

ELIOM: A Language for Modular Tierless Web Programming

GABRIEL RADANNE, University of Freiburg, Germany

JÉRÔME VOUILLON, Univ Paris Diderot, Sorbonne Paris Cité, France, BeSport, France, and CNRS, France

VINCENT BALAT, Univ Paris Diderot, Sorbonne Paris Cité, France and BeSport, France

Tierless Web programming languages allow programmers to combine client-side and server-side programming in a single program. Programmers can then define components with both client and server parts and get flexible, efficient and typesafe client-server communications. However, the expressive client-server features found in most tierless languages are not necessarily compatible with functionalities found in many mainstream languages. In particular, we would like to benefit from type safety, an efficient execution, static compilation, modularity and separate compilation.

In this paper, we propose ELIOM, an industrial-strength tierless functional Web programming language which extends OCAML with support for rich client/server interactions. It allows to build whole applications as a single distributed program, in which it is possible to define modular tierless libraries with both server and client behaviors and combine them effortlessly. ELIOM is the only language that combines type-safe and efficient client/server communications with a static compilation model that supports separate compilation and modularity. It also supports excellent integration with OCAML, allowing to transparently leverage its ecosystem.

To achieve all these features, ELIOM borrows ideas not only from distributed programming languages, but also from meta-programming and modern module systems. We present the design of ELIOM, how it can be used in practice and its formalization; including its type system, semantics and compilation scheme. We show that this compilation scheme preserves typing and semantics, and that it supports separate compilation.

CCS Concepts: • **Software and its engineering** → **Functional languages; Modules / packages; Formal language definitions; Compilers; Distributed programming languages;**

Additional Key Words and Phrases: Web, client/server, OCAML, ML, ELIOM, functional, module

ACM Reference Format:

Gabriel Radanne, Jérôme Vouillon, and Vincent Balat. 2019. ELIOM: A Language for Modular Tierless Web Programming. 1, 1 (February 2019), 87 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Writing websites nowadays requires the use of many different technologies: HTML and CSS for the content and visuals of the website, JAVASCRIPT for client interactions, any of the numerous server languages (PHP, Ruby, C#, ...), a database language such as SQL, etc. Not only do programmers need to use all these technologies, but they also need them to cooperate in an harmonious manner by ensuring that communications between these various components are correct. Such verification is often done manually by the programmer, which is both error-prone and time-consuming.

This work was partially performed at IRILL, center for Free Software Research and Innovation in Paris, France, <http://www.irill.org>.

Authors' addresses: Gabriel Radanne, University of Freiburg, Germany, radanne@informatik.uni-freiburg.de; Jérôme Vouillon, IRIF UMR 8243 CNRS, Univ Paris Diderot, Sorbonne Paris Cité, Paris, France, BeSport, Paris, France, CNRS, France, jerome.vouillon@irif.fr; Vincent Balat, IRIF UMR 8243, Univ Paris Diderot, Sorbonne Paris Cité, Paris, France, BeSport, Paris, France, vincent.balat@irif.fr.

© 2019 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in , <https://doi.org/10.1145/nnnnnnn.nnnnnnn>.

This juggling of many different technologies does not only make programming more complicated, it also imposes strong constraints on code organization and prevents modularity. Let us consider the case where we want to add a comment widget to a website. This widget requires JAVASCRIPT code for client interactions, for example a convenient editor. It also requires server code that will store and serve the comments and potentially some associated database queries. This code will thus be split in two codebases and up to three languages. Furthermore, in order to be properly integrated in the larger website, the internal communications between the client and server parts of the widget will be exposed to the rest of the program. All these programming constraints, which stem directly from the way Web programming is done currently, make it impossible to preserve abstraction and encapsulation for widgets who have both client and server aspects. This hinders safety and reusability of widgets by preventing programmers to create self-contained independent libraries that have both client and server aspects.

1.1 OCSIGEN

OCSIGEN provides a comprehensive set of tools and libraries for developing Web applications in OCAML. OCAML is an industrial strength statically typed functional programming language with a rich ecosystem. The OCSIGEN project includes the compiler JS_OF_OCAML [Vouillon and Balat 2014], a Web server, libraries for concurrency [Vouillon 2008] and HTML manipulation [TyXML 2017]. It also contains a complete framework to develop complex Web application with both client and server components. It includes modules for, among other things, RPCs, Web GUIs [Ocsigen Toolkit 2017], functional reactive Web programming and an advanced *service identification mechanism* [Balat 2014]. OCSIGEN is already used in production for a wide variety of websites [Besport 2017; Gencore 2017; Pumgrana 2017].

Based on our experience building such a framework, we realized we needed additional tools to bridge the gap between client and server components. Indeed, while OCAML provides excellent support for abstraction and modularity through its module language, it doesn't allow to talk directly about the relationship between client and server code. In particular, we needed language constructs that allows us to blur the boundaries between client and server. Such language constructs can notably be provided by tierless languages.

1.2 Tierless Web Programming Languages

Tierless languages [Chlipala 2015b; Cooper et al. 2006; Serrano et al. 2006] allow programmers to write in a composable way programs that have both client and server aspects. They provide language constructs to define on the server functions that create fragments of a Web page together with their associated client-side behavior. This is done by allowing to freely intersperse client and server expressions with seamless communication in a unique programming language. By grouping together client and server parts, they allow to encapsulate hybrid libraries and create more modular websites. To be executed, tierless programs are sliced in two: a part which runs on the server and a part which is compiled to JAVASCRIPT and runs on the client. This slicing can be done either dynamically, by generating JAVASCRIPT code at runtime, or statically, by cutting the program in two during compilation. Tierless languages can also leverage static typing to statically ensure that communications between client and server are always correct. Such tierless languages can also leverage both functional programming and static typing. Functional programming provides increased expressiveness and flexibility while static typing gives strong guarantees about client-server separation, in particular ensuring that communications are consistent across tiers.

To complete the design of OCSIGEN, we needed a language that not only provides such expressive tierless features, but also provides support for most features already supported by OCAML: type safety, efficiency, static compilation, modularity and separate compilation. Furthermore, Web

programming often relies on numerous external libraries (for encryption, HTML, logging, ...). We wanted to leverage all the existing tools and libraries developed either in the OCSIGEN project or the larger OCAML ecosystem.

1.3 The design of a Modular Tierless language

To build such a language with support for both modular and tierless features, we outline several goals that will guide the design of our language.

Explicit communications. Manual annotations should determine whether a piece of code is to be executed server- or client-side. This design decision stems from our belief that the programmer must be well aware of where the code is to be executed, to avoid unnecessary remote interaction. Explicit annotations also prevent ambiguities in the semantics, allow for more flexibility, and enable the programmer to reason about where the program is executed and the resulting trade-offs. Programmers can thus ensure that some data stays on the client or on the server, and choose how much communication takes place.

A simple and efficient execution model. The language should not incur additional back-and-forth communications between client and server. Furthermore, the execution model should be simple and predictable. Having a predictable execution model is essential in the context of a language which is not purely functional, like OCAML.

Leveraging the type system. The language should leverage the type system to allow composition and modularity of client-server programs while preserving type-safety and abstraction. This ensures, via the type-system, that client functions are not called by mistake inside server code (and conversely) and ensures the correctness of client-server communications.

Integration with the host language. Programmers must be able to leverage both the language and the ecosystem of OCAML. OCAML libraries can be useful on the server, on the client or on both. As such, any OCAML file, even when compiled with the regular OCAML compiler, should be a valid module in our language. Furthermore, we should be able to specify if we want to use a given library on the client, on the server, or everywhere.

Modularity and encapsulation. Module and type abstractions are very powerful programming tools. By only exposing part of a library, the programmer can safely hide implementation details and enforce specific properties. The language should leverage module abstraction to provide encapsulation and separation of concern for widgets and libraries. By combining module abstraction and tierless features, library authors can provide good APIs that do not expose the fine-grained details of client-server communication to the users.

1.4 Contributions and Plan

We propose ELIOM, a tierless web programming language that supports static typing, an efficient compilation scheme that avoids extra communications and a very powerful form of modularity inspired by ML-style languages. In particular, ELIOM is the only tierless programming language featuring efficient static separate compilation. ELIOM extends OCAML and can transparently leverage its complete ecosystem. This article presents the design of ELIOM, how it can be used in practice, and formalizes the underlying core language. Our contributions are the following:

- We designed ELIOM, a tierless web programming language that follows the set of goals presented above. We demonstrate the practicality of our language through examples (Section 2).
- We formalized the type system and the semantics of ELIOM, including both the expression and the module layers (Section 3). This formalization provides an *interpreted* semantics that is easy

to reason with and can be explained to programmers, but still maintains ELIOM's good properties regarding communications and typesafety. Using this semantics, we show that ELIOM supports separate typechecking ([Theorem 1](#)) and complete integration with OCAML ([Section 4](#) and [Theorem 2](#)).

- As highlighted before, ELIOM is not an interpreted language. We formalized the compilation model of ELIOM ([Section 5.2](#)). This compilation scheme turns one ELIOM program into two simpler OCAML programs; one for the client and one for the server. We show that such compilation scheme supports separate compilation and preserves both the typing and the semantics of ELIOM ([Theorems 3](#) and [4](#)).
- We implemented ELIOM as a patch on the OCAML compiler and a runtime library¹.

2 PROGRAMMING WITH ELIOM

An ELIOM application is composed of a single program that has both client and server behaviors. During compilation, the ELIOM compiler decomposes the program into two parts. The first part runs on a Web server, and is able to manage several connections and sessions at the same time. The second part, compiled statically to JAVASCRIPT, is sent to each client by the Web server together with the HTML page, in response to the initial HTTP request.

In this section, we give a tour of the ELIOM language through various examples. As a guiding example for our exploration of the ELIOM language, we consider the case of a commenting system similar to websites such as Reddit or Hackernews. A comment is a piece of HTML written by a user and identified by a unique identifier. Comments are initially stored on the server. Users can then manipulate them on the client by voting, hiding or searching them. Such a library features both server aspects (storing and rendering the comments) and client interactions (browsing and searching comments). We start by giving a reminder of some important concepts on OCAML modules before presenting the various ELIOM concepts. Even though ELIOM is based on OCAML, little knowledge of OCAML is required. We explicitly write some type annotations for illustration purposes but they are not mandatory.

2.1 Of comments and camels – A short introduction to OCAML modules

The OCAML module system forms a second language separate from the expression language. While the language of expressions allows programming “in the small”, the module language allows programming “in the large”. The ML flavor of module systems, which OCAML is part of, significantly extend usual module languages by providing module types (called signatures) and functions from modules to modules (called functors). The module system is implicitly used for any kind of OCAML or ELIOM programming: Each `.ml` and `.eliom` file form a structure containing the list of declarations included in the file. It is also possible to specify a signature for such module by adding a `.mli` or `.eliomi` file. We can do a lot more with OCAML modules. For example, let us say we are writing an HTML library. We want to gather the event related attributes in a single module. We can easily do so with the following construction.

```
1 module On = struct
2   let click = ...
3   let keypress = ...
4 end
```

These functions can then be used through qualified accesses:

```
1 open Html
2 let mywidget = div ~a:[On.click myclickhandler] [ ... ]
```

¹<https://github.com/ocsigen/ocaml-eliom> and <https://github.com/ocsigen/eliomlang>

Some users of our HTML library may want to experiment with new, custom-made HTML elements. They can easily do so by extending the `Html` module:

```
1 module HtmlPlus = struct
2   include Html
3   let blink elems = ...
4 end
```

Here, we declare a new module, `HtmlPlus`, in which we *include* `Html` and define the new `blink` function. The `include` operation simply takes all the fields of a module and adds them to the enclosing module. This way, we obtain a new module `HtmlPlus` which can be used anywhere `Html` can, but also includes the new function.

2.1.1 Abstraction and encapsulation. We now want to build a simple library to handle internet comments. In our library, comments are pieces of HTML (constructed with the `Html` module) identified by a unique number. We are not sure yet if we should use simple sequential IDs, date-based IDs or something else like UUIDs and [Hashids \[2017\]](#). Fortunately, we do not have to make this decision immediately! All we need in order to write the rest of our library is an interface for creating and using identifiers. We can declare such an interface in OCAML using a *signature*. In [Section 2.1.2](#), we declare the ID signature describing what a module implementing unique identifiers should look like. We then define two modules implementing this specification, `SequentialID` and `DateID`, which we can switch easily.

To the outer world, these two modules have exactly the same type and can not be distinguished. The type that implements the identifiers in the ID signature is abstract: its implementation is only visible inside the module and can not be used outside. It is also useful to note that such abstraction can be provided after the fact. Declaring a module and abstracting its interface are completely distinct operations.

Hiding the internal details of our ID modules is not only useful for modularity: it also allows to enforce abstraction boundaries. For example in the case of `SequentialID`, it is impossible to inadvertently use the ID as an integer, since the fact that it is an integer is not revealed! We can use this fact to enforce numerous complex properties, as we see in the rest of this section.

2.1.2 Functors. To implement our comment system, we sometimes need to find comments by their ID. The idiomatic OCAML solution is to use maps, also called dictionaries. Such maps are implemented with Binary Search Trees which require a comparison function on the keys of the map. `Map.Make` is a pre-defined functor in the OCAML standard library that takes a module implementing the `COMPARABLE` signature as argument and returns a module that implements dictionaries whose keys are of the type `t` in the provided module. In [Figure 3](#), we use this functor to create the `IDMap` module which defines dictionaries with IDs as keys. This is very easy, since the ID signature is already a super-set of the `COMPARABLE` signature. We then define `register`, a function which associates a fresh id to a comment `c`.

```
1 module type ID = sig
2   type t (* type of ids *)
3
4   val compare : t -> t -> int
5   val create : unit -> t
6   val to_string : t -> string
7 end
1 module SequentialID : ID = struct
2   type t = int
3
4   (* ... *)
5 end
1 module DateID : ID = struct
2   type t = date
3
4   (* ... *)
5 end
```

Fig. 1. The ID signature and two implementations.

The `Map.Make` functor uses abstraction in two important ways. First, since the type of the map is abstract, it is impossible to modify it through means not provided by the module. In particular, this enforces that the binary tree is always balanced. Second, since the comparison function is provided in advance by the argument of the functor, it is impossible to mix different comparison functions by mistake. Indeed, application of the functor to different modules would yield different types of maps.

```

1 module type COMPARABLE = sig
2   type t
3   val compare : t -> t -> int
4 end
5
6 module Make (Key : COMPARABLE) : sig
7   type 'a t
8   val empty : 'a t
9   val add : Key.t -> 'a -> 'a t -> 'a t
10  (* ... *)
11 end

```

```

1 module TheID = DateID (* The ID of our choice *)
2 module IDMap = Map.Make(TheID)
3
4 let register c map : Html.t IDMap.t =
5   let commentid = TheID.create () in
6   IDMap.add commentid c map

```

Fig. 3. Dictionaries from IDs to comments

Fig. 2. the `Map` module

2.2 Tierless widgets

Until now, we presented how to write various elements of libraries useful for our comment system by leveraging the power of the OCAML module system in various ways. We now want to write the widget that presents a comment. We thus need to define both client and server code, along with some client-server communication, which is precisely where tierless languages shine.

2.2.1 Declarations. Locations in ELIOM are explicit. Each declaration must be marked with an annotation that specifies whether a declaration is to be performed on the server or on the client as follows:

```

1 let%server s = ...
2 type%client t = ...

```

Every declaration, such as types, values and modules, can be annotated. These annotations allows to group related code in the same file, regardless of where it is executed. In the rest of this article, we use the following color convention: client is in **yellow** and server is in **blue**. Colors are however not mandatory to understand the rest of this article.

2.2.2 Fragments. While location annotations allow programmers to gather code across locations, they don't allow convenient communication. For this purpose, ELIOM allows to include client-side expressions inside server declarations: an expression placed inside `[%client ...]` will be computed on the client when the page is received; but the eventual client-side value of the expression can be passed around immediately as a black box on the server. These expressions are called *client fragments*.

