultaneously extend \leq so that it still contains the former. Hence, we do it by enlarging the parametric relation $\mathcal{R}_{\lesssim}^{\subseteq, \lesssim}$ with one new rule, **SUB-S-ZIPPER** given below.

4.1.1 Subtyping between zippers. A zipper signature $\langle A : \bar{D} \rangle S$ represents the projection on a field \mathbb{Z} of a signature of the form $\text{sig}_A \bar{D}, \text{module } \mathbb{Z} : S \text{ end}$. Therefore, subtyping between zippers could be defined by **SUB-S-ZIPPERWEAK**. While sound, this would be weaker than necessary for code-free subtyping \lesssim , as it would prevent reordering and deletion of floating fields. Since they are not present at run-time, reordering or deleting them is actually code free. Therefore, after expanding and simplifying the premises and using \subseteq for floating fields, we get the stronger²² rule **SUB-S-ZIPPER**.

$$\begin{array}{c} \text{SUB-S-ZIPPERWEAK} \\ \Gamma \vdash \text{sig}_A \bar{D}_1, \text{module } \mathbb{Z} : S_1 \text{ end} \\ \leq \text{sig}_A \bar{D}_2, \text{module } \mathbb{Z} : S_2 \text{ end} \\ \hline \Gamma \vdash \langle A : \bar{D}_1 \rangle S_1 \leq \langle A : \bar{D}_2 \rangle S_2 \\ \text{SUB-S-ZIPPER} \\ \bar{D}'_1 \subseteq \bar{D}_1 \\ \Gamma \uplus A : \bar{D}_1 \vdash \bar{D}'_1 // A \leq \bar{D}_2 \\ \Gamma \uplus A : \bar{D}_1 \vdash S_1 \leq S_2 \\ \hline \Gamma \vdash \langle A : \bar{D}_1 \rangle S_1 \leq \langle A : \bar{D}_2 \rangle S_2 \end{array}$$

4.1.2 Narrow subtyping. Code-free subtyping allows for over-abstractation, which we want to prevent, while equivalence prevents the removal of floating fields. We thus define *narrow subtyping* \lesssim as a *restriction* of code-free subtyping that prevents the loss of visible type equalities by the set of rules $\mathcal{R}_{\lesssim}^{\subseteq, \lesssim}$ except for **SUB-S-ZIPPER** and **SUB-S-SIG** which have been replaced by the following variants, restricted by additional premises:

$$\begin{array}{c} \text{SUB-S-SIGR} \\ \Gamma \uplus A : \bar{D}_1 \vdash \bar{D}_1 // A \lesssim \bar{D}_2 \\ \Gamma \uplus A : \bar{D}_2 \vdash \bar{D}_2 // A \lesssim \bar{D}_1 \\ \hline \Gamma \vdash \text{sig}_A \bar{D}_1 \text{ end} \lesssim \text{sig}_A \bar{D}_2 \text{ end} \\ \text{SUB-S-ZIPPERR} \\ \Gamma_1 = \Gamma \uplus A : \bar{D}_1, \text{module } \mathbb{Z} : S_1 \quad \Gamma_2 = \Gamma \uplus A : \bar{D}_2, \text{module } \mathbb{Z} : S_2, \bar{D}'_2 \\ \bar{D}'_1 \subseteq \bar{D}_1 \quad \Gamma_1 \vdash \bar{D}'_1 // A \lesssim \bar{D}_2, \bar{D}'_2 \quad \Gamma_1 \vdash S_1 \lesssim S_2 \\ \bar{D}'_2 \subseteq \bar{D}_2, \bar{D}'_2 \quad \Gamma_2 \vdash \bar{D}'_2 // A \lesssim \bar{D}_1 \quad \Gamma_2 \vdash S_2 // A.\mathbb{Z} \lesssim S_1 \\ \hline \Gamma \vdash \langle A : \bar{D}_1 \rangle S_1 \lesssim \langle A : \bar{D}_2 \rangle S_2 \end{array}$$

Rule **SUB-S-SIGR** for signatures disallow the removal and reordering of fields and requires field-by-field subtyping in both directions. Rule **SUB-S-ZIPPERR** still allows for \bar{D}_2 to contain fewer floating fields than \bar{D}_1 and appearing in a different order as Rule **SUB-S-ZIPPER** (premises on the second line). However, it also requires (premises on the last line) the converse subtyping up to the reintroduction of some floating fields \bar{D}'_2 that have been dropped from \bar{D}_1 . As will appear in §4.4, narrow subtyping is somehow the composition of an equivalence, preventing over-abstractation and reversible, with a projection, removing unreachable floating fields and irreversible.

4.2 Simplification

$$\Gamma \vdash S \rightsquigarrow S'$$

In this section, we present a set of *simplification rules* that transform a signature along the narrow subtyping relation \lesssim and can be used to reduce the size of the zipper. When no floating fields remain, we can remove the zipper altogether and this coincides with solving signature avoidance. Otherwise, simplification is partial, still delaying signature avoidance. At a high level, the simplification follows three elementary rules (*drop*, *move*, *split*), as it tries to simplify the right-most floating component, and one rule to reorganize zippers (*skip*). In this section, we consider a zipper signature $\langle A : D \rangle S$ where I is the identifier of D . The rules are given in Figure 11 and discussed below.

²²The second premise could even be $\Gamma \uplus A : \bar{D}_1 \vdash \bar{D}_0 // A \leq \bar{D}_2$, possibly allow coercion subtyping between floating fields, but we prefer to preserve them and only drop floating fields, as it suffices for simplifications.

$$\begin{array}{c}
\text{SIMP-EMPTY} \\
\Gamma \vdash \langle A : \emptyset \rangle S \rightsquigarrow S \\
\\
\text{SIMP-DROP} \\
\frac{I = \text{dom}(D_0) \quad A.I \notin \text{fv}(S)}{\Gamma \vdash \langle A : \bar{D}, D_0 \rangle S \rightsquigarrow \langle A : \bar{D} \rangle S} \\
\\
\text{SIMP-MOVE} \\
\frac{I = \text{dom}(D_0) \quad \text{anchor}(\Gamma, A.I, S) = \mathbb{Z}.P.I'}{\Gamma \vdash \langle A : \bar{D}, D_0 \rangle S \rightsquigarrow \langle A : \bar{D} \rangle S[A.I \leftarrow \mathbb{Z}.P.I']} \\
\\
\text{SIMP-SPLIT-TYPE} \\
\frac{\Gamma \vdash \langle A : \bar{D}, \text{type } t = A.t, \text{module } X : \text{sig}_B \bar{D}_0, \text{type } t = A.t, \bar{D}'_0 \text{ end} \rangle S \rightsquigarrow \langle A : \bar{D} \rangle S'}{\Gamma \vdash \langle A : \bar{D}, \text{module } X : \text{sig}_B \bar{D}_0, \text{type } t = B.t, \bar{D}'_0 \text{ end} \rangle S \rightsquigarrow \langle A : \bar{D} \rangle S'} \\
\\
\text{SIMP-SPLIT-MOD} \\
\frac{\Gamma \uplus A : \bar{D} \vdash S_1 \quad \Gamma \vdash \langle A : \bar{D}, \text{module } X'_1 : S_1 \text{ module } X : \text{sig}_B \bar{D}_0, \text{module } X_1 : (= A.X' < S_1), \bar{D}'_0 \text{ end} \rangle S \rightsquigarrow \langle A : \bar{D} \rangle S'}{\Gamma \vdash \langle A : \bar{D}, \text{module } X : (\text{sig}_B \bar{D}_0, \text{module } X_1 : S_1, \bar{D}'_0 \text{ end}) \rangle S \rightsquigarrow \langle A : \bar{D} \rangle S'} \\
\\
\text{SIMP-KEEP} \\
\frac{\text{dom}(D) = I \quad \Gamma \vdash \langle A : \bar{D}_0 \rangle ((\text{sig}_B D, \text{module } \mathbb{Z} : S \text{ end})[A.I \leftarrow B.I]) \rightsquigarrow \langle A : \bar{D}'_0 \rangle \text{sig}_B D', \text{module } \mathbb{Z} : S' \text{ end}}{\Gamma \vdash \langle A : \bar{D}_0, D \rangle S \rightsquigarrow \langle A : \bar{D}'_0, D' \rangle S'[B.I \leftarrow A.I]}
\end{array}$$