In the example below, the expression `1 + 3` will be evaluated on the client, but it's possible to refer server-side to the future value of this expression (for example, put it in a list). The variable `x` is only usable server-side, and has type `int fragment` which should be read "a fragment containing some integer". The value inside the client fragment cannot be accessed on the server.

```

1 let%server x : int fragment = [%client 1 + 3 ]

```

2.2.3 Injections. Fragments allow programmers to manipulate client values on the server. We also need the opposite direction. Values that have been computed on the server can be used on the client by prefixing them with the symbol `~%`. We call this an *injection*.

```
1 let%server s : int = 1 + 2
2 let%client c : int = ~%s + 1
```

Here, the expression `1 + 2` is evaluated and bound to variable `s` on the server. The resulting value `3` is transferred to the client together with the Web page. The expression `~%s + 1` is computed client-side.

An injection makes it possible to access client-side a client fragment which has been defined on the server. The value inside the client fragment is extracted by `~%x`, whose value is `4` here.

```
1 let%server x : int fragment = [%client 1 + 3 ]
2 let%client c : int = 3 + ~%x
```

2.2.4 Comment widget. These three constructions are sufficient to create complex client-server interactions such as the comment widget. The comment widget shows a single comment and can be used several times to show a list of comments. It not only shows the author and the content of the comment, but will also hide the content when the user clicks on it. Finally, we want the HTML to be generated server-side and sent to the client as a regular HTML page, which allows the comments to be accessible even when JAVASCRIPT cannot run. The implementation, the interface and the produced HTML fragment are shown in [Figure 4](#).

In order to implement our comment widget, we use an HTML DSL [[TyXML 2017](#)] that provides combinators such as `div` and `a_onclick` (which respectively create an HTML tag and an HTML attribute). In OCAML, `~a` denotes a named argument that is used here to provide the list of HTML attributes. We first create a `p` element which contains the text of the comment and a unique id. The text is included in a `div` which represents the comment. We then use a handler listening to the `onclick` event: since clicks are performed client-side, this handler needs to be a client function inside a fragment. Inside the fragment, an injection is used to access the argument `id` which contains the identifier of the comment. We then use this identifier to fetch the correct element and toggle the “hidden” CSS property, which hides it.

The signature of our widget function, shown [Figure 4b](#), does not expose the internal details of the widget’s behavior. In particular, the communication between server and client does not leak in the API: This provides proper encapsulation for client-server behaviors. Furthermore, this widget is easily composable: the embedded client state cannot affect nor be affected by any other widget and can be used to build larger widgets.

2.2.5 Notes on semantics. In the examples above, we showed that we can interleave client and server expressions and communications in fairly arbitrary manners. This would be costly if the communication between client and server were done naively.

Instead, the server only sends data once when the Web page is sent. In particular, in the comment widget presented above, the `id` of the comment is not sent for each click. This is made possible by the fact that client fragments are not executed immediately when encountered inside server code. Intuitively, the semantics is the following. When the server code is executed, the encountered client code is not executed right away; instead it is just registered for later execution. The client code is executed only once the Web page has been sent to the client. We also guarantee that client code, be it either client declarations or fragments, is executed in the order that it was encountered on the server.

This presentation might makes it seem as if we dynamically create the client code during execution of the server code. This is not the case. Like OCAML, ELIOM is statically compiled and separates

```

1 open Html
2 let%server make_comment commentid =
3   let content =
4     p ~a:[a_id commentid] [text (Comment.get commentid)]
5   in
6   let author =
7     text ("Author: " ^ Comment.author commentid)
8   in
9   let handler = [%client fun _ ->
10     let elem = get_element_by_id ~%commentid in
11     Css.toggle_hidden elem]
12 in
13 div ~a:[On.click handler] [author; content]

```

(a) Implementation

```

1 val%server make_comment :
2   id -> Html.element

```

(b) Interface

```

1 <div>
2   Author: The Ocsigen Team
3   <p id=42>I'm a composable widget!</p>
4 </div>

```

(c) Resulting HTML

Fig. 4. The comment widget

client and server code at compile time. During compilation, we statically extract the code included inside fragments and compile it as part of the client code to JAVASCRIPT. This allows us to provide both an efficient execution scheme that minimizes communication and preserves side-effect orders while still presenting an easy-to-understand semantics. We also benefit from optimizations done by the `JS_OF_OCAML` compiler, thus producing efficient and compact JAVASCRIPT code.

2.3 Hybrid data-structures

We want to add some buttons to our comment widget for various possible actions on comments such as “reply”, “permalink”, “report”, etc. We want our buttons to be created in a uniform manner as a list of symbolic actions. Some of our actions are normal links and some are client-side actions. Thanks to fragments, we can create hybrid data-structures that have both client and server parts. **Figure 5** implements such a hybrid data-structure to specify button actions. The `action` type is a server type that is either a normal link or a client-side action (here, a function from `unit` to `bool`) in a fragment. The `attr_of_action` function turns actions into HTML attributes, either a link or an on-click handler. Finally, `button_list` walks through a list of pairs of names and actions and returns an unordered list of `a` elements.

Creating such hybrid data-structures allows a great flexibility in how ELIOM code is organized by allowing to meld client and server code at any point in the application. Having explicit annotations for client and server code of is essential here. This would be quite difficult to achieve if the delimitation between client and server values were implicitly inferred.

```

1 type%server action =
2   | Link of Url.t
3   | Action of (unit -> bool) fragment
4
5 let%server attr_of_action action = match action with
6   | Link l -> a_href l
7   | Action f -> On.click f
8
9 let%server button_list (lst : (name * action) list) =
10   ul (List.map (fun (name, action) ->
11     li [a ~a:[attr_of_action action] [pdata name]])
12     lst)

```

Fig. 5. Function generating a list of buttons


```

1 type%server ('i,'o) t
2 type%client ('i,'o) t = 'i -> 'o
3
4 val%server create : ('i -> 'o) -> ('i, 'o) t

```

Fig. 6. Rpc signature

```

1 let%server fetch_comments
2   : (page_id, Html.t list) Rpc.t
3   = Rpc.create Comment.from_page
4
5 let%client button_load parent page_id =
6   let handler _ =
7     let comments = ~%fetch_comments page_id in
8     List.iter Dom.append_child parent comments
9   in
10  Html.button
11    ~a:[On.click handler]
12    [text "Load more comments"]

```

Fig. 7. Dynamic loading of comments with Rpc

2.3.1 Dynamic loading of comments. In order to handle pages with many comments, we want to only present a subset of the comment upfront and dynamically load additional comments when the user asks for it (or, alternatively, when the page scrolls). This requires dynamic communications which are not directly provided by fragments and injections, as indicated in the previous section. For such purposes, ELIOM provides the `Rpc` module whose API is presented in [Figure 6](#).

We implement a button which dynamically loads new comments in [Figure 7](#). We first create server-side an RPC endpoint `fetch_comments` with the function `Rpc.create`. This endpoint is of type `(page_id, Html.t list)Rpc.t`: it takes as argument the id of the current page and return the list of associated comments using the pre-defined function `Comments.from_page`. The type `Rpc.t` is abstract on the server, but is a synonym for a function type on the client. Of course, this function does not contain the actual implementation of the RPC handler, which only exists server-side. To use this API, we leverage injections. By using an injection on `~%fetch_comments` on [Line 7](#), we obtain *on the client* a value of type `Rpc.t`. We describe the underlying machinery that we leverage for converting RPC endpoints into client-side functions in [Section 2.3.2](#). What matters here is that we end up with a function that we can call like any other; calling it performs the remote procedure call. Once we have a way to fetch new comments, we can create a button `button_load` which, when clicked, loads the comments. This is done by creating a client-side handler and using the `On.click` function.

The RPC API proposed in [Figure 6](#) is “blocking”: the execution waits for the remote call to finish before pursuing, thus blocking the rest of the client program. Remote procedure calls should be made asynchronously: the client program keeps running while the call is made and the result is used when the communication is done. The actual implementation uses LWT [[Vouillon 2008](#)] to express asynchronous calls in a programmer-friendly manner through promises. The use of LWT is pervasive in the ELIOM ecosystem both on the server and on the client. In this article, we omit mentions of the LWT types and operators for pedagogic purposes.

2.3.2 Converters. In the RPC API, we associate two types with different implementation on the server and on the client. We rely on injections to transform the datastructure when moving from one side to the other. This ability to transform data before sending it to the client via an injection is made possible by the use of *converters*. [Figure 8](#) broadly presents the converter API. Given a serialization format `serial`, a converter is a pair of a *server* serialization function and a *client* de-serialization function. Note that the client and server types are not necessarily the same. Furthermore, we can arbitrarily manipulate the value before returning it. Several predefined converters are available for fragments, basic OCAML datatypes, and tuples in the module `Conv`. Implementation details about converters can be found in [Section 3.3.2](#).

```

1 type serial (* A serialization format *)
2 type%server ('a, 'b) converter = {
3   serialize : 'a -> serial ;
4   deserialize : (serial -> 'b) fragment ;
5 }

```

Fig. 8. Schematized API for converters

```

1 type%server ('i, 'o) t = {
2   url : string ;
3   handler: 'i -> 'o ;
4 }
5 type%client ('i, 'o) t = 'i -> 'o
6
7 let%server serialize t = serialize_string t.url
8 let%client deserialize x =
9   let url = deserialize_string x in (fun i -> XmlHttpRequest.get url i)
10
11 let conv = { serialize ; deserialize = [%client deserialize] ; }
12
13 let%server create handler =
14   let url = "/rpc/" ^ generate_new_id () in
15   serve url handler ;
16 { url ; handler }

```

Fig. 9. Simplified RPC implementation corresponding to Figure 6.

We can use converters to implement the RPC API (Figure 9). The server implementation of `Rpc.t` is composed of a handler, which is a server function, and a URL to which the endpoint answers. Our serialization function only sends the URL of the endpoint. The client de-serialization function uses this URL to create a function performing an HTTP request to the endpoint. This way, an RPC endpoint can be accessed simply with an injection. Thus, for the `create` function, we assume that we have a function `serve` of type `string -> (request -> answer) -> unit` that creates an HTTP handler at a specified URL. When `Rpc.create` is called, a unique identifier `id` is created, along with a new HTTP endpoint `"/rpc/id"` that invokes the specified function.

This implementation has the advantage that code using the `Rpc` module is completely independent of the actual URL used. The URL is abstracted away. Converters preserve abstraction by only exposing the needed information.

2.4 Modular Tierless Programming

We are now equipped with two tools. On one hand, we have a rich and expressive non-tierless module system, as presented in Section 2.1, which provides abstraction and modularity at the library level. On the other hand, we have a powerful tierless programming language, as presented in Section 2.2, which allows us to describe sophisticated client-server behaviors. In this section, we present how we can bring those two tools together and reap the numerous benefits of the OCAML module system in a tierless setting.

2.4.1 Interaction with OCAML. In the previous examples, we freely used OCAML modules in server and client contexts. In Figure 8, we even defined an API that has both normal OCAML declarations and server declarations! Web programming is never only about the Web. Web programmers need external libraries and a rich ecosystem that can not be provided by a fresh new language. The ability to transparently leverage the OCAML ecosystem in ELIOM is essential to write web applications productively.

This close integration is provided through the use of a third location called **base**. Code located on base can be used both on the client and on the server.

```
1 let%base f x = "Hello " ^ x ^ "!"
2 let%client a = f "client"
3 let%server b = f "server"
```

ELIOM-specific features such as fragments and injections are not allowed inside base code. In fact, base code corresponds exactly to OCAML code. This equivalence holds in theory but also in practice, meaning that any OCAML library compiled by the vanilla OCAML compiler can be directly reused by ELIOM as being on the base location. This allows a very smooth integration with the OCAML ecosystem. Furthermore, a given OCAML library can be loaded either on base, on the client or on the server, depending on what the user wants. For example, an OCAML library manipulating file descriptors might be better kept only on the server in order to avoid misuse. The type-checker then raises an error if the library is mistakenly used on the client.

2.4.2 Modules and locations. Since OCAML modules, such as the `Html` module defined earlier, are immediately available as ELIOM modules located on base, we can also use such modules on the client or on the server.

```
1 module%base TextHtml = Html
2 module%server ServerHtml = TextHtml
3 let%client l = Html.p [Html.text "Hello client!"]
```

Locations are checked by the compiler. For example, using a server module on the client is forbidden.

```
1 let%client x = ServerHtml.text "hello client!" (* X Error! *)
```

It is also possible to reuse OCAML module types freely. For example, we might want to define a client module `DomHtml` which shares the exact same API as the `Html` module, but is implemented using the Document Object Model that is available on the client. The type declaration for such a module would then be very simple, as shown below.

```
1 module%client DomHtml : Html.Signature = struct
2   (* ... *)
3 end
```

We can easily declare a new structure completely on one location. The constraint is that all the fields on such modules, including submodules, should be on the same location. For example, a client structure can only contain fields that are declared on the client. The following piece of code declares a `JsMap` client module containing various fields and implementing a dictionary data-structure with `JAVASCRIPT` strings.

```
1 module%client JsMap : sig
2   type 'a t
3
4   val empty : 'a t
5   val add : Js.string -> 'a -> 'a t -> 'a t
6   (* ... *)
7 end
```

We can also use functors in client and server code as we would in regular OCAML code. Consider the `JsMap` module above. The simplest way to obtain such a module would be to use the `Map.Make` functor presented in [Section 2.1.2](#). We could for example write a `JsDate` module which uses `JAVASCRIPT` native support for dates. We can then obtain the `JsDateMap` module simply by applying `Map.Make` to the module `JsDate` defined in [Figure 10](#). As expected, the module we obtain is directly on the client. We can thus mix and match client and server modules using the tierless

features and vanilla OCAML modules. This also works with all the other module features such as abstraction, high order functors and module inclusion. In all these cases, the ELIOM typechecker ensures that modules always end up on the appropriate location.

```

1 module%client JsDate = struct
2   type t = Js.date
3
4   (** Compare by timestamp *)
5   let compare x y = compare x##valueOf y##valueOf
6   end
7
8 module%client JsDateMap = Map.Make(JsDate)

```

Fig. 10. Definition of JsDate and JsDateMap

2.4.3 Abstraction and encapsulation across locations. A common idiom of web programming is to generate some HTML element on the server, add an `id` to it, and recover the element on the client through the `get_element_by_id` function. Indeed, this is exactly what we did in our comment widget in [Section 2.2.4](#). This is so common, in fact, that it could be considered the “id design pattern”. RPCs, channels and other communication APIs also follow the same mechanisms through the use of uniquely defined URLs. In all these cases, the means of identification for a given object is generally passed around directly, as a string, instead of being abstracted. Since client and server code are usually written separately, the programmer *must* expose the internal details to the outer world, including how to identify objects.

One solution, which we used in [Section 2.3.2](#), is to use converters and only use the ID as an implementation detail of the client-server communication. However, it is sometimes beneficial to keep an explicit ID. By combining tierless annotations and the abstraction capabilities provided by modules, we can have explicit ID that are abstract and type-safe. [Figure 11](#) presents an API that encapsulates unique ids for HTML elements. This new `HtmlID` is **mixed** and can contain client, server and base declarations. It is composed of an OCAML abstract type, `id`, and two operations. The server function `with_id` takes an HTML element, generates a fresh `id` and returns a pair composed of the HTML element with that `id` and the `id` itself. The client function `find` takes an `id` and retrieves the associated element as a DOM node on the client. The `id` type is abstract. Both the client and the server functions can use the real definition of `id` since they are both inside the module. The outer world, however, can not. Mixed modules allow abstraction to extends over the client-server boundary. This can provide further benefits in the case of more complex data-structures, as we will see in the next section.

2.4.4 Multi-tiers comments. We now want to implement a system of client-side search and filtering of comments. The user should be able to search and filter comments directly on the client, without the need to reload the page. For this purpose, we need to maintain the sets of comments both on the server and on the client. One simple way to do that is to create a replicated cache of comments which ensures that all the comments available on the server are also available on the client.