Fig. 11. Simplification rules

Given a signature $\langle \gamma \rangle S$, simplification can only *remove* floating fields from γ , but will not introduce new fields (even if this could result in a smaller zipper context). It does not move fields in the signature S (while narrow subtyping allows it), but it can rewrite the content of its fields, rewiring type and module equalities. Simplification uses a first-order criterion for applicative functors, similar to the anchoring restriction of [2], that only rewrites functor applications when aliases (for either a floating functor or a floating module argument) are available.

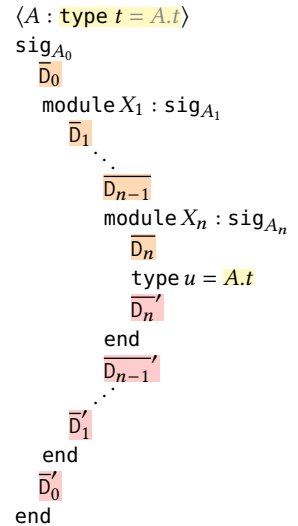
4.2.1 Dropping a floating field. When a floating field is useless, i.e., does not appear in the free variables of the signature, we may just remove it, as done by Rule **SIMP-DROP**. We may also use normalization to make floating module-type fields and floating concrete-type fields useless (as their name is replaced by their definition). Floating core-language value fields can always be dropped.

4.2.2 Moving a floating field. There are cases where the floating field D is not useless ($A.I$ appears in the signature even after normalization), but can still be removed. Indeed, there might be a declaration in S that can take up the same role as D , without loss of type information.

Anchoring points. We call such declaration an *anchoring point* for $A.I$ inside S , which we write $\text{anchor}(\Gamma, A.I, S)$ if it exists. It must validate three conditions: (1) when I is a type field, then D must be of the form $\text{type } t' = A.t$; when I is a module field, then D must be of the form $\text{module } X' : (= A.X < S)$ where S is equivalent to the signature of $A.X$ (not just a supertype). (2) it must be in a strictly positive position, inside neither a functor nor a module type. (3) it must come before any other occurrence of $A.I$ (which are called *usage points*). Formally, we have for a type field $\text{anchor}(\Gamma, A.t, S) = \mathbb{Z}.X_1.(...)X_n.u$ if and only if there exists n signatures S_1, \dots, S_n such that:

$$\begin{aligned}
S &= \text{sig}_{A_0} \bar{D}_0 \text{ module } X_1 : S_1 \bar{D}'_0 \text{ end} \wedge A.t \notin \text{fv}(\bar{D}_0) \\
S_1 &= \text{sig}_{A_1} \bar{D}_1 \text{ module } X_2 : S_2 \bar{D}'_1 \text{ end} \wedge A.t \notin \text{fv}(\bar{D}_1) \\
&\vdots \\
S_n &= \text{sig}_{A_n} \bar{D}_n \text{ type } u = A.t \bar{D}'_n \text{ end} \wedge A.t \notin \text{fv}(\bar{D}_n)
\end{aligned}$$

That is, there is a cascade of depth n of nested signatures where $A.t$ is not mentioned until the field $\text{type } u = A.t$. It is illustrated on the right for a type field.



$$\begin{aligned}
S[A.X] &\triangleq \text{sig type } u = A.X.t \text{ module } X' : (= A.X < \text{sig}_{A'} \text{ type } t = A'.t \text{ end}) \text{ end} & (1) \\
&\langle A : \text{type } t = A.t \text{ module } X : \text{sig type } t = A.t \text{ end} \rangle S[A.X] & (S_2) \\
&\approx \langle A : \text{type } t = A.t, \text{ module } X : \text{sig type } t = A.t \text{ end} \rangle & (S_3) \\
&\quad \text{sig type } u = A.t \text{ module } X' : (= A.X < \text{sig}_{A'} \text{ type } t = A'.t \text{ end}) \text{ end} \\
&\approx \langle A : \text{type } t = A.t \rangle \text{sig type } u = A.t \text{ module } X' : \text{sig type } t = A.t \text{ end} \text{ end} & (S_4) \\
&\approx \text{sig}_B \text{ type } u = B.u \text{ module } X' : \text{sig type } t = B.u \text{ end} \text{ end} & (S_5) \\
\text{S.-SPLIT-TYPE} &\frac{}{\langle A : \text{module } X : \text{sig}_{A'} \text{ type } t = A'.t \text{ end} \rangle S[A.X]} & (S_6) \\
&\approx \text{sig}_B \text{ type } u = B.u \text{ module } X' : \text{sig type } t = B.u \text{ end} \text{ end} & (S_5)
\end{aligned}$$

Fig. 12. Illustration of simplification by **SIMP-SPLIT-TYPE**. Given the definition (1), we cannot drop the module field in (S₆), nor move it, as a usage point of $A.X$ appears before an anchoring point. Yet, it can be simplified to (S₅) by splitting the type field, as (S₂) is a subtype of (S₅). Indeed, we can rearrange (S₂) into (S₃), then simplify it by moving away the module field $A.X$ in (S₄), then the outer type field $A.t$.

Contextual path substitution. If there exists an anchoring point for $A.t$ at $\mathbb{Z}.X_1.(...).X_n.u$, we may remove the floating field by (intuitively) replacing occurrences of $A.I$ by $\mathbb{Z}.X_1.(...).X_n.u$. However, the path to access u is not the same everywhere inside the signature. Therefore, we need a special form of substitution, called *contextual path substitution*, written $S[A.t \leftarrow \mathbb{Z}.X_1.(...).X_n.u]$ that proceeds as follows:

- we replace type $u = A.t$ by type $u = A_n.u$ (deep in the signature)
- in $\bar{D}_0, \dots, \bar{D}_n$, there is nothing to substitute as $A.t$ does not appear free.
- then u is accessible by $A_n.u$, so we may substitute $A.t$ by $A_n.u$ in the declaration \bar{D}_{n-1}' , by $A_{n-1}.X_n.u$ in the declaration \bar{D}_{n-1}' , and, finally, by $A_0.X_1.(...).X_n.u$ in the declarations \bar{D}_0 .

$$\begin{aligned}
&\text{sig}_{A_0} \\
&\bar{D}_0 \\
&\text{module } X_1 : \text{sig}_{A_1} \\
&\bar{D}_1 \\
&\dots \\
&\bar{D}_{n-1} \\
&\text{module } X_n : \text{sig}_{A_n} \\
&\bar{D}_n \\
&\text{type } u = A_n.u \\
&\bar{D}_n' [A.t \leftarrow A_n.u] \\
&\text{end} \\
&\bar{D}_{n-1}' [A.t \leftarrow A_{n-1}.X_n.u] \\
&\dots \\
&\bar{D}_1' [A.t \leftarrow A_1.X_2.(...).X_n.u] \\
&\text{end} \\
&\bar{D}_0' [A.t \leftarrow A_0.X_1.X_2.(...).X_n.u] \\
&\text{end}
\end{aligned}$$

Basically, when visiting S , contextual path substitution replaces $A.I$ by $A_1.X_2.(...).X_n.u$ stripped of its common prefix with the path of the current point, as illustrated on the right for a type field.

Floating module fields. For a floating module field $\langle A : \text{module } X : S \rangle$, the definition of the anchoring point is the same, except that it must be a module declaration with a transparent signature:

$$S_n = \text{sig}_{A_n} \bar{D}_n \text{ module } X' : (= A.X < S' \rangle \bar{D}_n' \text{ end} \wedge A.X \notin \text{fv}(\bar{D}_1) \dots \wedge A.X \notin \text{fv}(\bar{D}_n))$$

Besides, there is an additional equivalence condition to be satisfied:

$$\Gamma \uplus A_0 : \bar{D}_0 \uplus A_1 : \bar{D}_1 \uplus \dots \uplus A_n : \bar{D}_n \vdash S \approx (S' / A.X)$$

4.2.3 Splitting a floating module field. The splitting rules, **SIMP-SPLIT-TYPE** for a type field and **SIMP-SPLIT-MOD** for a module field, are meant to be used in conjunction with the rules for moving and dropping fields. They simply allow to *temporarily* introduce new floating fields if and only if those additional fields help the module field get simplified away and can themselves eventually be simplified. An example is given in Figure 12. In both rules we request the absence of the corresponding field on the right-hand side signature, therefore allowing only splitting when it helps simplification. Importantly, the rules apply only if the module X has a structural signature, not a functor signature. This is where the *first-order* restriction can be seen, as we do not try to *split* individual fields of functors. Simplification of functors can only use Rule **SIMP-SPLIT-MOD**. Implemented naively, splitting would require exponential backtracking and would be impractical.

4.2.4 Skipping a field. The previous simplification rules analyze on the right-most floating field. If the right-most field does not fall into one of the three cases above, we can keep it as a floating component and skip over. Rule **SIMP-KEEP** just amounts to consider the last field D as part of the visible signature, simplify the rest, and put D back into the zipper. Implemented naively, this would require two substitutions and the allocation of a signature for each skipped field.

4.3 Simplification algorithm

In this subsection we sketch an algorithm that simplifies a whole zipper using the simplification rules presented above. It works in three steps. First, *scanning* browses the zipper and the signature to collect information about the usage points of each floating field. Then, constraint resolution computes which floating fields are going to be *dropped*, *moved*, *split*, or *skipped*. Finally, a *simplification* pass revisits the signature to apply the corresponding transformations.

In the rest of this section, we consider the simplification of a zipper $\langle A : D_1, \dots, D_n \rangle S$, say S_0 , that we assume to be assigned to a variable \mathbb{Z} . We denote by I_k the identifier of the floating field D_k .

4.3.1 Scanning. We start with a depth-traversal of the signature, updating a mutable store $(\phi_k)^{k \in 1..n}$ that maps each k to a list of *usage points* of the identifier I_k in S_0 in order of appearance. Each usage point is one of three kinds:

- **zipper** $(A.I_\ell)$ indicates that $A.I_k$ is used in the floating field $A.D_\ell$.
- **anchor** $(A.I_k.\bar{X}.I, \mathbb{Z}.\bar{X}'.I')$ is used when the declaration at $\mathbb{Z}.\bar{X}'.I'$ could serve as an anchoring point for the subfield $A.I_k.\bar{X}.I$ if no occurrence of $A.I_k$ appears before $\mathbb{Z}.\bar{X}'.I'$.
- **usage** is used otherwise. For module fields where $A.I_k$ is used as a prefix, we store the whole path, as **usage** $(A.I_k.\bar{X}.I)$; otherwise, **usage** has no argument.

We then define a function $\text{visit}_P(\cdot)$ that visits the zipper S_0 recursively, in order of appearance, while updating the map ϕ (adding new elements to the tail of the list). P is an optional argument that is only passed when visiting the zipper body S . Initially, each ϕ_k is the empty list. When visiting the zipper context, only **zipper** (\cdot) usage points are used.

When visiting the signature S , the optional path argument P indicates the path to the root \mathbb{Z} if still accessible, or none otherwise. The initial call is $\text{visit}_{\mathbb{Z}} S$. A type declaration is first checked as a possible anchoring point when the path is nonempty. If not an anchoring point or if the path is empty, then it is a usage point. The path is set to none for recursive calls when entering: (1) a functor or a module-type; (2) a submodule with a transparent signature, as paths reaching inside the submodule are normalized away; and (3) a submodule with a module-type signature (as normalization should be used first).

4.3.2 Constraint resolution. For constraint resolution, we use the information collected in the first pass to compute the set of simplifications that can be applied to the signature. For that purpose, we introduce a single assignment map $(\Omega_k)^{k \in 1..n}$ initially undefined everywhere, where each Ω_k will be assigned an *action* **drop**, **move**, **split**, or **keep** exactly once. Constraint resolution updates Ω and builds a contextual path substitutions Θ , initially the identity, to be applied to S .

We iterate over k (hence over floating fields I_k) in reverse order, from n to 1. At each step, we consider the list ϕ_k of usage points collected in the first phase. We first remove from ϕ_k all usage points of the form **zipper** $(A.I_\ell)$ for which Ω_ℓ is **drop**, **move**, or **split**, since then I_ℓ will be removed during simplification. We then scan the list ϕ_k of remaining usage points in order. Anytime we set Ω_k , we proceed with the predecessor of k .

- if ϕ_k is empty, we set Ω_k to **drop**.
- if the head of ϕ_k is **anchor** (P, P') we set Ω_k to **move** and Θ to $[P \leftarrow P'] \circ \Theta$.

- Otherwise, the head of ϕ_k is a usage point. If it has no argument, then I_k is a type field and we set ϕ_k to **keep** (as it is used before an anchoring point); otherwise, I_k is a module field, which we try to split, as follows.

We consider a local assignment map ω for suffixes of $A.I_k$ which are paths of the form $A.I_k.\bar{X}.I'$.

We continue exploring the list ϕ_k with the following sub-cases:

- **usage** (P) : P must be a suffix of $A.I_k$. If a prefix of P (including P) is already set to **move** in ω , we remove the current usage point from ϕ_k and continue with the rest of ϕ_k ; otherwise, we set Ω_k to **keep**.
- **anchor** (P, P') : P must be a suffix of $A.I_k$. If no prefix of P is already set to **move** in ω , we set both $\omega(A.I_k.\bar{X}.I)$ to **move** and Θ to $[P \leftarrow P'] \circ \Theta$. In any case, we remove the current usage point from ϕ_k and continue with the rest of ϕ_k .
- If ϕ_k is empty, we set Ω_k to **split**.