We use the `Map` module as inspiration and create a functor that takes as argument a module describing the keys. The idea is that adding an entry to a server-side table also adds the element to the client-side table. Consequently, the server-side representation of a table needs to include a client-side one.

```

1 module%mixed HtmlID : sig
2   type id
3
4   val%server with_id : Html.t -> Html.t * id
5
6   val%client get : id -> DomHtml.t
7 end

```

(a) Interface

```

1 module%mixed HtmlID = struct
2   type id = string
3
4   let%server with_id elem =
5     let myid = random_string () in
6     let elem_with_id = Html.add_id elem myid in
7     (elem_with_id, myid)
8
9   let%client find myid = get_element_by_id myid
10 end

```

(b) Implementation

Fig. 11. Abstract HTML ids for client/server communications

The result API is shown in [Figure 12](#). The resulting module contains both a client and a server side types, both named 'a table, which represent the local table. The module also exposes traditional Map functions. The implementation, shown in [Figure 13](#), is more interesting. We exploit the fact that client and server namespaces are distinct, and name both client and server map modules M. On the server, the cache is implemented as a pair of a server-side and a client-side dictionary. The server-side add implementation stores a new value locally in the expected way, but additionally builds a fragment that has the side-effect of performing a client-side addition. The retrieval operation (find) returns a shared value that contains both the server side version and the client side. On the client, however, we can directly use the local values. Since the client-side type exactly corresponds to a regular map, we can directly use the usual definitions for the various map operations. This is done by including the client M module on the client.

Note that this functor cannot be implemented in a decomposed way without sacrificing either abstraction or modularity. Indeed, the server implementation relies on the client-side version of the functor argument (Comparable) to implement proper usage of the keys. Furthermore, the signature of the functor ensures that the server-side and client-side parts of the cache are in sync without leaking any implementation details. Separating this mixed functors in two would require exposing the guts of the data-structure. Abstraction also makes it easy to extend such modules with new features. For example, it would possible to add full-blown replication through “push” or “pull” communications between the client and the server. Thanks to the abstraction provided by the signature of the module, this can even be done while keeping the API of the functor unchanged.

We can now use this cache for our comment system by using, for example, the DateID module for the keys. This is done in [Figure 14](#). Adding a new comment to the page is done through the add_comment server function. This function creates the associated HTML using the widget defined in [Section 2.2.4](#) and adds it to the cache. We can then create the webpage containing all the comments simply by collecting all the comments and putting them inside a div. This is done by the generate_page server function. Finally, the client function filter_comments filters the shown comments on the client. It takes as argument a predicate function and the current client cache. It uses this predicate function to filter the cache, using the function CommentCache.filter, which directly uses the equivalent function from the Map module. We then find the HTML element containing all the elements and replace them by the updated list. Since the CommentCache module also contains a client-side add function, this approach works very well the dynamic loading of comment presented in [Section 2.3.1](#): Once the RPC returns the new comments, we can add them to the client-side cache directly.

```

1 module%mixed MakeCache (Key : COMPARABLE) : sig
2   type%client 'a t
3   type%server 'a t
4
5   val%client add : Key.t -> 'a -> 'a t -> 'a t
6   val%server add : Key.t -> 'a -> 'a t -> 'a t
7   (* ... *)
8 end

```

Fig. 12. Interface of MakeCache

```

1 module%mixed MakeCache (Key : COMPARABLE) = struct
2   module%client M = Map.Make(Key)
3   module%server M = Map.Make(Key)
4
5   include%client M
6
7   type%server 'a table = 'a M.t * 'a M.t fragment
8   let%server add id v (tbl_server, tbl_client) =
9     [%client M.add ~%id ~%v ~%tbl_client ];
10    M.add id v tbl_server
11   (* ... *)
12 end

```

Fig. 13. Implementation of MakeCache

```

1 module%mixed DateKey = DateID
2 module%mixed CommentCache = MakeCache(DateKey)
3
4 let%server add_comment id cache =
5   let html = make_comment id in
6   CommentCache.add id html cache
7
8 let%server generate_page cache =
9   Html.div
10    ~a:[a_id "comments"]
11    [CommentCache.elements cache]
12
13 let%client filter_comments predicate cache =
14   let filtered_cache =
15     CommentCache.filter predicate cache
16   in
17   let comment_container =
18     get_element_by_id "comments"
19   in
20   Dom.replace_children
21     comment_container
22     (CommentCache.elements filtered_cache)

```

Fig. 14. Using MakeCache

A note on mixed functors. Mixed functors unfortunately have some limitations. Arguments must themselves be mixed modules and injections inside client-side bindings can only reference elements outside of the functor. Additionally, there are some limitations regarding nesting of mixed structures and functors. These limitations are formalized and discussed more precisely in [Section 3.4.2](#). Mixed functors are nevertheless sufficient for a large class of complex client-server APIs that are useful in practice, as demonstrated with the MakeCache functor.

2.5 Going further

Through these various examples, we demonstrated how we can combine traditional tierless features with an advanced module system to create powerful and expressive APIs. On one hand, tierless languages traditionally allows for complex interplay of client and server code. Module systems, on the other hand, allows to manipulate large pieces of code while preserving abstraction, encapsulation and modularity. ELIOM allows to preserve these abstraction capabilities while enjoying the free-form tierless programming style. The OCSIGEN ecosystem uses these concepts to provide numerous additional tools such as advanced communication patterns, shared HTML across tiers or even distributed Functional Reactive Programming [[Radanne et al. 2016a](#); [Radanne and Vouillon 2018](#); [Tutorial 2017](#)]. All these advanced mechanisms are used in production by ELIOM users.

3 THE ELIOM_ε CALCULUS

We now formalize ELIOM as an extension of ML with both an expression and a module language. To emphasize the new elements introduced by ELIOM, these additional elements will be colored in blue. This is only for ease of reading and is not essential for understanding the formalization.

While ELIOM is an extension of OCAML, ELIOM_ε is an extension of a simpler ML calculus, which we present quickly in [Section 3.1](#). We then present the location system in [Section 3.2](#), followed by the expression and the module languages in [Sections 3.3](#) and [3.4](#). Finally, we describe its semantics in [Section 3.6](#).

Syntactic considerations. Let us first define some notations and meta-syntactic variables. As a general rule, the expression language is in lowercase (e) and the module language is in uppercase (M). Module types are in calligraphic letters (\mathcal{M}). More precisely: e are expressions, x are variables, p are module paths, X are module variables, τ are type expressions and t are type constructors. x_i , X_i and t_i are identifiers (for values, modules and types). Identifiers (such as x_i) have a name part (x) and a stamp part (i) that distinguish identifiers with the same name (following Leroy [1995]): α -conversion should keep the name intact and change only the stamp. Sequences are noted with a star; for example τ^* is a sequence of type expressions. Indexed sequences are noted (τ_i) , with an implicit range. Substitution of a by b in e is noted $e[a \mapsto b]$. Repeated substitution of each a_i by the corresponding b_i is noted $e[a_i \mapsto b_i]_i$.

3.1 A crash course in the ML calculus

There are many variants of the ML calculus. For the purpose of ELIOM_e , we consider a core calculus with polymorphism, let bindings and parametrized datatypes in the style of Wright and Felleisen [1994], accompanied by a fully featured module system with separate compilation and applicative functors in the style of Leroy [1994, 1995]. The goal of this calculus is to closely module the most relevant features of OCAML, while being sufficiently simple to be extended and reasoned on formally.

In this section, we will simply give a quick overview of the syntax and an informal reminder of the most distinctive or unusual points. The complete treatment of the language is given in Appendix A. Each set of typing or reduction rules can also be obtained by considering the ELIOM_e rules in the next section and ignoring the parts written in blue.

3.1.1 Syntax. The complete syntax is presented in Figure 42 and follows the syntax of OCAML closely. The expression language is a fairly simple extension of the lambda calculus with a fixpoint combinator (Y) and let bindings ($\text{let } x = e_1 \text{ in } e_2$). The language is parametrized by a set of constants *Const*. Variables can be qualified by a module path p . Paths can be either module identifiers such as X_i , a submodule access such as $X_i.Y$, or a path application such as $X_i(Y_j.Z)$. Note that, as said earlier, that fields of modules are only called by their name, without stamp.

The module language is composed of functors, module application ($M_1(M_2)$), type constraints ($(M : \mathcal{M})$), and structures containing a list of value, types or module definitions ($\text{struct let } x_i = 2 \text{ end}$). Programs are lists of definitions. Module types can be either functor types ($\text{functor}(X_i : \mathcal{M}_i)\mathcal{M}_2$) or a signature ($\text{sig val } x_i : \text{int end}$), which contains a list of value, types and module descriptions. Type descriptions can expose their definition or can be left abstract. Typing environments are simply module signatures. We note them Γ for convenience.

3.1.2 Type system. As a general rule, judgements are denoted with the symbol \triangleright for expressions and \blacktriangleright for modules. More precisely:

$\Gamma \triangleright e : \tau$	Expression e has type τ in Γ
$\Gamma \triangleright \tau_1 \approx \tau_2$	τ_1 and τ_2 are equivalent in Γ
$\Gamma \blacktriangleright M : \mathcal{M}$	module M has type \mathcal{M} in Γ
$\Gamma \blacktriangleright \mathcal{M} <: \mathcal{M}'$	Module type \mathcal{M} is a subtype of \mathcal{M}' in Γ
$\Gamma \vDash \tau$ and $\Gamma \vDash \mathcal{M}$	The type τ (resp. module type \mathcal{M}) is well formed.

The type system for expression is the standard Hindley-Milner type system with polymorphism, generalization and parametrized datatype. The module system, which follows Leroy [1994, 1995], is less usual. We will focus on three specific features: qualified accesses, applicative functors and strengthening.

Qualified accesses. Let us consider the module X with the module type (sig type t ; val $a : t$ end). X contains a type t and a value a of that type. We wish to typecheck $X.a$. One expected type for this expression is $X.t$. However, the binding of v in X gives the type t , with no mention of X . In order to make the type of $X.a$ expressible in the current scope, we prefix the type variable t by the access path X . This is done in the rule QUALMODVAR by the substitution $\mathcal{M}[n_i \mapsto p.n \mid n_i \in \text{BV}(\mathcal{S}_1)]$ which prefixes all the bound variables of \mathcal{S}_1 , noted $\text{BV}(\mathcal{S}_1)$, by the path p . We thus obtain the following type derivation

$$\text{QUALVAR} \frac{\text{MODVAR} \frac{(\text{module } X : \dots) \in (\text{module } X : \dots)}{(\text{module } X : \dots) \blacktriangleright X : \text{sig type } t; \text{ val } a : t \text{ end}}}{(\text{module } X : \dots) \blacktriangleright X.a : X.t} \text{ WITH } X.t = t[t \mapsto p.t]$$

Applicative functors. In [Section 2.1](#), we used the `Map.Make` functor to create a new dictionary data-structure using the provided type as keys. One might wonder what happens if `Map.Make` is applied to the same module twice. Are the dictionaries thus produced compatible? More formally, Does $M_1 = M_2$ implies $F(M_1) = F(M_2)$? When this holds, we say that functors are “applicatives”. Otherwise, we say they are “generatives”. OCAML functors are applicatives by default. In our simple ML calculus, all functors are applicatives. This is implemented by enriching type constructors with types of the form $F(M).t$. If we consider the modules $N_i = \text{Map.Make}(M_i)$, then we have $N_1.t = \text{Map.Make}(M_1).t = \text{Map.Make}(M_2).t = N_2.t$.

Strengthening. Let us consider a module X answering the signature shown [Example 20a](#). We want to pass this module to the functor F shown in [Example 20b](#). This functor application is expected to succeed, since $X.t = t$, where t comes from X . However, the definition of t in X is abstract, and does not have a definition allowing us to check that X is indeed compatible with the proposed signature. For this purpose, we *strengthen* the type of the module X with additional type equalities. If X is the type of X , we note X/X the strengthened signature. The result is the signature (sig type $t = X.t$; val $v : t$ end). This signature is now trivially included in the argument type of F .

<pre> module X : sig type t val v : t end </pre>	<pre> module F (E : sig type t = X.t val v : t end) = struct ... end module X = F(X) </pre>
(a) Typing environment	(b) Application of multi-argument functor using manifests

Fig. 15. Program using functors and manifest types

3.1.3 Semantics. We use a rule-based big-step semantics with traces. Big-step semantics are more amenable to reasoning with modules, as they allow to easily introduced local bindings. We record traces in order to reason about execution order. We note v for values in the expression language and V for values in the module language. Values in the expression language can be either constants or lambdas. Module values are either structures, which are list of bindings of values, or functors. Execution environments, noted ρ , are a list of value bindings. We note the concatenation of environment $+$. Environment access is noted $\rho(x) = v$ where x has value v in ρ . The same notation is also used for structures. Traces are lists of messages. The empty trace is noted $\langle \rangle$. Concatenation of traces is noted $@$.

$$\ell ::= s \mid c \mid b \qquad \zeta ::= m \mid \ell$$

Fig. 16. Grammar of locations – ℓ and ζ

$$m > s \qquad m > c \qquad b > s \qquad b > c \qquad b > m \qquad \forall \ell \in \{s, c, m, b\} \ell > \ell$$

Fig. 17. “can be used in” relations on locations – $\zeta > \zeta'$

$$m <: s \qquad m <: c \qquad m <: b \qquad \forall \ell \in \{s, c, m, b\} \ell <: \ell$$

Fig. 18. “can contain” relation on locations – $\zeta <: \zeta'$

Given an expression e (resp. a module M), an execution environment ρ , a value v (resp. V) and a trace θ ,

$$e \xRightarrow{\rho} v, \theta$$

means that e reduces to v in ρ and prints θ . The rules are composed of traditional call-by-value lambda-calculus with a fixed execution order that is represented by the traces.

3.1.4 Formal results. The ML calculus considered here has been shown to support for separate compilation ([Theorem 5](#)) and representation independence [[Leroy 1994](#)]. It doesn't have a proof of soundness, which discussed in [Appendix A.3.3](#).

3.2 Locations

Before introducing the ELIOM_ε language constructs, we consider the notation for *locations* annotations which indicate “where the code runs”. The grammar of locations is given in [Figure 16](#). There are three core locations: **server**, **client** or **base**. The base side represents expressions that are “location-less”, that is, which can be used everywhere. We use the meta-variable ℓ for an unspecified core location. There is a fourth location that is only available for modules: **mixed**. A mixed module can have client, server and base components. We use the meta-variable ζ for locations that are either **m** or one of the core locations. In most contexts, locations are annotated with subscripts.

We also introduce two relations:

$\zeta > \zeta'$ defined in [Figure 17](#), means that a variable (either values, types or modules) defined on location ζ can be used on a location ζ' . For example, a base type can be used in a client context. Base declarations are usable everywhere. Mixed declarations are not usable in base code.

$\zeta <: \zeta'$ defined in [Figure 18](#), means that a module defined on location ζ can contain component on location ζ' . In particular, the mixed location m can contain any component, while other location can contain only component declared on the same location.

Both relations are reflexive: For instance, it is always possible to use client declarations when you are on the client.

3.3 Expressions

ELIOM_ε 's expression language is based on a simple ML language, as presented in [Section 3.1](#), extended with ELIOM specific constructs: fragments and injections. This expression language can be seen as an alternative formulation of the core language proposed by [Radanne et al. \[2016b\]](#), but that is compatible with modules.

A *client fragment* $\{\{ e \}\}$ can be used on the server to represent an expression that will be computed on the client, but whose future value can be referred to on the server.

An injection $f\%v$ can be used on the client to access values defined on the server. An injection must make explicit use of a converter f that specifies how to send the value. In particular, this should involve a serialization step, executed on the server, followed by a deserialization step executed on the client. For ease of presentation, injections are only done on variables and constants. In the implementation, this restriction is removed by adding a lifting transformation. For clarity, we sometime distinguish injections *per se*, which occur outside of fragments, and escaped values, which occur inside fragments. The syntax of types is also extended with two constructs. A *fragment type* $\{\tau\}$ is the type of a fragment. A *converter type* $\tau_s \rightsquigarrow \tau_c$ is the type of a converter taking a server value of type τ_s and returning a client value of type τ_c . All type variables α_ℓ are annotated with a core location ℓ . There are now three sets of constants: client, server and base.