4.3.3 Simplification. The simplification of S_0 is the zipper $\langle A : (\bar{D}_k)^{k \in K} \rangle (S\Theta)$ that just retains the floating fields $(\bar{D}_k)^{k \in K}$ where K is the subdomain of Ω that is mapped to **keep**, after applying the path contextual substitution Θ to S .

4.4 Preservation of typability

In this section we show that narrow subtyping (which includes simplification) never prevents further typings, i.e., if a typing derivation of M inferred a type S for a subexpression N that could be simplified into S' , then we could have continued the typing derivation by taking S' instead of S for the type of N without losing information (up to further simplifications). We prove an almost equivalent, simpler, *principality* property: *If $\Gamma \vdash M : S$ is derivable (with or without simplifications) and Γ' is a typing environment where simplifications have been performed, which we write $\Gamma \approx \Gamma'$, then we have $\Gamma' \vdash M : S'$ for some S' that is equivalent to a simplification of S .*

We do so by establishing a correspondence between Γ and Γ' , keeping track of floating fields that have been simplified, instead marked as *eventually removable* fields. This enables a dual view, where removable fields can either be seen or ignored. Technically, we introduce ZIPML^F , a conservative extension of ZIPML with *full zippers* $\langle A : \bar{D} \mid \bar{D}' \rangle S$. While a simple zipper $\langle A : D \rangle S$ stands, intuitively, for a signature $\text{sig}_A \bar{D}, \mathbb{Z} : S, \bar{D}'$ end that has been projected on field \mathbb{Z} , keeping the left floating fields \bar{D} but forgetting about the right floating fields \bar{D}' that follow the signature S of field \mathbb{Z} , a full zipper also retains the right floating fields, which may then refer to all other floating fields via the label A and to the signature S via the path $A.\mathbb{Z}$. Hence, a full zipper $\langle A : \bar{D} \mid \bar{D}' \rangle S$ is well-formed whenever the signature $\text{sig}_A \bar{D}, \mathbb{Z} : S, \bar{D}'$ end is. A simple zipper $\langle A : D \rangle S$ is then just a particular case $\langle A : D \mid \emptyset \rangle S$ of a full zipper with an empty sequence of right floating fields.

Strengthening is extended to full zippers in the obvious way, while most definitions treat zippers abstractly and need not be redefined,²³ except for subtyping: we replace Rule **SUB-S-ZIPPER** (which will be derivable), by the following Rule **SUB-S-ZIPPERF** that allows subtyping between full zippers:

$$\frac{\text{SUB-S-ZIPPERF} \quad \Gamma_1 = \Gamma \uplus A : (\bar{D}_1, \mathbb{Z} : S, \bar{D}'_1) \quad \bar{D}'_1 \subseteq \bar{D}_1, \bar{D}'_1 \quad \Gamma_1 \vdash \bar{D}'_1 // A \preccurlyeq \bar{D}_2, \bar{D}'_2 \quad \Gamma_1 \vdash S_1 \preccurlyeq S_2}{\Gamma \vdash \langle A : \bar{D}_1 \mid \bar{D}'_1 \rangle S_1 \preccurlyeq \langle A : \bar{D}_2 \mid \bar{D}'_2 \rangle S_2}$$

Interestingly, dropping right floating fields is possible as a particular case of subtyping. In particular, $\Gamma \vdash \langle A : \bar{D} \mid \bar{D}' \rangle S \preccurlyeq \langle A : \bar{D} \rangle S$ always holds. We define the *projection* of a signature S , and write $[S]$ for the signature obtained by the removal of all right floating fields of S (including deep zippers). While toplevel projection preserves well-formedness, deeper projection may not—as is usually the case with subtyping. We say that a signature is *projectable* if its projection is well-formed. In

²³Definition should respect the unzipped signature view of full zippers.

particular, it means that the right floating fields of its zippers are not accessed by the rest of the signature (but can be accessed by other right floating fields).

The key point is that narrow subtyping in ZIPML can be seen in ZIPML^F as an equivalence \approx followed by a projection. Therefore, to show that narrow subtyping preserves typing, we *delay* the projection in ZIPML^F, keeping the projected fields in the right-hand side. This enables reasoning by equivalences in ZIPML^F, and, as long as we preserved projectability, project the judgments back into ZIPML at the end.

Technically, we link a typing environment Γ of ZIPML (without full zippers) with an *equivalent* typing environment Γ' of ZIPML^F, where some signatures have been rewritten along the equivalence of ZIPML^F. We write $\Gamma \vdash S \approx S'$ to mean $\Gamma \vdash S \approx S'$ in ZIPML^F when S is in ZIPML and S' is projectable. If Γ is in ZIPML, this coincides with narrow subtyping. We extend the notation \approx to declarations in the obvious way. We then extend it typing environments with the rule given on the right. By construction, $\Gamma \approx \Gamma'$ implies that Γ is in ZIPML. The premise $[\Gamma'] \vdash [\bar{D}']$ ensures that $[\Gamma' \uplus A : \bar{D}']$ is well-formed.

$$\frac{\Gamma \approx \Gamma' \quad \Gamma' \vdash \bar{D} \approx \bar{D}' \quad [\Gamma'] \vdash [\bar{D}']}{\Gamma \uplus A : \bar{D} \approx \Gamma' \uplus A : \bar{D}'}$$

To establish a tight correspondence between ZIPML and ZIPML^F, we must not only project signatures but also judgments—and derivations. We write $\Gamma' \vdash S$ to mean that both $\Gamma' \vdash S$ holds in ZIPML^F and $[\Gamma'] \vdash [S]$ holds in ZIPML²⁴. We similarly define $\Gamma' \vdash M : S$ and $\Gamma' \vdash S \leq S'$. The principality property is the projection in ZIPML of the following lemma:

LEMMA 4.1. *If $\Gamma \vdash M : S$ and $\Gamma \approx \Gamma'$, then there exists S' such that $\Gamma' \vdash M : S'$ and $\Gamma' \vdash S \approx S'$.*

Its proof is a combination of reasoning along equivalences, relying on general meta-theoretical properties of ZIPML^F, followed by more specific, but easy results about projectability, ensuring that not only judgments can be projected, but their derivation can be done ZIPML. In particular:

- If $\Gamma \vdash S$ and $\Gamma \vdash S'$ and $\Gamma \vdash S \leq S'$, then $\Gamma \vdash S \leq S'$. (When \leq ranges over \lesssim and \leq .)
- If $\Gamma \vdash S$ and $\Gamma \vdash M : S$, then $\Gamma \vdash M : S$.

Simplifications in ZIPML^F. So far, we have not changed the simplification rule **TYP-M-SIMPL**. Hence, we cannot do more simplifications in ZIPML^F that in ZIPML, although it would be safe to do so. In particular, any simplification along the code-free equivalence is safe. Doing this, simplifications in ZIPML^F could keep removable floating fields as right floating fields, preparing and enabling their eventual removal later on just by projection.

5 Properties

The main property, type soundness, is proved by elaboration of ZIPML into M^ω [2], which is sound, while preserving the underlying untyped semantics. This is done in §5.1. In §5.3, we present some arguments supporting a completeness conjecture, i.e., that any M^ω program can be elaborated into an equivalent ZIPML program. However, since ZIPML has a module-level notion of sharing, the comparison should be made with version M_{id}^ω of M^ω instrumented with module-level sharing, obtained by the composition of the source-to-source transformation that introduces identity tags at the level of modules prior to typing in M^ω , as described in [2, §3.6], or equivalently, using derived rules operating directly on source terms without their translation. We choose the later below, but just keep writing M^ω for M_{id}^ω . By contrast with [2], we use pairs to (τ, C) to represent tagged signatures $\text{sig type } id = \tau \text{ module } Val : C$ end and keep them concise and distinct from regular signatures. We still keep letter C to range over either regular or tagged signatures.

²⁴That is, the derivation can be conducted without using full zippers.

5.1 Soundness of ZIPML by Elaboration in M^ω

We expect the reader to be familiar with elaboration of ML modules into F^ω , ideally M^ω [2] (or *F-ing* [28]). This elaboration serves as a proof of soundness of ZIPML. Elaborating in M^ω rather than directly in F^ω allows to benefit from the sound extrusion and skolemization of M^ω . For lack of space, we must refer the reader to [2] for a precise definition of the typing judgments of M^ω .

As both ZIPML and M^ω have the same source language and the same untyped semantics, the goal is to show the following theorem, so that ZIPML inherits type soundness from M^ω :

THEOREM 5.1 (SOUNDNESS). *Every module expression that typechecks in ZIPML (without the simplification rule) also typechecks in M^ω . That is, $\vdash^\diamond M : S$ implies $\Vdash M : \exists^\diamond \bar{\alpha}. C$ for some $\bar{\alpha}$ and C .*

Since simplifications can always be postponed, as shown in §4.4, we may reason using the subtyping relation of §3 without zippers on the right. We write \Vdash for judgments in M^ω , as opposed to \vdash for judgments in ZIPML, and we use a light red background. In order to prove this result by induction on the typing derivations we extend it to a non-empty typing environment and link the two output signatures. To that aim, we introduce an elaboration of source typing environments Γ into M^ω ones, written $\Vdash \Gamma \rightsquigarrow \Gamma_\omega$. The induction hypothesis for the soundness statement becomes:

$$\Gamma \vdash^\diamond M : S \wedge \Vdash \Gamma \rightsquigarrow \Gamma_\omega \quad \Longrightarrow \quad \exists \bar{\alpha}, C. \Gamma_\omega \Vdash S : \lambda \bar{\alpha}. C \wedge \Gamma_\omega \Vdash M : \exists^\diamond \bar{\alpha}. C$$

Yet, we need to extend M^ω elaboration of signatures to support zippers for this statement to make sense. The rest of this section is composed as follows: (1) we first explain the treatment of zippers; (2) we discuss the elaboration of abstract types in the environment and extend the induction hypothesis; (3) we state the elaboration of strengthening, normalization, subtyping, and resolution; (4) we state the elaboration of path typing, signature typing, and module typing.

5.1.1 Treatment of zippers. The first difficulty in the soundness statement is that the language of inferred signatures of ZIPML is larger than the source signatures of M^ω , with the introduction of zippers signatures $\langle \gamma \rangle S$ and zipper-context accesses in paths $P.A$. Zippers are removed at runtime, and could be removed in M^ω . However, to establish a one-to-one correspondence on inferred signatures, we need to keep zippers in both inferred signatures and the environment.

Therefore, we define M_{zip}^ω , an extension of M^ω with zippers. We add zippers to M^ω signatures with the following elaboration rule, along with appropriate path typing extension. However, and this is a key point, those zippers are only used during the proof by induction for typing the inferred signatures of ZIPML. They should be seen as *decorations* on types and contexts to build a typing derivation in M^ω , after which zippers can be erased. In a typing derivation of M_{zip}^ω that does not use **M-TYP-S-ZIP** nor **M-TYP-P-ZIP**, we can erase all zippers and obtain a normal derivation of M^ω .

$$\begin{array}{c} \text{M-TYP-S-ZIP} \\ \frac{\Gamma \Vdash_A \bar{D} : \lambda \bar{\alpha}. \bar{D} \quad \Gamma, A : \bar{D} \Vdash S : \lambda \bar{\beta}. C}{\Gamma \Vdash \langle A : \bar{D} \rangle S : \lambda \bar{\alpha}, \bar{\beta}. \langle A : \bar{D} \rangle C} \\ \\ \text{M-TYP-P-ZIP} \\ \frac{\Gamma \Vdash P : \langle A : \bar{D} \rangle C}{\Gamma \Vdash P.A : \text{sig } \bar{D} \text{ end}} \end{array}$$

5.1.2 Elaboration of environments and strengthening. When elaborating a typing environment Γ of ZIPML into a typing environment Γ_ω of M^ω , a technical point arises from the fact that declarations and signatures in Γ are always strengthened, and are therefore self-referring. By contrasts, abstract types variables are inserted in Γ_ω before declarations and signatures.

We define a (non-algorithmic) rule **M-TYP-E-DECL** that captures the *self-referring* representation of abstract types in ZIPML. It is best understood by its key property, shown on the right:

$$\begin{array}{c} \text{M-TYP-E-DECL} \\ \frac{\Vdash \Gamma \rightsquigarrow \Gamma_\omega \quad \Gamma, \bar{\alpha}, A : \mathcal{D} \Vdash_A D : \mathcal{D}}{\Vdash (\Gamma, A : D) \rightsquigarrow (\Gamma_\omega, \bar{\alpha}, A : \mathcal{D})} \quad \frac{\Vdash \Gamma \rightsquigarrow \Gamma_\omega \quad \Gamma_\omega \Vdash_A D : \lambda \bar{\alpha}. \mathcal{D}}{\Vdash (\Gamma \uplus A : D) \rightsquigarrow (\Gamma_\omega, \bar{\alpha}, A : \mathcal{D})} \end{array}$$

The rule is not algorithmic as it requires *guessing* the set of abstract type variables $\bar{\alpha}$ and the result of the elaboration \mathcal{D} . However, as shown by the derived rule on the right, $\bar{\alpha}$ and \mathcal{D} can be obtained by elaborating the unstrengthened declaration D_0 , when it is added to the context. This is sufficient to conduct our proof, as declarations are only added to the context with \uplus . We have similar rules and properties for adding functor parameters and zippers to the environment.

We have the following property for strengthening:

LEMMA 5.2 (ELABORATION OF STRENGTHENING). *Given a path P and a signature S , such that $\llbracket \Gamma \rrbracket \Vdash P : C'$ and $\llbracket \Gamma \rrbracket \Vdash S : \lambda \bar{\alpha}. C$ and $\llbracket \Gamma \rrbracket \Vdash C' \leq C[\bar{\alpha} \leftarrow \bar{\tau}]$, we have $\llbracket \Gamma \rrbracket \Vdash S // P : (\lambda \bar{\alpha}. C) \bar{\tau}$.*

5.1.3 *Elaboration of judgments.* As normalization, path resolution, path typing, and subtyping are mutually recursive, we state four combined properties that are proved by mutual induction. Type and module type definitions are kept as such and only inlined on demand in ZIPML, while they are immediately inlined in M^ω . Therefore, ZIPML normalization becomes the identity in M^ω .