3.3.1 Typing rules. Typing judgements for ELIOM_ℓ are annotated with a location that specifies where the given code should be typechecked. For the expression language, we only consider core locations ℓ which are either base, client or server. The typing judgment, defined in Figure 20, is noted $\Gamma \triangleright_\ell e : \tau$ where e is of type τ in the environment Γ on the location ℓ . We also define type validity and equivalence judgements in Figures 21 and 22. We note $\text{TypeOf}_\ell(c)$ the type of a given constant c on the location ℓ . Binding in typing environments, just like in signatures, are annotated with a location. The first three kind of bindings, corresponding to the core language, can only appear on core locations: s , c or b . Modules can also be of mixed location m . Names are namespaced by locations, which means it is valid to have different client and server values with the same name.

Most of the rules are straightforward adaptations of traditional ML rules and are described in Section 3.1. We focus on rules that demonstrate some particular features of ELIOM_ℓ .

Rule VAR contains a significant difference compared to the traditional ML rules. As described earlier, bindings in ELIOM_ℓ are located. Since access across sides are explicit, we want to prevent erroneous cross-location accesses. For example, client variables can not be used on the server. To use a variable x bound on location ℓ' in a context ℓ , we check that location ℓ' can be used in location ℓ , noted $\ell' > \ell$ and defined in Figure 17. Base elements b are usable everywhere. Mixed elements m are usable in both client and server.

Type variables and types constructors are also annotated with a location and follow the same rules. Using type variables from the client on the server, for example, is disallowed. Such constraints are enforced by the type validity and equivalence judgements through rules such as FRAGMENTVAL and CONVVAL which tracks locations across type expressions.

Expressions		Path	
$e ::= c$	(Constant)	$p ::= X_i \mid p.X \mid p_1(p_2)$	
$\mid x_i \mid p.x$	(Variables)	Type Schemes	
$\mid Y$	(Fixpoint)	$\sigma ::= \forall(\alpha_\ell)^*.\tau$	
$\mid (e e)$	(Application)	Type Expressions	
$\mid \lambda x.e$	(Function)	$\tau ::= \alpha_\ell$	(Type variables)
$\mid \text{let } x = e \text{ in } e$	(Let binding)	$\mid \tau \rightarrow \tau$	(Function types)
$\mid \{\{ e \}\}$	(Fragment)	$\mid (\tau^*)_{t_i} \mid (\tau^*)p.t$	(Type constructors)
$\mid f\%x$	(Injection)	$\mid \{\tau\}$	(Fragment)
$f ::= p.v \mid v_i \mid c$	(Converter)	$\mid \tau \rightsquigarrow \tau$	(Converter types)
$c \in \text{Const}$	(Constants)		

Fig. 19. Grammar for ELIOM_ℓ expressions

Rule **FRAGMENT** is for the construction of client fragments and can only be applied on the server. If e is of type τ on the client, then $\{\{ e \}\}$ is of type $\{\tau\}$ on the server. Since no other typing rule involves client fragments, it is impossible to deconstruct them on the server.

Rule **INJECTION** is for the communication from the server to the client and can only be applied on the client. If e is of type τ_s on the server and f is a converter of type $\tau_s \rightsquigarrow \tau_c$ on the server, then $f\%e$ is of type τ_c on the client. Note that injections always require a converter, which we describe in greater details now.

Common rules		
$\frac{\text{VAR} \quad (\text{val}_{\ell'} x : \sigma) \in \Gamma \quad \ell' > \ell \quad \sigma > \tau}{\Gamma \triangleright_{\ell} x : \tau}$	$\frac{\text{LAM} \quad \Gamma; (\text{val}_{\ell} x : \tau_1) \triangleright_{\ell} e : \tau_2}{\Gamma \triangleright_{\ell} \lambda x. e : \tau_1 \rightarrow \tau_2}$	$\frac{\text{CONST} \quad \text{TypeOf}_{\ell}(c) > \tau}{\Gamma \triangleright_{\ell} c : \tau}$
$\frac{\text{LETIN} \quad \Gamma \triangleright_{\ell} e_1 : \tau_1 \quad \Gamma; (\text{val}_{\ell} x : \text{Close}(\tau_1, \Gamma)) \triangleright_{\ell} e_2 : \tau_2}{\Gamma \triangleright_{\ell} \text{let } x = e_1 \text{ in } e_2 : \tau_2}$	$\frac{\text{EQUIV} \quad \Gamma \triangleright_{\ell} e : \tau_1 \quad \Gamma \triangleright_{\ell} \tau_1 \approx \tau_2}{\Gamma \triangleright_{\ell} e : \tau_2}$	
$\frac{\text{APP} \quad \Gamma \triangleright_{\ell} e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \triangleright_{\ell} e_2 : \tau_1}{\Gamma \triangleright_{\ell} (e_1 e_2) : \tau_2}$	$\frac{\text{Y}}{\Gamma \triangleright_{\ell} Y : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2}$	
$\frac{\text{QUALVAR} \quad \Gamma \triangleright_{\ell} p : (\text{sig } S_1; \text{val}_{\ell'} x_i : \tau; S_2 \text{ end}) \quad \ell' > \ell}{\Gamma \triangleright_{\ell} p.v : \tau[n_i \mapsto_{\ell} p.n \mid n_i \in \text{BV}_{\ell}(S_1)]}$		
Server rules	Client rules	
$\frac{\text{FRAGMENT} \quad \Gamma \triangleright_c e : \tau}{\Gamma \triangleright_s \{\{ e \}\} : \{\tau\}}$	$\frac{\text{INJECTION} \quad \Gamma \triangleright_s f : \tau_s \rightsquigarrow \tau_c \quad \Gamma \triangleright_s e : \tau_s}{\Gamma \triangleright_c f\%e : \tau_c}$	

$\text{Close}(\tau, \Gamma) = \forall \alpha_0 \dots \alpha_n. \tau$ with $\{\alpha_0, \dots, \alpha_n\} = \text{FreeTypeVar}(\tau) \setminus \text{FreeTypeVar}(\Gamma)$

Fig. 20. ELIOM_e expression typing rules – $\Gamma \triangleright_{\ell} e : \tau$

3.3.2 Converters. To transmit values from the server to the client, we need a serialization format. We assume the existence of a type `serial` in `Constb` which represents the serialization format. The actual format is irrelevant. For instance, one could use JSON or XML.

Converters are special values that describe how to move a value from the server to the client. A converter can be understood as a pair of functions. A converter f of type $\tau_s \rightsquigarrow \tau_c$ is composed of a server-side encoding function of type $\tau_s \rightarrow \text{serial}$, and a client-side decoding function of type $\text{serial} \rightarrow \tau_c$. We assume the existence of two built-in converters:

- The `serial` converter of type $\text{serial} \rightsquigarrow \text{serial}$. Both sides are the identity.
- The `frag` converter of type $\forall \alpha_c. (\{\alpha_c\} \rightsquigarrow \alpha_c)$.

3.3.3 Type universes. It is important to note that there is no identity converter (that would be of type $\forall \alpha. (\alpha \rightsquigarrow \alpha)$). Indeed the client and server type universes are distinct and we cannot translate arbitrary types from one to the other. Some types are only available on one side: database handles, system types, JAVASCRIPT API types. Some types, while available on both sides (because they are in `base` for example), are simply not transferable. For example, functions cannot be serialized in

$$\begin{array}{c}
\text{TYPEVAL} \\
\frac{(\text{type}_{\ell}(\alpha_{\ell_i})t) \in \Gamma \quad \forall i, \Gamma \models_{\ell_i} \tau_i}{\Gamma \models_{\ell}(\tau_i)t} \\
\\
\text{ARROWVAL} \\
\frac{\Gamma \models_{\ell} \tau_1 \quad \Gamma \models_{\ell} \tau_2}{\Gamma \models_{\ell} \tau_1 \rightarrow \tau_2} \\
\\
\text{VARVAL} \\
\frac{}{\Gamma \models_{\ell} \alpha_{\ell}} \\
\\
\text{FRAGVAL} \\
\frac{\Gamma \models_c \tau}{\Gamma \models_s \{\tau\}} \\
\\
\text{QUALIFIEDVAL} \\
\frac{\Gamma \triangleright_{\ell} p : (\text{sig } \mathcal{S}_1; \text{type}_{\ell'}(\alpha_{\ell_i})t; \mathcal{S}_2 \text{ end}) \quad \forall i, \Gamma \models_{\ell_i} \tau_i \quad \ell' > \ell}{\Gamma \models_{\ell}(\tau_i)p.t} \\
\\
\text{CONVVAL} \\
\frac{\Gamma \models_s \tau_1 \quad \Gamma \models_c \tau_2}{\Gamma \models_s \tau_1 \rightsquigarrow \tau_2}
\end{array}$$

Fig. 21. Type validity rules – $\Gamma \models_{\ell} \tau$

$$\begin{array}{c}
\text{REFLEQ} \\
\frac{}{\Gamma \triangleright_{\ell} \tau \approx \tau} \\
\\
\text{TRANSEQ} \\
\frac{\Gamma \triangleright_{\ell} \tau_1 \approx \tau_2 \quad \Gamma \triangleright_{\ell} \tau_2 \approx \tau_3}{\Gamma \triangleright_{\ell} \tau_1 \approx \tau_3} \\
\\
\text{COMMEQ} \\
\frac{\Gamma \triangleright_{\ell} \tau_2 \approx \tau_1}{\Gamma \triangleright_{\ell} \tau_1 \approx \tau_2} \\
\\
\text{FUNEQ} \\
\frac{\Gamma \triangleright_{\ell} \tau_1 \approx \tau'_1 \quad \Gamma \triangleright_{\ell} \tau_2 \approx \tau'_2}{\Gamma \triangleright_{\ell} \tau_1 \rightarrow \tau_2 \approx \tau'_1 \rightarrow \tau'_2} \\
\\
\text{DEFTYPEEQ} \\
\frac{(\text{type}_{\ell'}(\alpha_{\ell_i})t = \tau) \in \Gamma \quad \ell' > \ell}{\Gamma \triangleright_{\ell}(\tau_i)t \approx \tau[\alpha_i \mapsto_{\ell_i} \tau_i]_i} \\
\\
\text{ABSTYPEEQ} \\
\frac{(\text{type}_{\ell'}(\alpha_{\ell_i})t) \in \Gamma \quad \forall i, \Gamma \triangleright_{\ell_i} \tau_i \approx \tau'_i \quad \ell' > \ell}{\Gamma \triangleright_{\ell}(\tau_i)t \approx (\tau'_i)t} \\
\\
\text{QUALDEFTYPEEQ} \\
\frac{\Gamma \triangleright_{\ell} p : (\text{sig } \mathcal{S}_1; \text{type}_{\ell'}(\alpha_{\ell_i})t = \tau; \mathcal{S}_2 \text{ end}) \quad \forall i, \Gamma \triangleright_{\ell_i} \tau_i \approx \tau'_i \quad \ell' > \ell}{\Gamma \triangleright_{\ell}(\tau_i)p.t \approx \tau[n_i \mapsto_{\ell} p.n \mid n_i \in \text{BV}_{\ell}(\mathcal{S}_1)][\alpha_i \mapsto_{\ell_i} \tau_i]_i} \\
\\
\text{QUALABSTYPEEQ} \\
\frac{\Gamma \triangleright_{\ell} p : (\text{sig } \mathcal{S}_1; \text{type}_{\ell'}(\alpha_{\ell_i})t; \mathcal{S}_2 \text{ end}) \quad \forall i, \Gamma \triangleright_{\ell_i} \tau_i \approx \tau'_i \quad \ell' > \ell}{\Gamma \triangleright_{\ell}(\tau_i)p.t \approx (\tau'_i)p.t} \\
\\
\text{FRAGMENTEQ} \\
\frac{\Gamma \triangleright_c \tau \approx \tau'}{\Gamma \triangleright_s \{\tau\} \approx \{\tau'\}} \\
\\
\text{CONVEQ} \\
\frac{\Gamma \triangleright_s \tau_s \approx \tau'_s \quad \Gamma \triangleright_c \tau_c \approx \tau'_c}{\Gamma \triangleright_s \tau_s \rightsquigarrow \tau_c \approx \tau'_s \rightsquigarrow \tau'_c}
\end{array}$$

Fig. 22. Type equivalence rules – $\Gamma \triangleright_{\ell} \tau \approx \tau'$

general. Another example is file handles: they are available both on the server and on the client, but moving a file handle from server to client seems adventurous. Finally, some types may share a semantic meaning, but not their actual representation. This is the case where converters are used, as demonstrated in [Section 2.3.2](#).

3.3.4 Mixed datatypes. The version of ML we consider supports an interesting combination of three features: abstract datatypes, parametrized datatypes and separate compilation at the module level. ELIOM_{ϵ} , as an extension of ML, also supports these features. These three features have non-trivial interactions that need to be accounted for, in particular when introducing properties on types, such as locations.

Let us consider the module shown in [Example 1](#). We declare a server datatype t with two parameters and we hide the definition in the signature. We now want to check that $(t1, t2)t$ is a correct type expressions. However, without the type definition, we don't know if $t1$ and $t2$ are base, client or server types. In order to type check the type sub-expressions, we need more

<pre> 1 module M : sig 2 type%server ('a, 'b) t 3 end = struct 4 type%server ('a, 'b) t = 'a fragment * 'b 5 end </pre> <p style="text-align: center;">(a) Incorrect abstract datatype</p>	<pre> 1 module M : sig 2 type%server ('a[@client], 'b) t 3 end = struct 4 type%server ('a[@client], 'b) t = 'a fragment * 'b 5 end </pre> <p style="text-align: center;">(b) Correct abstract datatype</p>
--	--

Example 1. A module with an abstract datatype.

Figure 1a does not expose information about acceptable sides for 'a and 'b. In **Figure 1b**, annotations specifying the side of type variables are exposed in the interface.

information about the definition of t . The solution, much like variance, is to annotate type variables in datatypes with extra information. This is done in the syntax for type declarations given in **Figure 19**. Each type parameter is annotated with a location. Type variables can only be used on the right location. This ensures proper separation of client and server type variables and their proper usage.

These annotations can be thought of as a simplistic kind system. One could also consider 'a as a constrained type variable, in the style of ML^F [Botlan and Rémy 2003].

3.4 Modules

We now present the module language part of $ELIOM_\epsilon$ as an extension of Leroy [1995] which model the bulk of the OCAML module language. The syntax, presented in **Figure 23**, is composed of the usual module constructs: functors, module constraints, functor application and structures. A structure is composed of a list of components. Similarly, module types are composed of functors and signatures which are a list of signature components. Components can be declaration of values, types or modules. A type in a signature can be declared abstract or not.

The main difference between $ELIOM_\epsilon$ and ML is that structure and signature components are annotated with locations. Value and type declarations can be annotated with a core location ℓ which is either b , s or c . Module declarations can also have one additional possible location: the mixed location m . We use ζ for locations that can be m , b , s or c . Only modules on location m can have subfields on different locations. We also introduce mixed functors, noted $\text{functor}_m(X : \mathcal{M})\mathcal{M}$, which body can contain both client and server declarations. A program is a list of declarations including a client value declaration `return` which is the result of the program.

We first introduce the various features of our module system along with some motivating examples. We then detail how those features are enforced by the typing rules.

3.4.1 Base location and specialization. In **Section 2.1**, we presented an example where a base functor `Map.Make`, is applied to a client module to obtain a new client module. As `Map.Make` is a module provided by the standard library of OCAML, it is defined on location b . In particular, its input signature has components on location b , thus it would seem a module whose components are on the client or the server should not be accepted. We would nevertheless like to create maps of elements that are only available on the client. To do so, we introduce a specialization operation, defined in **Figure 26**, that allows to use a base module in a client or server scope by replacing instances of the base location with the current location.

The situation is quite similar to the application of a function of type $\forall \alpha. \alpha \rightarrow \alpha$ to an argument of type `int`: we need to instantiate the function before being able to use it. Mixed modules only offer a limited version of polymorphism for locations: there is only one "location variable" at a time, and it's always called b . The specialization operation simply rewrites a module signature by substituting all instances of the location b or m by the specified c or s location. Note that before being specialized, a module should be accessible according to the "can be used" relation defined

Module Expressions

$M ::= X_i \mid p.X$ (Variables)
 $\mid (M : \mathcal{M})$ (Type constraint)
 $\mid M_1(M_2)$ (Functor application)
 $\mid \text{functor}(X_i : \mathcal{M})M$ (Functor)
 $\mid \text{functor}_m(X_i : \mathcal{M})M$ (Mixed functor)
 $\mid \text{struct } S \text{ end}$ (Structure)

Structure body

$S ::= \varepsilon \mid D; S$

Structure components

$D ::= \text{let}_\ell x_i = e$ (Values)
 $\mid \text{type}_\ell t_i = \tau$ (Types)
 $\mid \text{module}_\zeta X_i = M$ (Modules)

Programs

$P ::= \text{prog } S \text{ end}$

Module types

$\mathcal{M} ::= \text{sig } S \text{ end}$ (Signature)
 $\mid \text{functor}(X_i : \mathcal{M}_1)\mathcal{M}_2$ (Functor)
 $\mid \text{functor}_m(X_i : \mathcal{M}_1)\mathcal{M}_2$ (Mixed functor)

Signature body

$S ::= \varepsilon \mid \mathcal{D}; \mathcal{S}$

Signature components

$\mathcal{D} ::= \text{val}_\ell x_i : \tau$ (Values)
 $\mid \text{type}_\ell t_i = \tau$ (Types)
 $\mid \text{type}_\ell t_i$ (Abstract types)
 $\mid \text{module}_\zeta X_i : \mathcal{M}$ (Modules)

Environments

$\Gamma ::= \mathcal{S}$

Fig. 23. ELIOM_m's grammar for modules

Figure 17. This means that we never have to specialize a server module on the client (or conversely). Specialization towards location b has no effect since only base modules are accessible on location base. Specialization towards the location m has no effect either: since all locations are allowed inside the mixed location, no specialization is needed. Mixed functors are handled in a specific way, as we see in the next section.

3.4.2 Mixed Functors. Mixed functors are functors that take as input a mixed module and return a mixed module. We note $\text{functor}_m(X_i : \mathcal{M})M$ the mixed functor that takes an argument X_i of type \mathcal{M} and return a module M . They can contain both client and server declarations (or mixed submodules). Mixed functors and regular functors have different types that are not compatible. We saw in [Section 2.4.4](#) an example of usage for mixed functors. Mixed functors have several restrictions compared to regular functors which we now detail using various examples.

Specialization. A naive implementation of specialization of mixed functors would be to specialize on both side of the arrow and apply the resulting functor. Let us see on an example why this solution does not work. In [Example 2](#), the functor F takes as argument a module containing a base declaration and uses it on both sides. If the type of the functor parameter were specialized, the functor application in [Example 2b](#) would be well-typed. However, this makes no sense: $M.y$ is supposed to represent a fragment whose content is the client value of b , but this value doesn't exist since b was declared on the server. There is no value available to inject in the declaration of y' .

The solution here is that specialization on mixed functors should only specialize the return type, not the argument.

Injections. Injections inside client sections (as opposed to escaped values inside client fragments) are fairly static: the value might be dynamic, but the position of the injection and its use sites are statistically known and does not depend on the execution of the program. In particular, injections are independent of the control flow. We can just give a unique identifier to each injection, and use

```

1 module%mixed F (A : sig val b : int end)
2 = struct
3   let%server x = A.b
4   let%server y = [%client A.b]
5 end

```

(a) A mixed functor using a base declaration

```

1 module%server M =
2   F(struct let%server b = 2 end)
3 let%client y' = ~%M.y

```

(b) An ill-typed application of F

Example 2. A mixed functor using base declaration polymorphically

that unique name for lookup on the client. This property comes from the fact that injected server identifiers cannot be bound in a client section.

Unfortunately, this property does not hold in the presence of mixed functor when we assume the language can apply functor at arbitrary positions, which is the case in OCAML. Let us consider [Example 3](#). The functor F takes a structure containing a server declaration x holding an integer and returns a structure containing the same integer, injected in the client. In [Example 3b](#), the functor is used on A or B conditionally. The issue is that the client integer depends both on the server integer and on the local *client* control flow. Lifting the functor application at toplevel would not preserve the semantics of the language, due to side effects. Thus, we avoid this kind of situation by forbidding injections that access dynamic names inside mixed functors.

```

1 module%mixed F
2 (A : sig val%server x : int end)
3 = struct
4   let%client x' = ~%A.x
5 end

```

(a) An problematic mixed functor with an injection

```

1 module%mixed A = struct let%server x = 2 end
2 module%mixed B = struct let%server x = 4 end
3 let%client a =
4   if Random.bool ()
5   then let module M = F(A) in M.x'
6   else let module M = F(B) in M.x'

```

(b) A pathological functor application

Example 3. Problematic example of injection inside a mixed functor

In order to avoid this situation, we add the constraints that injections inside the body of a mixed functors can only refer to outside of the functor. Escaped values, which are injections inside client fragments, are still allowed. The functor presented in [Example 4a](#) is not allowed while the one in [Example 4b](#) is allowed. Formally, this is guaranteed by the MIXEDFUNCTOR rule, where each injection is typechecked in the outer typing environment.

```

1
2 module%mixed F
3 (A:sig val%server x : int end)
4 = struct
5   let%client y = ~%A.x + 2
6 end
7

```

(a) An ill-typed mixed functor using an injection

```

1 let%server x = 3
2 module%mixed F
3 (A:sig val%server y : int end)
4 = struct
5   let%client z = ~%x
6   let%server z' = [%client ~%A.y + 1]
7 end

```

(b) A well-typed mixed functor using an injection

Example 4. Mixed functor and injections

Functor application. Mixed functors can only be applied to mixed structures. This means that in a functor application $F(M)$, M must be a structure defined by a `modulem` declaration. Note that this breaks the property that the current location of an expression or a module can be determined syntactically: The location inside `F(struct ... end)` can be either mixed or not, depending on

F. This could be mitigated by using a different syntax for the application of mixed functor. The justification for this restriction is detailed in [Section 3.6](#).

3.4.3 Type rules. We now review how these various language constructs are reflected in the rules of our type system. As before, the ELIOM module system is built on the ML module system. We extend the typing, validity and subtyping judgments by a location annotation that specifies the location of the current scope. The program typing judgments don't have a location, since a program is always considered mixed. Most rules are fairly straightforward adaptations of the ML rules, annotated with locations.

The typing rules `MODVAR` and `QUALMODVAR` follow the usual rules of ML modules with two modifications: We first check that the module we are looking up can indeed be used on the current location. This is done by the side condition $\zeta' > \zeta$ where ζ is the current location and ζ' is the location where the identifier is defined. This allows, for instance, to use *base* identifiers in a *client* scope. We also specialize the module type of the identifier towards the current location ζ . The specialization operation, which was described in [Section 3.4.1](#), is noted $[M]_{\zeta}$ and is defined in [Figure 26](#).

There are two new typing rules compared to ML: the rules `MIXEDFUNCTOR` and `MIXEDAPPLICATION` define mixed functor definition and application. We use `INJS(·)` which returns the set of all injections in client declarations.

3.4.4 Subtyping and equivalence of modules. Subtyping rules are given in [Figure 25](#). For brevity, we note $\zeta <:(\zeta_1 > \zeta_2)$ as a shorthand for $\zeta <:\zeta_1 \wedge \zeta <:\zeta_2 \wedge \zeta_1 > \zeta_2$, that is, both ζ_1 and ζ_2 are valid locations for components of a module on location ζ and location ζ_1 encompasses location ζ_2 . Note that the following holds:

$$\Gamma \triangleright_{\zeta} \text{struct val}_b t_i : \text{int end} <: \text{struct val}_c t_i : \text{int end}$$

This is perfectly safe, since for any identifier x_i on base, $\text{let}_c x'_j = x_i$ is always valid. This allows programmers to declare some code on base (and get the guarantee that the code is only using usual OCAML constructs) but to expose it as client or server in the module type.

3.5 Separate typechecking

Our `ELIOMe` calculus inherit all the properties of the ML calculus regarding separate compilation. In particular, we show here that it trivially supports separate typechecking, where each module is typechecked independently and only requires knowledge of the module *types* (but not the implementations!) of its dependencies.

Theorem 1 (Separate Typechecking). Given a list of module declarations that form a typed program, there exists an order such that each module can be typechecked with only knowledge of the type of the previous modules.

More formally, given a location ζ and a list of n declarations D_i and a signature \mathcal{S} such that

$$\triangleright_{\zeta}(D_1; \dots; D_n) : \mathcal{S}$$

then there exists n definitions \mathcal{D}_i and a permutation π such that

$$\forall \zeta \forall i < n, \mathcal{D}_1; \dots; \mathcal{D}_i \triangleright_{\zeta} \mathcal{D}_{i+1} : \mathcal{D}_{i+1} \quad \triangleright_{\zeta} \mathcal{D}_{\pi(1)}; \dots; \mathcal{D}_{\pi(n)} <: \mathcal{S}$$

PROOF. Proceed exactly as the ML proof for [Theorem 5](#) in [Appendix A.2.3](#). □

$$\begin{array}{c}
\text{MODVAR} \\
\frac{(\text{module}_{\zeta'} X_i : \mathcal{M}) \in \Gamma \quad \zeta' > \zeta}{\Gamma \triangleright_{\zeta} X_i : \lfloor \mathcal{M} \rfloor_{\zeta}}
\end{array}
\quad
\frac{\text{QUALMODVAR}}{\Gamma \triangleright_{\zeta} p : (\text{sig } \mathcal{S}_1; \text{module}_{\zeta'} X_i : \mathcal{M}; \mathcal{S}_2 \text{ end}) \quad \zeta' > \zeta}
\frac{\Gamma \triangleright_{\zeta} p.X : \lfloor \mathcal{M} [n_i \mapsto_{\zeta'} p.n \mid n_i \in \text{BV}_{\zeta'}(\mathcal{S}_1)] \rfloor_{\zeta}}{\Gamma \triangleright_{\zeta} p.X : \lfloor \mathcal{M} [n_i \mapsto_{\zeta'} p.n \mid n_i \in \text{BV}_{\zeta'}(\mathcal{S}_1)] \rfloor_{\zeta}}$$

$$\frac{\text{STRENGTH}}{\Gamma \triangleright_{\zeta} p : \mathcal{M}}
\frac{\Gamma \triangleright_{\zeta} M : \mathcal{M}' \quad \Gamma \triangleright_{\zeta} \mathcal{M}' <: \mathcal{M}}{\Gamma \triangleright_{\zeta} M : \mathcal{M}}$$

$$\frac{\Gamma \triangleright_{\zeta} M_1 : \text{functor}(X_i : \mathcal{M})\mathcal{M}' \quad \Gamma \triangleright_{\zeta} M_2 : \mathcal{M}}{\Gamma \triangleright_{\zeta} M_1(M_2) : \mathcal{M}' [X_i \mapsto_{\zeta} M_2]}$$

$$\frac{\Gamma \models_{\ell} \mathcal{M} \quad X_i \notin \text{BV}_{\ell}(\Gamma) \quad \Gamma; (\text{module}_{\ell} X_i : \mathcal{M}) \triangleright_{\ell} M : \mathcal{M}'}{\Gamma \triangleright_{\ell} \text{functor}(X_i : \mathcal{M})M : \text{functor}(X_i : \mathcal{M})\mathcal{M}'}
\quad
\frac{\Gamma \models_{\zeta} \mathcal{M} \quad \Gamma \triangleright_{\zeta} M : \mathcal{M}}{\Gamma \triangleright_{\zeta} (M : \mathcal{M}) : \mathcal{M}}$$

$$\frac{\text{MIXEDFUNCTOR}}{\Gamma \models_m \text{sig } \mathcal{S} \text{ end}}
\frac{\Gamma; (\text{module}_m X_i : \text{sig } \mathcal{S} \text{ end}) \triangleright_m M : \mathcal{M}' \quad \forall f_j \% X_j \in \text{INJS}(M), \Gamma \triangleright_c f_j \% X_j : \tau_j}{\Gamma \triangleright_m \text{functor}_m(X_i : \text{sig } \mathcal{S} \text{ end})M : \text{functor}_m(X_i : \text{sig } \mathcal{S} \text{ end})\mathcal{M}'}$$

$$\frac{\text{MIXEDAPPLICATION}}{\Gamma \triangleright_{\zeta} M_1 : \text{functor}_m(X_i : \mathcal{M})\mathcal{M}' \quad \Gamma \triangleright_m M_2 : \mathcal{M} \quad m > \zeta}
\frac{\Gamma \triangleright_{\zeta} M_1(M_2) : \mathcal{M}' [X_i \mapsto_{\zeta} M_2]}{\Gamma \triangleright_{\zeta} M_1(M_2) : \mathcal{M}' [X_i \mapsto_{\zeta} M_2]}$$

$$\frac{\Gamma \triangleright_{\ell} e : \tau \quad x_i \notin \text{BV}_{\ell}(\Gamma) \quad \Gamma; (\text{val}_{\ell} x_i : \text{Close}(\tau, \Gamma)) \triangleright_{\zeta} S : \mathcal{S} \quad \zeta <: \ell}{\Gamma \triangleright_{\zeta} (\text{let}_{\ell} x_i = e; s) : (\text{val}_{\ell} x_i : \tau; S)}$$

$$\frac{\Gamma \models_{\ell} \tau \quad t_i \notin \text{BV}_{\ell}(\Gamma) \quad \Gamma; (\text{type}_{\ell} t_i = \tau) \triangleright_{\zeta} S : \mathcal{S} \quad \zeta <: \ell}{\Gamma \triangleright_{\zeta} (\text{type}_{\ell} t_i = \tau; s) : (\text{type}_{\ell} t_i = \tau; S)}$$

$$\frac{\Gamma \triangleright_{\zeta} M : \mathcal{M} \quad X_i \notin \text{BV}_{\zeta}(\Gamma) \quad \Gamma; (\text{module}_{\zeta} X_i : \mathcal{M}) \triangleright_{\zeta'} S : \mathcal{S} \quad \zeta' <: \zeta \quad \forall \zeta'' \in \text{locations}(\mathcal{M}), \zeta'' > \zeta}{\Gamma \triangleright_{\zeta'} (\text{module}_{\zeta} X_i = M; s) : (\text{module}_{\zeta} X_i : \mathcal{M}; S)}$$

$$\frac{\Gamma \triangleright_{\zeta} S : \mathcal{S}}{\Gamma \triangleright_{\zeta} \text{struct } S \text{ end} : \text{sig } \mathcal{S} \text{ end}}
\quad
\frac{}{\Gamma \triangleright_{\zeta} \varepsilon : \varepsilon}$$