LEMMA 5.3 (ELABORATION OF NORMALIZATION). *If S normalizes to S' , i.e., $\Gamma \vdash S \downarrow S'$, and if we have $\Vdash \Gamma \rightsquigarrow \Gamma_\omega$, then both signatures have the same elaboration in M^ω . That is, $\Gamma_\omega \Vdash S : \lambda \bar{\alpha}. C$ implies $\Gamma_\omega \Vdash S' : \lambda \bar{\alpha}. C$.*

LEMMA 5.4 (ELABORATION OF PATH-TYPING). *If $\Gamma \vdash P : S$ and $\Vdash \Gamma \rightsquigarrow \Gamma_\omega$ hold, then the typing of P and of its signature S coincide in M^ω , i.e., we have both $\Gamma_\omega \Vdash P : (\tau, C)$ and $\Gamma_\omega \Vdash S : (\tau, C)$.*

LEMMA 5.5 (ELABORATION OF PATH RESOLUTION). *Resolving a path P allows to fetch either the identity or the content of the signature. Assuming $\Vdash \Gamma \rightsquigarrow \Gamma_\omega$, we have:*

- *If $\Gamma \vdash P \triangleright S$ then we have $\Gamma_\omega \Vdash P : (_ , C)$ and $\Gamma_\omega \Vdash S : \lambda \alpha. (_ , C)$;*
- *If $\Gamma \vdash P \triangleright P'$ then we have $\Gamma_\omega \Vdash P : (\tau, C)$ and $\Gamma_\omega \Vdash P' : (\tau, C')$ and $\Gamma_\omega \Vdash C' \leq C$.*

LEMMA 5.6 (ELABORATION OF SUBTYPING). *The elaboration preserves the subtyping relationship: If $\Gamma_\omega \Vdash S : \lambda \bar{\alpha}. C$ and $\Gamma_\omega \Vdash S' : \lambda \bar{\alpha}'. C'$ then $\Gamma \vdash S \leq S'$ implies $\Gamma_\omega \Vdash \lambda \bar{\alpha}. C \leq \lambda \bar{\alpha}'. C'$.*

LEMMA 5.7 (ELABORATION OF SIGNATURE TYPING). *Well-typed signatures of ZIPML can be elaborated in M^ω . If $\Gamma \vdash S$ and $\Vdash \Gamma \rightsquigarrow \Gamma_\omega$ then $\exists \bar{\alpha}, C. \Gamma_\omega \Vdash S : \lambda \bar{\alpha}. C$.*

Soundness as stated by [Theorem 5.1](#) relies on [Lemmas 5.3](#) to [5.6](#), which are proved by mutual induction over the typing derivations of their premises.

5.2 Zipper renaming, introduction, and extrusion

These operations are not directly available in ZIPML. Yet, we could add them as additional typing rules.

5.2.1 *Renaming of zipper labels.* When a zipper is introduced during projection, it turns a self-reference into a label whose name is chosen arbitrarily and cannot later be changed. Hence, it must be chosen fresh so as to avoid shadowing of other labels of the same name, which well-formedness forbids. Still, labels can be renamed *consistently*. For example, if $\Gamma \vdash M : \langle A : D \rangle S$ and A does not appear free in Γ , then we also have $\Gamma \vdash M : \langle A' : D[A \leftarrow A'] \rangle S[A \leftarrow A']$ when A' is chosen fresh for Γ , D , and S . This is a lemma and not a *derivable* typing rule. However, we can then safely add renaming as an *admissible* typing rule.

A logically cleaner solution, worth considering in the future, would be to introduce a binder construct $\nu A. S$ in signatures that limits the scope of A to S . However, this would be more involved and invasive as (1) typing rules should then be adjusted to extrude binders and gather them at to the toplevel of typing judgments and in the codomains of generative functors, and (2) equivalence (an other judgments) should also be redefined up to the consistent renaming of binders.

5.2.2 Introduction of zippers. In ZIPML, zippers are only introduced during projections and indirectly transported by other operations. They can be simplified at toplevel. Yet, they cannot be moved, in particular, they cannot be extruded. In ZIPML^F, we may allow the introduction of a right-zipper as a toplevel typing rule. This is always safe, as the extra fields \bar{D} , are not currently accessible from S and, intuitively, will not be accessible by the rest of the program.

Interesting, the rule may then be followed by an equivalence that moves the right floating fields on the left and do some “rewiring” in S , accordingly. For example, we could turn a module of signature $\text{sig}_A \bar{D}$ end into one of type $\langle A : \bar{D} \rangle \text{sig } \bar{D} // A$ end by first introducing $\langle A : \emptyset | \bar{D} \rangle S$, then using equivalence and simplification. This could also be used for zipper extrusion.

For example, if $\Gamma \vdash M : S_0$ where S_0 is sig_A module $X : \langle A_0 : D \rangle S$ end and A does not appear in D , then we could extrude $A_0 : D$ by giving M the successive types:

- $S_0 \triangleq \text{sig}_A$ module $X : \langle A_0 : D \rangle S$ end *initially*
- $\langle A_1 : \emptyset | D // A_1.Z.X.A_0 \rangle \text{sig}_A$ module $X : \langle A_0 : D \rangle S$ end *by introduction of a zipper*
- $\langle A_1 : D \rangle \text{sig}_A$ module $X : \langle A_0 : D // A_1 \rangle S$ end *by equivalence*
- $\langle A_1 : D \rangle \text{sig}_A$ module $X : \langle A_0 : D // A_1 \rangle S[A.X.A_0 \leftarrow A_1]$ end *by equivalence*
- $\langle A_1 : D \rangle \text{sig}$ module $X : \langle A_0 : \emptyset | D // A_1 \rangle S[A.X.A_0 \leftarrow A_1]$ end *by equivalence*
- $\langle A_1 : D \rangle \text{sig}$ module $X : S[A.X.A_0 \leftarrow A_1]$ end *by projection*
- $S_1 \triangleq \langle A_0 : D \rangle \text{sig}$ module $X : S[A.X.A_0 \leftarrow A_0]$ end *by renaming of A_1 .*

5.3 Completeness of ZIPML with respect to M^ω

The reader may wonder whether the type system of ZIPML is powerful enough to encode the F^ω -based systems like M^ω [2]. We *conjecture* that it is indeed the case. In this section, we present some key properties supporting this claim. We reason in ZIPML^F to benefit from full zippers as an intermediate step, with the extension of §5.2 to enable zipper renaming and extrusion. Then, we remark that: (1) floating fields can be used to encode existentially quantified variables; (2) zipper introduction and code-free equivalence on full zippers can simulate the M^ω extrusion and skolemization of variables; and (3) universally and lambda bound variables need not be encoded.

Quantified variables are α -convertible in M^ω , but labels of zipper fields are not. Label renaming is not part of equivalence, but still possible on the right of typing judgments, which suffices.

5.3.1 Encoding existentially quantified types with floating fields. Our claim is that *floating fields can encode all existentially quantified variables of M^ω* . We first consider first-order type variables, then module identities, and, finally, higher-order types.