Fig. 24. Module typing rules – $\Gamma \triangleright_{\zeta} m : \mathcal{M}$

3.6 Interpreted semantics

While ELIOM, just like OCAML, is a compiled language, it is desirable to present a semantics that does not involve complex program transformation. The reason is two-fold: First, this simple semantics should be reasonably easy to explain to users. Indeed, this semantics is the one used to present ELIOM in Section 2. However, we must also show that this semantics is correct, in that it does actually corresponds to our compilation scheme. This is done in Section 5.5. As presented in

$$\begin{array}{c}
\text{SUBSTRUCT} \\
\frac{\pi : [1; m] \rightarrow [1; n] \quad \forall i \in [1; m], \Gamma; \mathcal{D}_1; \dots; \mathcal{D}_n \triangleright_{\zeta} \mathcal{D}_{\pi(i)} <: \mathcal{D}'_i}{\Gamma \triangleright_{\zeta} (\text{sig } \mathcal{D}_1; \dots; \mathcal{D}_n \text{ end}) <: (\text{sig } \mathcal{D}'_1; \dots; \mathcal{D}'_m \text{ end})} \\
\\
\frac{\Gamma \triangleright_{\ell_2} \tau_1 \approx \tau_2 \quad \zeta <: (\ell_1 > \ell_2)}{\Gamma \triangleright_{\zeta} (\text{val}_{\ell_1} x_i : \tau_1) <: (\text{val}_{\ell_2} x_i : \tau_2)} \quad \frac{\Gamma \triangleright_{\zeta_2} \mathcal{M}_1 <: \mathcal{M}_2 \quad \zeta <: (\zeta_1 > \zeta_2)}{\Gamma \triangleright_{\zeta} (\text{module}_{\zeta_1} X_i : \mathcal{M}_1) <: (\text{module}_{\zeta_2} X_i = \mathcal{M}_2)} \\
\\
\frac{\Gamma \triangleright_{\ell} \mathcal{M}'_a <: \mathcal{M}_a \quad \Gamma, (\text{module}_{\ell} X : \mathcal{M}'_a) \triangleright_{\ell} \mathcal{M}_r <: \mathcal{M}'_r}{\Gamma \triangleright_{\ell} \text{functor}(X : \mathcal{M}_a) \mathcal{M}_r <: \text{functor}(X : \mathcal{M}'_a) \mathcal{M}'_r} \\
\\
\frac{\Gamma \triangleright_m \mathcal{M}'_a <: \mathcal{M}_a \quad \Gamma, (\text{module}_m X : \mathcal{M}'_a) \triangleright_m \mathcal{M}_r <: \mathcal{M}'_r}{\Gamma \triangleright_m \text{functor}_m(X : \mathcal{M}_a) \mathcal{M}_r <: \text{functor}_m(X : \mathcal{M}'_a) \mathcal{M}'_r} \\
\\
\frac{\Gamma \triangleright_{\ell_2} \tau_1 \approx \tau_2 \quad \zeta <: (\ell_1 > \ell_2)}{\Gamma \triangleright_{\zeta} (\text{type}_{\ell_1} t_i = \tau_1) <: (\text{type}_{\ell_2} t_i = \tau_2)} \quad \frac{\zeta <: (\ell_1 > \ell_2)}{\Gamma \triangleright_{\zeta} (\text{type}_{\ell_1} t_i) <: (\text{type}_{\ell_2} t_i)} \\
\\
\frac{\Gamma \triangleright_{\ell_2} t_i \approx \tau \quad \zeta <: (\ell_1 > \ell_2)}{\Gamma \triangleright_{\zeta} (\text{type}_{\ell_1} t_i) <: (\text{type}_{\ell_2} t_i = \tau)} \quad \frac{\zeta <: (\ell_1 > \ell_2)}{\Gamma \triangleright_{\zeta} (\text{type}_{\ell_1} t_i = \tau_1) <: (\text{type}_{\ell_2} t_i)}
\end{array}$$

Fig. 25. Module subtyping rules – $\Gamma \triangleright_{\zeta} M <: M'$

$$\begin{array}{l}
[\mathcal{M}]_b = \mathcal{M} \qquad [\mathcal{M}]_m = \mathcal{M} \\
[\text{sig } \mathcal{S} \text{ end}]_l = \text{sig } [\mathcal{S}]_l \text{ end} \quad [\text{functor}_m(X_i : \mathcal{M}) \mathcal{M}'_i]_l = \text{functor}_m(X_i : \mathcal{M}) [\mathcal{M}'_i]_l \\
[\varepsilon]_l = \varepsilon \quad [\text{functor}(X_i : \mathcal{M}) \mathcal{M}'_i]_l = \text{functor}(X_i : [\mathcal{M}]_l) [\mathcal{M}'_i]_l \\
\\
[\text{val}_{\ell} x_i : \tau; \mathcal{S}]_l = \begin{cases} \text{val}_l x_i : \tau; [\mathcal{S}]_l & \text{when } \ell > l \\ [\mathcal{S}]_l & \text{otherwise} \end{cases} \\
[\text{type}_{\ell} t_i = \tau; \mathcal{S}]_l = \begin{cases} \text{type}_l t_i = \tau; [\mathcal{S}]_l & \text{when } \ell > l \\ [\mathcal{S}]_l & \text{otherwise} \end{cases} \\
[\text{type}_{\ell} t_i; \mathcal{S}]_l = \begin{cases} \text{type}_l t_i; [\mathcal{S}]_l & \text{when } \ell > l \\ [\mathcal{S}]_l & \text{otherwise} \end{cases} \\
[\text{module}_{\zeta} X_i : \mathcal{M}; \mathcal{S}]_l = \begin{cases} \text{module}_l X_i : [\mathcal{M}]_l; [\mathcal{S}]_l & \text{when } \zeta > l \\ [\mathcal{S}]_l & \text{otherwise} \end{cases}
\end{array}$$

Where l is either c or s .Fig. 26. Module specialization operation – $[\mathcal{M}]_{\zeta}$

$$\begin{aligned}
& \varepsilon/p = \varepsilon \\
& (\text{sig } \mathcal{S} \text{ end})/p = \text{sig } \mathcal{S}/p \text{ end} \\
& (\text{module}_{\zeta} X_i = \mathcal{M}; \mathcal{S})/p = \text{module}_{\zeta} X_i = \mathcal{M}/p; \mathcal{S}/p \\
& (\text{type}_{\ell} t_i = \tau; \mathcal{S})/p = \text{type}_{\ell} t_i = (\alpha^*)p.t; \mathcal{S}/p \\
& (\text{type}_{\ell} t_i; \mathcal{S})/p = \text{type}_{\ell} t_i = (\alpha^*)p.t; \mathcal{S}/p \\
& (\text{val}_{\ell} x_i : \tau; \mathcal{S})/p = \text{val}_{\ell} x_i : \tau; \mathcal{S}/p \\
& (\text{functor}(X_i : \mathcal{M})\mathcal{M}')/p = \text{functor}(X_i : \mathcal{M})(\mathcal{M}'/p(X_i)) \\
& (\text{functor}_m(X_i : \mathcal{M})\mathcal{M}')/p = \text{functor}_m(X_i : \mathcal{M})(\mathcal{M}'/p(X_i))
\end{aligned}$$

Fig. 27. Module strengthening operation – M/p

Section 2, ELIOM execution proceeds in two steps: The server part of the program is executed first. This creates a client program, which is then executed.

Let us first introduce a few notations. Generated client programs are noted μ . Server expressions (resp. declarations) that do not contain injections are noted \bar{e} (resp. \bar{D}). Values are the same as for ML: constants, closures, structures and functor closures. We consider a new class of identifiers called “references” and noted in bold, such as \mathbf{r} or \mathbf{R} . We assume the existence of a name generator that can create arbitrary new fresh \mathbf{r} identifiers at any point of the execution. References are used as global identifiers that ignore scoping rules. References can also be qualified as “reference paths”, noted $\mathbf{X.r}$. This is used for mixed functors, in particular. We use γ to note the global environment where such references are stored.

We now introduce a new reduction relation, \Rightarrow_{ζ} , which is the reduction over ELIOM constructs on side ζ . The notation \Rightarrow_{ζ} actually represents several reduction relations which are presented in [Figures 29, 32 and 33](#). Four of these relations reduce the server part of the code and emit a client program. We note $e \xrightarrow{\rho}_{\iota} v, \mu, \theta$ the reduction of a server expression e inside a context ι in the environment ρ . It returns the value v , the client program μ and emits the trace θ . The context ι can be either base (b), server (s), server code inside client contexts (c/s) or server code inside mixed contexts (m). We also have a client reduction, noted $e \xrightarrow{\rho | \gamma \rightarrow \gamma'}_c v, \theta$ which reduces a client expression e inside an environment ρ , returns a value v and emits a trace θ . It also updates a global environment from γ to γ' .

Note that the first family of relation executes *only* the server part of a program and returns a client program, which is then executed by \Rightarrow_c . This is represented formally by the PROGRAM rule. In order to reduce an ELIOM program P , we first reduce the server part using \Rightarrow_m . This returns no value and a client program μ which we execute. We now look into each specific feature in greater detail

3.6.1 Generated client programs. Let us first describe evaluation rules for generated client programs. Generated client programs are ML programs with some additional constructions which are described in [Figure 28](#). The new evaluation rules are presented in [Figure 29](#). The construction $\text{bind env } \mathbf{f}$ binds the current accessible environment to \mathbf{f} in the global environment γ . This is implemented by the BINDERV rule. $\text{bind } \mathbf{r} = e$ with \mathbf{f} computes e in the environment previously associated to \mathbf{f} . The result is then stored as \mathbf{r} in γ . This construction is also usable for module expressions and is implemented by the BIND and BIND_m rules. All these constructions also accept paths of references such as $\mathbf{R.f}$.

The new bind constructs are similar to the ones used in languages with continuations in the catch/throw style. Instead of storing both an environment and the future computation, we store only the environment. This will allow us to implement closures across locations, in particular the case where fragments are used inside a server closure.

The client reduction relation also inherits the ML rules (rule CLIENTCODE). In such a case, the global environment is passed around following an order compatible with traces. For example, the LETIN rule for let expression would be modified like so:

$$\frac{e' \xrightarrow{\rho | \gamma \rightarrow \gamma'}_c v', \theta \quad e \xrightarrow{\rho + \{x \mapsto v'\} | \gamma' \rightarrow \gamma''}_c v, \theta'}{(\text{let } x = e' \text{ in } e) \xrightarrow{\rho | \gamma \rightarrow \gamma''}_c v, \theta @ \theta'}$$

Here, e' is evaluated first (since θ is present first in the resulting traces), hence it uses the initial environment γ and returns the environment γ' , which is then passed along.

From now on, we use f to denote the reference associated to fragments closures and r to denote the reference associated to a specific value of a fragment.

$$\begin{aligned} \mathbf{p} &::= (\mathbf{X}.)^* \mathbf{x} && \text{(Reference path)} \\ D_c &::= D_{\text{ML}} \\ &| \text{bind env } \mathbf{p} && \text{(Env binding)} \\ &| \text{bind } \mathbf{p} = e \text{ with } \mathbf{p}' \mid \text{bind } \mathbf{p} = M \text{ with } \mathbf{p}' && \text{(Global binding)} \end{aligned}$$

Fig. 28. Grammar of client programs

$$\begin{array}{c} \text{BIND} \\ \frac{e \xrightarrow{\gamma(\mathbf{p}_f) | \gamma \rightarrow \gamma'}_c v, \theta \quad S \xrightarrow{\rho | (\gamma' + \{\mathbf{p} \mapsto v\}) \rightarrow \gamma''}_c V, \theta'}{(\text{bind } \mathbf{p} = e \text{ with } \mathbf{p}_f; S) \xrightarrow{\rho | \gamma \rightarrow \gamma''}_c V, \theta @ \theta'} \\ \\ \text{BIND}_m \\ \frac{M \xrightarrow{\gamma(\mathbf{p}_f) | \gamma \rightarrow \gamma'}_c V, \theta \quad S \xrightarrow{\rho | (\gamma' + \{\mathbf{p} \mapsto V\}) \rightarrow \gamma''}_c V', \theta'}{(\text{bind } \mathbf{p} = M \text{ with } \mathbf{p}_f; S) \xrightarrow{\rho | \gamma \rightarrow \gamma''}_c V', \theta @ \theta'} \end{array} \quad \begin{array}{c} \text{BINDENV} \\ \frac{S \xrightarrow{\rho | (\gamma + \{\mathbf{p}_f \mapsto \rho\}) \rightarrow \gamma'}_c V, \theta}{(\text{bind env } \mathbf{p}_f; S) \xrightarrow{\rho | \gamma \rightarrow \gamma'}_c V, \theta} \\ \\ \text{CLIENTCODE} \\ \text{Inherit the rules} \\ \text{from ML} \end{array}$$

Fig. 29. Semantics for client generated programs – $e \xrightarrow{\rho | \gamma \rightarrow \gamma'}_c v, \theta$

3.6.2 Base, Client and Server declarations. We now consider the case of base, client and server declarations. The rules are presented in [Figure 32](#). Let us first describe the execution of complete ELIOM_e programs (rule PROGRAM). A program P reduces to a client value v if and only if we can first create a *server* reduction of P that produces no value, emits a client program μ and a trace θ_s . We can then create a reduction of μ that reduces in v with a trace θ_c . The trace of the program is the concatenation of the traces. We see that the execution of ELIOM_e program is split in two as described earlier. Let us now look in more details at various construction of the ELIOM_e language.

Base. The base reduction relation corresponds exactly to the ML reduction relation, and always returns empty programs (rule BASECODE). When reducing a base declaration in a mixed context, we both reduce the declaration using \Rightarrow_b , but also add the declaration to the emitted client program (rule BASEDECL). As we can see, base declarations are executed twice: once on the server and once on the client.

Client contexts and injections. The goal of the client reduction relation $\Rightarrow_{c/s}$ is *not* to reduce client programs. It only reduces server code contained by injections inside client code. It returns a client expression without injections, a client program and a trace. Since we don't want to execute client code, it does not inherit the reduction rules for ML. Given an injection $f\%e$, the rule INJECTION reduces the server side expression $(f^s e)$ to a value v . We then transform the server value v into a client value using the \downarrow operator presented Figure 31. We then returns the client expression $(f^c \downarrow v)$ without executing it. This expression will be executed on the client side, to deserialize the value. The value injection operator, noted \downarrow represents the serialization of values from the server to the client and is the identity over constants in $Const_b$ and references, and fail on any other values. According to the definition of converters, if f is a converter $\tau_s \rightsquigarrow \tau_c$, then f^s is the server side function of type $\tau_s \rightarrow \text{serial}$ and v should be of type serial . Since serial is defined on b , the injection of values should be the identity.

The rule CLIENTCONTEXT defines the evaluation of server expression up to client contexts. Client contexts are noted $E[e_1, \dots, e_n]$ and are defined in Figure 30. A client context can have any number of holes which must all contain injections. The rest of the context can contain arbitrary ML syntax that are not injections. Evaluation under a multi-holed context proceed from left to right. The resulting programs and traces are constructed by concatenation of each program and trace.

In order to evaluate client declarations, the rule CLIENTDECL uses $\Rightarrow_{c/s}$ to evaluate the server code present in the declaration D_c which returns a declaration without injections \overline{D}_c and a client program μ . We then return the client program composed by the concatenation of the two. We demonstrate this in Example 5. The ELIOM _{ϵ} program is presented on the left side. It first declares the integer a on the server then inject it on the client and returns the result. The emitted code, shown in the middle, contains an explicit call to the int^c deserializer while the rest of the client code is unchanged. The returned value is shown on the right.

$$\begin{array}{l} \text{let}_s a = 3 \\ \text{let}_c \text{return} = \text{int}\%x + 1 \end{array} \Longrightarrow_m \text{let return} = (\text{int}^c 3) + 1 \Longrightarrow_c 4, \langle \rangle$$

Example 5. Execution of a client declaration

$$\begin{array}{ll} E_e ::= [f\%e] \mid e \mid (E_e E_e) \mid \lambda x. E_e \mid \text{let } x = E_e \text{ in } E_e & \downarrow c = c \quad \text{when } c \in Const_b \\ E_M ::= M \mid (E_M : \mathcal{M}) \mid E_M(E_M) & \downarrow p = p \\ \quad \mid \text{functor}(X_i : \mathcal{M}) E_M \mid \text{struct } (E_D)^* \text{ end} & \downarrow v = \perp \quad \text{otherwise} \\ E_D ::= D \mid \text{let}_c x_i = E_e \mid \text{module}_c X_i = E_M & \end{array}$$

Fig. 30. Execution contexts for injections – $E[\cdot]$

Fig. 31. Injections
of values – $\downarrow v$

Server code and fragments. The server reduction relation reduces server code and emits the appropriate client program associated to client fragments. Since client program are mostly ML programs, it inherits the ML reduction rules (rule SERVERCODE) where client programs are concatenated in the same order as traces. Client fragments are handled by the rule FRAGMENT. Let us consider a fragment $\{\{ e \}\}$, this evaluation proceeds in two steps: first, we evaluate all the injections inside the client expression e using the relation $\Rightarrow_{c/s}$ described in the previous section. We thus obtain an expression without injection \overline{e} and a client program μ .

The second step is to register \overline{e} to be evaluated in the client program. One could propose to simply consider client fragments as values. This is however quite problematic, as it could lead to

duplicated side effects. Consider the program presented on the left side of [Example 6](#). If we were simply to simply pass fragments along, the `print` statement would be evaluated twice. Instead, we create a fresh identifier `r` that will be globally bound to \bar{e} in the client program, as shown in rule `FRAGMENT`. This way, the client expression contained inside the fragment will be executed once, in a timely manner. The execution rule for fragment is demonstrated in [Example 6](#). As before, the ELIOM_e program is presented on the left, the emitted client program is shown in the middle and the returned value is on the right. Note that both frag^s and frag^c are the identity function.

$$\begin{array}{l}
 \text{let}_s x = \{ \{ (\text{print } 3) \} \} \\
 \text{let}_c \text{return} = \\
 \quad \text{frag}^s x + \text{frag}^c x
 \end{array}
 \Longrightarrow_m
 \begin{array}{l}
 \text{bind env } f \\
 \text{bind } r = (\text{print } 3) \text{ with } f \\
 \text{let return} = \\
 \quad (\text{frag}^c r) + (\text{frag}^c r)
 \end{array}
 \Longrightarrow_c 6, \langle 3 \rangle$$

Example 6. Execution of a fragment containing side-effects

Closures and fragments. In the client program above, we also use a reference `f` and the `bind env` construct. To see why this is necessary, we now consider a case where fragments are used inside closures. This is presented in [Example 7](#). The ELIOM_e program, presented on the left, computes $1 + 3 + 2$ on the client (although in a fairly obfuscated way). We first define the client variable `a` as 1. We then define a server closure `f` containing a client fragment capturing `a`. We then define a new variable also named `a` and call $(f \ 3)$, inject the results and returns. When evaluating the definitions of `f`, since it contains syntactically a client fragment, we will emit the client instruction `bind env f`, where `f` is a fresh identifier. This will capture the local environment, which is $\{a \mapsto 1\}$ at this point of the client program. When we execute $(f \ 3)$, we will finally reduce the client fragment and emit the `(bind r = (intc 3) + a with f)` instruction. On the client, this will be executed in the `f` environment, hence `a` is 1 and the result is 4. Once this is executed, we move back to the regular environment, where `a` is 2, and proceed with the execution.

Thanks to this construction, the capturing behavior of closures is preserved across location boundaries. The `bind env` construct is generated by the `SERVERDECL` rule. $\text{FRAGS}(D_s)$ returns the fragments syntactically present in D_s . For each fragment, the local environment is bound to the associated reference.

$$\begin{array}{l}
 \text{let}_c a = 1 \\
 \text{let}_s f x = \{ \{ \text{int}^c x + a \} \} \\
 \text{let}_c a = 2 \\
 \text{let}_s y = (f \ 3) \\
 \text{let}_c \text{return} = \text{frag}^c y + a
 \end{array}
 \Longrightarrow_m
 \begin{array}{l}
 \text{let } a = 1 \\
 \text{bind env } f \\
 \text{let } a = 2 \\
 \text{bind } r = (\text{int}^c 3) + a \text{ with } f \\
 \text{let return} = (\text{frag}^c r) + a
 \end{array}
 \Longrightarrow_c 6, \langle \rangle$$

Example 7. Execution of a fragment inside a closure

Fragment annotations. In the previous examples, we presented the server reduction rules where, for each syntactic fragment, a fresh reference `f` is generated and bound to the environment. In the rest of this thesis, we will simply assume that all fragments syntactically present in the program are annotated with a unique reference. Such annotation is purely syntactic and can be done by walking the syntax tree of the program. Annotated fragments are noted $\{ \{ \dots \} \}_f$.

Mixed structures syntactically present in the program are also annotated in a similar manner with a unique module reference. Annotated mixed structures are noted `struct ... endF`.

3.6.3 *Mixed modules.* Let us now describe the reduction relation for mixed modules. The mixed reduction relation is presented in Figure 33 and, just like the server relation, has for goal to evaluate all the server code and emit a client program to be later evaluated by the client relation. Mixed modules can be composed of either mixed functors, functor applications or structures. The mixed relation contains various rules that are similar to the ML reduction rules for modules. The notable novel aspect of mixed functor is that they both have a client part and a server part. This is different from client fragments, which only have a client part that can be manipulated on the server via an identifier. The server part of mixed modules also need to indicate its client part. In order to do this, each mixed structure will contains an additional field called Dyn which contains a module identifier. The identifier points to a globally bound module on the client which is the result of the client-side evaluation.

Let us first demonstrate these features in Example 8. In this example, we declare a mixed module X containing a fragment x and an integer y . We then declare another mixed module Y containing a submodule. The structure of the emitted client code mimics closely the structure of the server code. In particular, the `bind` operation is nested inside the mixed module X that is emitted on the client. The exact same names are reused on the client. We also register each structure in the global environment using the annotated identifier of the structure. Here, we use the `bind` construct as a shorthand for `bind with` that doesn't change the environment. The shape of the program is kept intact thanks to the `MIXEDMODVAR`, `MIXEDQUALMODVAR` and `MIXEDSTRUCT` rules. The first

$$\begin{array}{c}
 \text{Server code inside client contexts} \\
 \text{INJECTION} \qquad \qquad \qquad \text{CLIENTCONTEXT} \\
 \frac{(f^s e) \xrightarrow{\rho}_s v, \mu, \theta}{f\%e \xrightarrow{\rho}_{c/s} (f^c \downarrow v), \mu, \theta} \qquad \frac{\forall i, e_i \xrightarrow{\rho}_{c/s} v_i, \mu_i, \theta_i}{E[e_1, \dots, e_n] \xrightarrow{\rho}_{c/s} E[v_1, \dots, v_n], \mu_1; \dots; \mu_n, @_i \theta_i} \\
 \text{Server code} \qquad \qquad \qquad \text{Base code} \\
 \text{FRAGMENT} \qquad \qquad \qquad \text{SERVERCODE} \qquad \qquad \qquad \text{BASECODE} \\
 \frac{e \xrightarrow{\rho}_{c/s} \bar{e}, \mu, \theta \quad r \text{ fresh}}{\{\{ e \}\}_f \xrightarrow{\rho}_s r, (\mu; \text{bind } r = \bar{e} \text{ with } f), \theta} \quad \text{Inherit the rules from ML} \qquad \qquad \frac{\text{BASECODE}}{\xrightarrow{\rho} \equiv \xrightarrow{\rho}_b} \\
 \text{Declarations} \\
 \text{BASEDECL} \qquad \qquad \qquad \text{CLIENTDECL} \\
 \frac{D_b \xrightarrow{\rho}_b V, \varepsilon, \theta \quad S \xrightarrow{\rho+V}_m V', \mu', \theta'}{D_b; S \xrightarrow{\rho}_m V + V', (D_b; \mu'), \theta @ \theta'} \qquad \frac{D_c \xrightarrow{\rho}_{c/s} \overline{D}_c, \mu, \theta \quad S \xrightarrow{\rho}_m V, \mu', \theta'}{D_c; S \xrightarrow{\rho}_m V, (\mu; \overline{D}_c; \mu'), \theta} \\
 \text{SERVERDECL} \\
 \frac{\text{FRAGS}(D_s) = \{\{ e_i \}\}_{f_i} \quad D_s \xrightarrow{\rho}_s V, \mu, \theta \quad S \xrightarrow{\rho+V}_m V', \mu', \theta'}{D_s; S \xrightarrow{\rho}_m V + V', (\text{bind env } f_i; \mu; \mu'), \theta @ \theta'} \\
 \text{PROGRAM} \\
 \frac{P \xrightarrow{\rho}_m (), \mu, \theta_s \quad \mu \xrightarrow{\rho | \varepsilon \rightarrow \gamma}_c v, \theta_c}{P \xrightarrow{\rho}_s v, \theta_s @ \theta_c}
 \end{array}$$

Fig. 32. Semantics for base, client and server sections – $e \xrightarrow{\rho}_{c/s} v, \mu, \theta$

two are similar to the non mixed version, but the last one deserves some explanation. First, it prefixes all the fragment references inside the body of the structure. This is for consistency with functors, as we will see later. It then adds the Dyn field to the returned structure, as discussed before. Finally, it emits a bind on the client and returns the module reference. Each structure is thus bound appropriately, even when nested.

Module identifiers are not used in the present program, but they are used in the case of mixed functors, as we will see now.

<pre> module_m X = struct let_s x = {{ 1 }} let_c y = 2 + frag%<i>x</i> end_X module_m Y = struct module_m A = X end_Y let_c return = Y.A.y </pre>	\Longrightarrow_m	<pre> bind X = struct bind env X.f bind r = 1 with X.f let y = 2 + (frag^c r) end module X = X bind Y = struct module A = X end module Y = Y let return = Y.A.y </pre>	\Longrightarrow_c	3, ⟨⟩
--	---------------------	--	---------------------	-------

Example 8. Execution of mixed modules

<pre> let_s x = 1 module_m F(X : \mathcal{M}) = struct let_c b = X.a + int%<i>x</i> end_Y module_m Y = struct let_c a = 2 end_Y module_c Z = F(Y) let_c return = Z.b </pre>	\Longrightarrow_m	<pre> module F(X : \mathcal{M}) = struct let b = X.a + (int^c 1) end_Y bind Y = struct let a = 2 end module Y = Y module Z = F(Y) let return = Z.b </pre>	\Longrightarrow_c	3, ⟨⟩
--	---------------------	--	---------------------	-------

Example 9. Execution of mixed functors with injections

Mixed functors, injections and client side application. Before exposing the complex interaction of mixed functors and fragment, let us illustrate various details about mixed functors in [Example 9](#). The server code proceed in the following way: we first define a server variable x followed by a mixed functor F containing an injection. We then define a mixed module Y and executes *on the client* the functor application $F(Y)$.

First, let us recall that injections inside mixed functors can only refer to elements outside the functor. This means that injections inside functors can be reduced as soon as we consider a functor. In particular, we do not wait for functor application. This can be seen in the MODCLOSURE rule which returns a functor closure on the server side and emit the client part of the functor on the client side. We then take the client part of the body of the functor (noted $M|_c$) and applies the $\Longrightarrow_{c/s}$ reduction relation, which executes injections inside client code. In this example, it results in the injection $\text{int}\%x$ being resolved immediately in the client-side version of the functor.

Mixed functor application can be done in client and server contexts. When it is done in a client context, we simply call the client-side definition and omits the server-side execution completely.


```

modulem F(X : M) = struct
  lets x = {{ X.a + int%X.b }}fx
endF
modulem Y = struct
  letc a = 4
  lets b = 2
endY

modulem Z = F(Y)

letc return = frag%Z.x

bind env F
module F(X : M) = struct end
bind Y = struct
  let a = 4
end
module Y = Y

⇒m bind RZ = struct
  module X = Y
  bind env RZ.fx
  bind rx = Y.a + (intc 2)
  with RZ.fx
end with F
module Z = F(Y)

let return = (fragc rx)

```

Example 10. Execution of mixed functors with fragments

Hence we can simply emit the client-code $F(Y)$. Execution is done through the usual rules for client sections. This is always valid since each mixed declaration emits a client declaration with the same name and the same shape.

Mixed functors and fragments. The difficulty of the reduction of mixed functor containing fragments is that the server-side application of a mixed functor should result in both server and client effects. This makes the reduction rules for mixed functor application quite delicate. We illustrate this with [Example 10](#). In this example, we define a functor F contains only the server declaration x . The argument of the functor simply contains two integers, one on the server and one on the client. In the fragment bound to x , we add the two integers (using an escaped value). The interesting aspect here is that the body of the client fragment depends on both the client and the server side of the argument, even if there is no actual client side for the functor F . The rest of the program is composed of a simple mixed module Y and the mixed functor application $F(Y)$.

The first step of the execution is to define the client side part of F and Y , as demonstrated in the previous example. In this case, since F only contains a server side declaration, the client part of the functor returns an empty structure. We then have to execute $F(Y)$. This is done with the STRUCTBETA rule. When reducing a mixed functor application, we first generate a fresh identifier (R_Z here) and prefix all the fragment closure identifiers. We then evaluate the body of the functor on the server, which gives us both the server module value and the generated client code. In this case, we simply obtain the binding of r_x . Note that this reference is not prefixed by R_Z since it is freshly generated at runtime. If the functor was applied again, we would simply generated a new one. In order for functor arguments to be properly available on the client, we need to introduce additional bindings. For this purpose, we lookup the Dyn field for each module argument and insert the additional binding. In this case, module $X = Y$. This gives us a complete client structure which we can bind to R_Z .

We see here that the body of functors allows to emit client code in a dynamic but controlled way. Generated module references used on the client are remembered on the server using the Dyn field while closure identifiers ensure that the proper environment is used. One problematic aspect

Mixed module expressions

$$\begin{array}{c}
\text{MIXEDSTRUCT} \\
\frac{S[\mathbf{f}_i \mapsto \mathbf{X.f}_i]_i \xRightarrow{\rho}_m V, \mu, \theta \quad V' = V + \{\text{Dyn} \mapsto \mathbf{X}\}}{\text{struct } S \text{ end}_X \xRightarrow{\rho}_m V', \mathbf{X}, \text{bind } \mathbf{X} = \text{struct } \mu \text{ end}, \theta} \\
\\
\text{MIXEDMODVAR} \\
\frac{\rho(X) = V}{X \xRightarrow{\rho}_m V, \mathbf{X}, \varepsilon, \langle \rangle} \\
\\
\text{APP} \\
\frac{M \xRightarrow{\rho}_m V, M_c, \mu, \theta \quad M' \xRightarrow{\rho}_m V', M'_c, \mu', \theta' \quad V(V') \xRightarrow{\rho}_m V'', \mu'', \theta''}{M(M') \xRightarrow{\rho}_m V'', M_c(M'_c), \mu; \mu'; \mu'', \theta @ \theta' @ \theta''} \\
\\
\text{STRUCTBETA} \\
\frac{\mathbf{R} \text{ fresh} \quad V_f = \text{functor}_m(\rho')(X_i : \mathcal{M}_i)_i \text{struct } S \text{ end}_F \\
V_i(\text{Dyn}) = \mathbf{R}_i \quad S[\mathbf{f}_i \mapsto \mathbf{R.f}_i]_i \xRightarrow{\rho' + \{X_i \mapsto V_i\}}_m V, \mu, \theta}{V_f(V_1) \dots (V_n) \xRightarrow{\rho}_m V + \{\text{Dyn} \mapsto \mathbf{R}\}, \left(\begin{array}{l} \text{bind } \mathbf{R} = \text{struct} \\ \text{(module } X_i = \mathbf{R}_i;)_i \\ \mu \\ \text{end with } F \end{array} \right), \theta} \\
\\
\text{NOTSTRUCTBETA} \\
\frac{V = \text{functor}_m(\rho')(X_i : \mathcal{M}_i)_i M \\
M \xRightarrow{\rho' + \{X_i \mapsto V_i\}}_m V_r, \mu, \theta}{V(V_1) \dots (V_n) \xRightarrow{\rho}_m V_r, \mu, \theta,} \\
\\
\text{MIXEDQUALMODVAR} \\
\frac{p \xRightarrow{\rho}_m V, \mu, \theta}{p.X \xRightarrow{\rho}_m V(X), p.X, \mu, \theta} \\
\\
\text{EMPTY} \\
\frac{}{\varepsilon \xRightarrow{\rho}_m \{\}, \varepsilon, \langle \rangle} \\
\\
\text{MODCLOSURE} \\
\frac{M|_c \xRightarrow{\rho}_{c/s} \overline{M}, \mu, \theta}{\text{functor}_m(X : \mathcal{M})M \xRightarrow{\rho}_m \text{functor}_m(\rho)(X : \mathcal{M})M, \text{functor}(X : \mathcal{M}|_c)\overline{M}, \mu, \theta} \\
\\
\text{Mixed declarations} \\
\text{MIXEDMODDECL} \\
\frac{\text{MixedStructIds}(M) = \overline{\mathbf{F}}_i \quad M \xRightarrow{\rho}_m V, M^c, \mu, \theta \quad S \xRightarrow{\rho + \{X \mapsto V\}}_m V', \mu', \theta'}{\text{module}_m X = M; S \xRightarrow{\rho}_m \{X \mapsto V\} + V', (\overline{\text{bind env } \mathbf{F}_i; \mu; \text{module } X = M^c; \mu'}, \theta @ \theta')} \\
\\
\text{Fig. 33. Semantics for mixed modules - } M \xRightarrow{\rho}_\zeta V, \mu, \theta
\end{array}$$

of this method is that it leads to two executions of the client side. A partial work-around is to keep track of loaded client module to avoid duplicating side-effects.