First-order type variables. Let us consider a M^ω -signature with a single quantified type variable of the base kind \star , which is of the form $\exists \alpha. C$. Inside C , the variable α is accessible everywhere. This is quite similar to a zipper signature $\langle A : \text{type } t = A.t \rangle S$, where $A.t$ is a type accessible everywhere inside S . Therefore, we may translate the M^ω -signature into a zipper signature by introducing a floating field with a fresh name for every quantified variable. We may define a reverse elaboration judgment $\Gamma_\omega \Vdash \exists \alpha. C \hookrightarrow \langle \gamma \rangle S$, where we extend the environment to attach the name of a floating field to every (existentially) quantified abstract type. We would have rules of the form:

$$\frac{\text{REV-S-STAR} \quad \Gamma_\omega, \alpha \hookrightarrow A.t \Vdash C \hookrightarrow S \quad A \text{ fresh}}{\Gamma_\omega \Vdash \exists \alpha. C \hookrightarrow \langle A : \text{type } t = A.t \rangle S} \quad \text{REV-T-ABSTYPE} \quad \frac{\alpha \hookrightarrow A.t \in \Gamma_\omega}{\Gamma_\omega \Vdash \alpha \hookrightarrow A.t}$$

Module identities. Module identities of M^ω can be encoded as floating module fields. However, there is a difficulty: what is the attached signature of an identity floating field? The simplest answer is to extend ZIPML with a *bottom signature* \perp that is a subtype of all signatures. Doing so, we would

get rules of the following form:

$$\frac{\text{REV-S-MOD} \quad \Gamma, \alpha_{id} \hookrightarrow A.X \Vdash C \hookrightarrow S \quad A \text{ fresh}}{\Gamma \Vdash \exists \alpha_{id}. C \hookrightarrow \langle A : \text{module } X : \perp \rangle S} \quad \frac{\text{REV-S-TRANSPARENT} \quad \alpha_{id} \hookrightarrow A.X \in \Gamma_\omega \quad \Gamma \Vdash C \hookrightarrow S}{\Gamma \Vdash (\alpha_{id}. C) \hookrightarrow (= A.X < S)}$$

However, as hinted in by [2, Theorem 3.1], module identities of M^ω are always attached to signatures that have a common ancestor in the subtyping order. Therefore, rather than extending ZIPML with a bottom signature \perp , we could instrument the typing rules of M^ω to obtain the common ancestor signature associated with each module identity and use it in the corresponding floating field.

Higher-order types. Higher-order types of M^ω can be encoded as floating functors producing a single type field. Again, there is a subtlety: what is the signature of the domain of such a functor? The simplest answer is to extend ZIPML with a *top signature* \top that is a supertype of all signatures. Using it, we would get rules

$$\frac{\text{REV-S-HIGHERORDER} \quad \Gamma, \varphi \hookrightarrow A.F(\cdot).t \Vdash C \hookrightarrow S}{\Gamma \Vdash \exists^\forall \varphi. C \hookrightarrow \langle A : \text{module } F : (Y : \top) \rightarrow \text{sig}_B \text{ type } t = B.t \text{ end} \rangle S} \quad \frac{\text{REV-T-TRANSPARENT} \quad \varphi \hookrightarrow A.F(\cdot).t \in \Gamma_\omega \quad \alpha_{id} \hookrightarrow B.X}{\Gamma \Vdash \varphi(\alpha_{id}) \hookrightarrow A.F(B.X).t}$$

Higher-order module identities would work similarly. Following the same reasoning as for module identities, we conjecture that, rather than extending ZIPML with a top signature \top , we could instrument the typing rules of M^ω to obtain the domain of the functor that originally introduced φ , which is a supertype of all use-cases.

Universally and lambda quantified types. Our argument applies to existentially quantified types. But the universal quantification and lambda quantification of M^ω are always used for signatures that come from elaboration of the source (namely, functor parameter and module type definitions), and can therefore also be represented in ZIPML, without using floating field to encode type variables.

5.3.2 Extrusion. Floating fields can be extruded, similarly to existential types. For instance, if floating fields have been introduced in front of submodules, we could extrude them:

```
if M : sigA module X1 : ⟨A1 : type t = A1.t⟩ sig type t = A1.t end
      module X2 : ⟨A2 : type t = A2.t⟩ sig type t = A2.t end end
then M : ⟨A1 : type t = A1.t ; A2 : type t = A2.t⟩
        sigA module X1 : sig type t = A1.t end module X2 : sig type t = A2.t end end
```

This also applies to *skolemization*, as we could also have:

```
if M : (Y : Sa) → ⟨B : type t = B.t⟩ sig type t = B.t end
then M : ⟨B : module F : (Y0 : Sa) → sigD type t = D.t end⟩ (Y : Sa) → sig type t = B.F(Y).t end
```

where we would first introduce the outer zipper as a right floating field, apply an equivalence, and remove the inner zipper by projection. (Here, for the domain of the floating functor, we could also use the top signature \top instead of the signature S_a of the functor we extruded from, as we did.)

In summary, the extrusion and skolemization mechanisms of M^ω could be mimicked in ZIPML using extrusion of floating fields. This should suffice to maintain a compositional correspondence between typings in M^ω and typings in ZIPML.

6 Missing features and Conclusion

We have presented ZIPML, a source system for ML modules that uses a new feature, signature zippers, to delay and resolve instances of signature avoidance. ZIPML also models transparent ascription, delayed strengthening, applicative and generative functors, and parsimonious inlining of signatures. Several features of OCAML are still missing in ZIPML.

The **open** and **include** constructs allow users to access or inline a given module. While not problematic when used on paths, OCAML also allows opening structures [16], which may trigger signature avoidance, as shown in §2 on a restricted case. We expect floating fields to easily model the opening of structures although some adjustments will be needed. In particular, type checking of signatures will have to become an elaboration judgment as mentioned in §3.5.

OCAML allows *abstract signatures* which amounts to quantifying over signatures in functors. This feature, while rarely used in practice, unfortunately makes the system undecidable [25, 33]. In our context, as ZIPML expands module type names only by need. We conjecture that abstract signatures could be added to the system, as they should not impact zippers. However, the undecidability of subtyping should be addressed, maybe by restricting their instantiation.

Typechecking of recursive modules raises the question of *double vision* [6, 23]. By contrast, OCAML requires full type annotations, along with an initialization semantics which can fail at runtime. All these solutions are compatible with ZIPML.

We leave these explorations to future work. Implementing zipper signatures, transparent ascription and early-lazy strengthening in OCAML remains to be done to appreciate the gain in expressiveness and the impact on typechecking speed. A mechanization of ZIPML metatheory, for which we only have paper-sketched proofs, would also be worth doing and would fit well with other efforts towards a mechanized specification of OCAML and formal proofs of properties of OCAML programs.

References

- [1] Sandip K. Biswas. 1995. Higher-Order Functors with Transparent Signatures. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) (POPL '95). Association for Computing Machinery, New York, NY, USA, 154–163. <https://doi.org/10.1145/199448.199478>
- [2] Clément Blaudeau, Didier Rémy, and Gabriel Radanne. 2024. Fulfilling OCaml Modules with Transparency. *Proc. ACM Program. Lang.* 8, OOPSLA1, Article 101 (apr 2024), 29 pages. <https://doi.org/10.1145/3649818>
- [3] Luca Cardelli and Xavier Leroy. 1990. Abstract types and the dot notation. In *Proceedings IFIP TC2 working conference on programming concepts and methods*. North-Holland, 479–504.
- [4] Karl Crary. 2017. Modules, Abstraction, and Parametric Polymorphism. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL 2017)*. ACM, New York, NY, USA, 100–113. <https://doi.org/10.1145/3009837.3009892> event-place: Paris, France.
- [5] Karl Crary. 2020. A focused solution to the avoidance problem. *Journal of Functional Programming* 30 (2020), e24. <https://doi.org/10.1017/S0956796820000222>
- [6] Derek Dreyer. 2007. Recursive type generativity. *Journal of Functional Programming* 17, 4-5 (2007), 433–471. <https://doi.org/10.1017/S0956796807006429>
- [7] Derek Dreyer, Karl Crary, and Robert Harper. 2003. A type system for higher-order modules. In *Conference Record of POPL 2003: The 30th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, New Orleans, Louisiana, USA, January 15-17, 2003*, Alex Aiken and Greg Morrisett (Eds.). ACM, 236–249. <https://doi.org/10.1145/604131.604151>
- [8] Jacques Garrigue and Leo White. 2014. Type-level module aliases: independent and equal (*ML Family/OCaml Users and Developers workshops*). <https://www.math.nagoya-u.ac.jp/~garrigue/papers/modalias.pdf>
- [9] Robert Harper and Mark Lillibridge. 1994. A Type-Theoretic Approach to Higher-Order Modules with Sharing. In *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Portland, Oregon, USA) (POPL '94). Association for Computing Machinery, New York, NY, USA, 123–137. <https://doi.org/10.1145/174675.176927>
- [10] Robert Harper, John C. Mitchell, and Eugenio Moggi. 1989. Higher-Order Modules and the Phase Distinction. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) (POPL '90). Association for Computing Machinery, New York, NY, USA, 341–354. <https://doi.org/10.1145/96709.96744>
- [11] Robert Harper and Christopher A. Stone. 2000. A type-theoretic interpretation of standard ML. In *Proof, Language, and Interaction*. <https://api.semanticscholar.org/CorpusID:9208816>
- [12] Gérard Huet. 1997. The Zipper. *Journal of Functional Programming* 7, 5 (1997), 549–554. <https://doi.org/10.1017/S0956796897002864>

- [13] Xavier Leroy. 1994. Manifest Types, Modules, and Separate Compilation. In *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Portland, Oregon, USA) (POPL '94). Association for Computing Machinery, New York, NY, USA, 109–122. <https://doi.org/10.1145/174675.176926>
- [14] Xavier Leroy. 1995. Applicative functors and fully transparent higher-order modules. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages - POPL '95*. ACM Press, San Francisco, California, United States, 142–153. <https://doi.org/10.1145/199448.199476>
- [15] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, Kc Sivaramakrishnan, and Jérôme Vouillon. 2023. *The OCaml system release 5.1: Documentation and user's manual*. Intern report. Inria. <https://inria.hal.science/hal-00930213>
- [16] Runhang Li and Jeremy Yallop. 2017. Extending OCaml's 'open'. In *Proceedings ML Family / OCaml Users and Developers workshops, ML/OCaml 2017, Oxford, UK, 7th September 2017 (EPTCS, Vol. 294)*, Sam Lindley and Gabriel Scherer (Eds.), 1–14. <https://doi.org/10.4204/EPTCS.294.1>
- [17] David MacQueen, Robert Harper, and John Reppy. 2020. The history of Standard ML. *Proc. ACM Program. Lang.* 4, HOPL, Article 86 (jun 2020), 100 pages. <https://doi.org/10.1145/3386336>
- [18] David B. MacQueen. 1986. *Using Dependent Types to Express Modular Structure*. Association for Computing Machinery, New York, NY, USA, 277–286. <https://doi.org/10.1145/512644.512670>
- [19] Anil Madhavapeddy, Richard Mortier, Charalampos Rotsos, David J. Scott, Balraj Singh, Thomas Gazagnaire, Steven Smith, Steven Hand, and Jon Crowcroft. 2013. Unikernels: library operating systems for the cloud. In *Architectural Support for Programming Languages and Operating Systems, ASPLOS 2013, Houston, TX, USA, March 16-20, 2013*, Vivek Sarkar and Rastislav Bodik (Eds.). ACM, 461–472. <https://doi.org/10.1145/2451116.2451167>
- [20] Robin Milner, Mads Tofte, and Robert Harper. 1990. *The Definition of Standard ML (revised)*. MIT Press, Cambridge, MA, USA.
- [21] Robin Milner, Mads Tofte, and David Macqueen. 1997. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA. <https://doi.org/10.7551/mitpress/2319.003.0001>
- [22] John C. Mitchell and Gordon D. Plotkin. 1985. Abstract Types Have Existential Types. In *Proceedings of the 12th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (New Orleans, Louisiana, USA) (POPL '85). Association for Computing Machinery, New York, NY, USA, 37–51. <https://doi.org/10.1145/318593.318606>
- [23] Keiko Nakata and Jacques Garrigue. 2006. Recursive Modules for Programming. (2006), 13.
- [24] Norman Ramsey. 2005. ML Module Mania: A Type-Safe, Separately Compiled, Extensible Interpreter. In *Proceedings of the ACM-SIGPLAN Workshop on ML, ML 2005, Tallinn, Estonia, September 29, 2005 (Electronic Notes in Theoretical Computer Science, Vol. 148)*, Nick Benton and Xavier Leroy (Eds.). Elsevier, 181–209. <https://doi.org/10.1016/J.ENTCS.2005.11.045>
- [25] Andreas Rossberg. 1999. Undecidability of OCaml type checking. <https://sympa.inria.fr/sympa/arc/caml-list/1999-07/msg00027.html>.
- [26] Andreas Rossberg. 2018. 1ML - Core and modules united. *J. Funct. Program.* 28 (2018), e22. <https://doi.org/10.1017/S0956796818000205>
- [27] Andreas Rossberg and Derek Dreyer. 2013. Mixin'Up the ML Module System. *ACM Trans. Program. Lang. Syst.* 35, 1 (April 2013), 2:1–2:84. <https://doi.org/10.1145/2450136.2450137>
- [28] Andreas Rossberg, Claudio Russo, and Derek Dreyer. 2014. F-ing modules. *Journal of Functional Programming* 24, 5 (Sept. 2014), 529–607. <https://doi.org/10.1017/S0956796814000264>
- [29] Claudio V. Russo. 2000. First-Class Structures for Standard ML. In *Programming Languages and Systems*, Gerhard Goos, Juris Hartmanis, Jan van Leeuwen, and Gert Smolka (Eds.). Vol. 1782. Springer Berlin Heidelberg, Berlin, Heidelberg, 336–350. https://doi.org/10.1007/3-540-46425-5_22 Series Title: Lecture Notes in Computer Science.
- [30] Claudio V. Russo. 2001. Recursive structures for standard ML. *SIGPLAN Not.* 36, 10 (oct 2001), 50–61. <https://doi.org/10.1145/507669.507644>
- [31] Claudio V. Russo. 2004. Types for Modules. *Electronic Notes in Theoretical Computer Science* 60 (2004), 3–421. [https://doi.org/10.1016/S1571-0661\(05\)82621-0](https://doi.org/10.1016/S1571-0661(05)82621-0)
- [32] Zhong Shao. 1999. Transparent Modules with Fully Syntactic Signatures. In *Proceedings of the fourth ACM SIGPLAN International Conference on Functional Programming (ICFP '99), Paris, France, September 27-29, 1999*. ACM, 220–232. <https://doi.org/10.1145/317636.317801>
- [33] Leo White. 2015. Girard Paradox implemented in OCaml using abstract signatures.

Received 2024-07-11; accepted 2024-11-07