4 INTEGRATION WITH OCAML

We now formalize more precisely the interaction between ELIOM_ε , its location system, and the vanilla ML language. In particular, [Theorem 2](#) shows that ML modules can be completely embedded in ELIOM_ε programs without changes.

Let us note $\mathcal{M}_{[\ell \mapsto \ell']}$ the substitution on locations in an ELIOM module type \mathcal{M} . Given an ML module M , we note $M_{[\text{ML} \mapsto \ell]}$ the ELIOM module where all the module components have been annotated with location ℓ . Given an ELIOM module M , we note $M_{[\ell \mapsto \text{ML}]}$ the ML module where all the location have been erased. We extend these notations to module types and environments.

4.1 Properties of specialization

Let us first clarify the behavior of specialization with regard to mono-located ELIOM_ℓ module types.

Proposition 1. Given an ELIOM module type \mathcal{M} and a location $\ell \in \{b, c, s\}$, if $\Gamma \models_\ell \mathcal{M}$, then $\lfloor \mathcal{M} \rfloor_\ell = \mathcal{M}$.

PROOF. By definition of $\langle \cdot \rangle$, \mathcal{M} can only contain declarations on ℓ . This means that, by reflexivity of $\langle \cdot \rangle$, only specialization rules that leave the declaration unchanged are involved. \square

Proposition 2. Given an ELIOM module type \mathcal{M} and a location $\ell \in \{c, s\}$, if $\Gamma \models_b \mathcal{M}$, then $\lfloor \mathcal{M} \rfloor_\ell = \mathcal{M}_{[b \mapsto \ell]}$.

PROOF. We remark that for all $\ell \in \{c, s\}$, $b \succ \ell$. Additionally, mixed functors cannot appear on base (since $m \neq b$). We can then proceed by induction over the rules for specialization. \square

4.2 Interaction between ML and ELIOM_ℓ

We can now relate ML code to equivalent ELIOM_ℓ code that has been annotated with a single location.

Proposition 3. Given ML type τ , expression e , module M and module type \mathcal{M} and locations ℓ, ℓ' :

$$\begin{aligned} \Gamma \models_{\text{ML}} \tau &\implies \Gamma_{[\text{ML} \mapsto \ell']} \models_\ell \tau && \text{Where } \ell' \succ \ell \\ \Gamma \triangleright_{\text{ML}} e : \tau &\implies \Gamma_{[\text{ML} \mapsto \ell']} \triangleright_\ell e : \tau && \text{Where } \ell' \succ \ell \\ \Gamma \models_{\text{ML}} \mathcal{M} &\implies \Gamma_{[\text{ML} \mapsto \ell']} \models_\ell \mathcal{M}_{[\text{ML} \mapsto \ell]} && \text{Where } \ell' \succ \ell \\ \Gamma \blacktriangleright_{\text{ML}} M : \mathcal{M} &\implies \Gamma_{[\text{ML} \mapsto \ell']} \blacktriangleright_\ell M_{[\text{ML} \mapsto \ell]} : \mathcal{M}_{[\text{ML} \mapsto \ell]} && \text{Where } \ell' \succ \ell \end{aligned}$$

PROOF. We remark that each syntax, typing rule or well formedness rule for ML has a direct equivalent rule in ELIOM . We can then simply rewrite the proof tree of the hypothesis to use the ELIOM type and well-formedness rules. We consider only some specific cases:

- By **Proposition 1** and since the modules are of uniform location, the specialization operation in VAR and MODVAR are the identity.
- The side conditions $\ell' \prec \ell$ are always respected since the modules are of uniform location and by reflexivity of $\langle \cdot \rangle$.
- The side conditions $\ell' \succ \ell$ are respected by hypothesis. \square

Proposition 4. Given ML type τ , expression e , module m and module type \mathcal{M} :

$$\begin{aligned} \Gamma \models_b \tau &\implies \Gamma_{[\mapsto \text{ML}]} \models_{\text{ML}} \tau && \Gamma \models_b \mathcal{M} \implies \Gamma_{[\mapsto \text{ML}]} \models_{\text{ML}} \mathcal{M}_{[\mapsto \text{ML}]} \\ \Gamma \triangleright_b e : \tau &\implies \Gamma_{[\mapsto \text{ML}]} \triangleright_{\text{ML}} e : \tau && \Gamma \blacktriangleright_b M : \mathcal{M} \implies \Gamma_{[\mapsto \text{ML}]} \blacktriangleright_{\text{ML}} M_{[\mapsto \text{ML}]} : \mathcal{M}_{[\mapsto \text{ML}]} \end{aligned}$$

PROOF. We first remark that the following features are forbidden in the base part of the language: injections, fragments, mixed functors and any other location than base. The rest of the language contains no tierless features and coincides with ML. We can then proceed by induction over the proof trees. \square

Proposition 5. Given an ML module M (resp. expression e), an execution environment ρ , a location ℓ , a value V (resp. v) and a trace θ :

$$\begin{aligned} M &\xRightarrow{\rho} V, \theta \iff M_{[\text{ML} \mapsto \ell]} \xRightarrow{\rho_{[\text{ML} \mapsto \ell]}}_\ell V_{[\text{ML} \mapsto \ell]}, \varepsilon, \theta \\ e &\xRightarrow{\rho} v, \theta \iff e_{[\text{ML} \mapsto \ell]} \xRightarrow{\rho_{[\text{ML} \mapsto \ell]}}_\ell v_{[\text{ML} \mapsto \ell]}, \varepsilon, \theta \end{aligned}$$

Furthermore, given an ML program P , an execution environment ρ , a value v and a trace θ :

$$\begin{aligned} P \xRightarrow{\rho} v, \theta &\iff P_{[\text{ML} \mapsto i]} \xRightarrow{\rho_{[\text{ML} \mapsto i]}} v_{[\text{ML} \mapsto i]}, \theta && i \in \{c, s\} \\ P \xRightarrow{\rho} v, \theta &\iff P_{[\text{ML} \mapsto b]} \xRightarrow{\rho_{[\text{ML} \mapsto b]}} v_{[\text{ML} \mapsto b]}, \theta @ \theta \end{aligned}$$

PROOF. Let us first note that the ML reduction relation is included in the base, the server and the client-only relations. Additionally, the considered programs, modules or expressions can not contain fragments, injections or binds. The additional rules in the server and client-only relations are only used for these additional syntactic constructs. For the first three statements, we can then proceed by induction. For the last statement, we remark that base code is completely copied to the client during server execution. Using rule PROGRAM, we execute the program twice, which returns the same value but duplicates the trace. \square

Theorem 2 (Base/ML correspondance). ELIOM modules, expressions and types on base location b correspond exactly to the ML language.

PROOF. By Propositions 3, 4 and 5. \square

Thanks to Theorem 2, we can completely identify the language ML and the part of ELIOM on base location. This is of course by design: the base location allows us to reason about the host language, OCAML, inside the new language ELIOM. It also provides the guarantee that anything written in the base location does not contain any communication between client and server. In the rest of the thesis, we omit location substitutions of the form $[\text{ML} \mapsto b]$ and $[b \mapsto \text{ML}]$.

Proposition 3 also has practical consequences: Given a file previously typechecked by an ML typechecker, we can directly use the module types either on base, but also on the client or on the server, by simply annotating all the signature components. This give us the possibility, in the implementation, to completely reuse compiled objects from the OCAML typechecker and load them on an arbitrary location. In particular, it guarantees that we can reuse pure OCAML libraries safely and freely.

5 COMPILATION OF ELIOM PROGRAMS

In Section 3, we gave a tour of the ELIOM _{ϵ} language from a formal perspective, providing a type system and an interpreted semantics. ELIOM, however, is not an interpreted language. The interpreted semantics is here both as a formal tool and to make the semantics of the language more approachable to users, as demonstrated in Section 2. In the implementation, ELIOM programs are compiled to two programs: one server program (which is linked and executed on the server) and a client program (which is compiled to JAVASCRIPT and executed in a browser). The resulting programs are efficient and avoid unnecessary back-and-forth communications between the client and the server.

Description of the complete compilation toolchain, including emission of JAVASCRIPT code, is out of scope of this article (see Vouillon and Balat [2014]). Instead, we describe the compilation process in term of emission of client and server programs in an ML-like language equipped with additional primitives. Hence, we present the typing and execution of compiled programs, in Section 5.1 and the compilation process, in Section 5.2. We then show that our compilation process preserves both the typing and the semantics in Section 5.5.

5.1 Target languages ML_s and ML_c

We introduce the two target languages ML_c and ML_s as extensions of ML. The additions in these two new languages are highlighted in Figure 34. Typing is provided in Section 5.1.5. The semantics

is provided in [Section 5.1.6](#). As before, we use globally bound identifiers, which we call “references” and note in bold: \mathbf{r} . References can also be paths, such as $\mathbf{X.r}$. In some contexts, we accept a special form of reference path, noted Dyn.x which we explain in [Section 5.1.4](#). In practice, these references are implemented with uniquely generated names and associative tables. Contrary to the interpreted semantics, references are also used to transfer values from the server to the client and can appear in expressions. A reference used inside an expression is always of type `serial`.

ML_s grammar		ML_c grammar	
$\mathbf{p} ::= \text{Dyn.f} \mid (\mathbf{X.})^* \mathbf{f}$	(Reference path)	$\mathbf{p} ::= \text{Dyn.f} \mid (\mathbf{X.})^* \mathbf{f}$	(Reference path)
$\tau ::= \dots \mid \text{frag}$	(Fragment type)	$e ::= \dots \mid \mathbf{x}$	(Reference)
$\mathbf{v} ::= \dots \mid \mathbf{r}$	(Reference)	$M ::= \dots \mid \mathbf{X}$	(Module reference)
$e ::= \dots$		$D ::= \dots$	
$\mid \text{fragment } \mathbf{p} \ e^*$	(Fragment call)	$\mid \text{bind } \mathbf{p} = e$	(Fragment closure)
$M ::= \dots$		$\mid \text{bind}_m \mathbf{p} = M$	(Functor fragment)
$\mid \mathbf{p}.\text{Dyn}$	(Dynamic field)	$\mid \text{exec } ()$	(Fragment execution)
$\mid \text{fragment}_m \mathbf{p} (M^*)$	(Fragment)		
$D ::= \dots$			
$\mid \text{injection } \mathbf{x} \ e$	(Injection)		
$\mid \text{end } ()$	(End token)		

Fig. 34. Grammar for ML_s and ML_c as extensions of ML_e

5.1.1 Converters. For each converter f , we note f^s and f^c the server side encoding function and the client side decoding function. If f is of type $\tau_s \rightsquigarrow \tau_c$, then f^s is of type $\tau_s \rightarrow \text{serial}$ and f^c is of type $\text{serial} \rightarrow \tau_c$. We will generally assume that if the converter f is available in the environment, then f^c and f^s are available in the client and server environment respectively.

5.1.2 Injections. For injections, we associate server-side the injected value e to a reference \mathbf{v} using the construction `injection $\mathbf{v} \ e$` , where e is of type `serial`. When the server execution is over, a mapping from references to injected values is sent to the client. \mathbf{v} is then used client-side to access the value.

An example is given in [Example 11](#). In this example, two integers are sent from the server to the client and add them on the client. We suppose the existence of a base abstract type `int`, a converter `int` of type $(\text{int} \rightsquigarrow \text{int})$ and the associated encoding and decoding functions. The server program, in [Example 11a](#), creates two injections, \mathbf{v}_1 and \mathbf{v}_2 and does not expose any bindings nor return any values. These injections hold the serialized integers 4 and 2. The client program, in [Example 11b](#), uses these two injections, deserialize their values, adds them, and returns the result. Note that injection is *not a network operation*. It simply stores a mapping between references (i.e., names) and serialized values. The mapping generated at the end of the server execution is shown in [Example 11c](#). After the execution of the server code, this mapping is sent to the client, and used to execute the client code.

5.1.3 Fragments. The primitive related to fragments also relies on shared references between the server program and the client program. However, these references allow to uniquely identify

<pre>injection v₁ (int^s 4); injection v₂ (int^s 2);</pre>	<pre>let return = (int^c v₁) + (int^c v₂);</pre>	<pre>v₁ ↦ 4 v₂ ↦ 2</pre>
(a) Server program	(b) Client program	(c) Mapping of injections

Example 11. Client-server programs calling and using injections

functions that are defined on the client but are called on the server. To implement this, we use the following primitives:

- In ML_c structures, `bind p = e` declares a new client function bound to the reference `p`. The function `e` takes an arbitrary amount of argument of type `serial` and returns any type.
- In ML_s expressions, `fragment p e1 ... en` is a delayed function application which registers that, on the client, the function associated to `p` will be applied to the arguments `ei`. All the arguments must be of type `serial`. It returns a value of type `frag`, which holds a unique identifier referring to the result of this application.

Here again, none of these primitives are network communication primitives. While the API is similar to Remote Procedure Calls, the execution is very different: `fragment` only accumulates the function call in a list, to be executed later. When the server execution is over, the list of calls is sent to the client, and used during the client execution. OCAML, and consequently ELIOM, are impure languages: the order of execution of statement is important. In order to control the order of execution, we introduce two additional statements: `end ()`, on the server, introduces an end marker in the list of calls. `exec ()`, on the client, executes all the calls until the next end token.

Example 12 presents a pair of programs which emit the client trace $\langle 2; 3; 3 \rangle$, but in such a way that, while the client does the printing, the values and the execution order are completely determined by the server. The server code (**Example 12a**) calls `f` with 2 as argument, injects the result and then calls `f` with 3 as argument. The client code, in **Example 12b**, declares a fragment closure `f`, which simply adds one to its arguments, and `exec` both fragments. In-between both executions, it prints the content of the injection `v`. During the execution of the server, the list of calls (**Example 12c**) and the mapping of injections (**Example 12d**) are built. First, when fragment `f (ints 2)` is executed, a fresh reference `r1` is generated, the call to the fragment is added to the list and `r1` is returned. The injection adds the association `v1 ↦ r1` to the mapping of injections. The call to `end ()` then adds the token `end` to the list of fragments. The second fragment proceeds similarly to the first, with a fresh identifiers `r2`. Once server execution is over, the newly generated list of fragments and mapping of injections are sent to the client. During the client execution, the execution of the list is controlled by the `exec` calls. First, `(f 2)` emits $\langle 2 \rangle$ and is evaluated to 3, and the mapping `r1 ↦ 3` is added to a global environment. Then `v1` is resolved to `r1` and printed (which shows $\langle 3 \rangle$). Finally `(f 3)` emits $\langle 3 \rangle$ and is evaluated to 4.

The important thing to note here is that both the injection mapping and the list of fragments are completely dynamic. We could add complicated control flow to the server program that would drive the client execution according to some dynamic values. The only static elements are the names `f` and `v1`, the behavior of `f` and the number of call to `exec ()`. We cannot, however, make the server-side control flow depend on client-side values, since we obey a strict phase separation between server and client execution.

Finally, remark that we do not need the `bind env` construct introduced in [Section 3.6.1](#). Instead, we directly capture the environment using closures that are extracted in advance. We will see how this extraction works in more details while studying the compilation scheme, in [Section 5.2](#).

<pre>let x₁ = fragment f (int^s 2); injection v₁ (frag^s x₁); end (); let x₂ = fragment f (int^s 3); end ();</pre>	<pre>bind f = λx.((print (int^c x)) + 1); exec (); let a = (print (frag^c v₁)); exec (); let return = a</pre>	<pre>{r₁ ↦ (f 2)}; end; {r₂ ↦ (f 3)}; end; v₁ ↦ r₁</pre>
(a) Server program	(b) Client program	(c) List of fragments (d) Mapping of injections

Example 12. Client-server program defining and calling fragments

5.1.4 Modules. We introduce three new module-related construction that are quite similar to fragment primitives:

- $\text{bind}_m \mathbf{p} = M$ is equivalent to bind for modules. It is a client instruction that associates the module or functor M to the reference \mathbf{p} .
- $\text{fragment}_m \mathbf{p} (\mathbf{R}_1) \dots (\mathbf{R}_n)$ is analogous to $\text{fragment } \mathbf{p} e$ for modules. It is a delayed functor application that is used on the server to register that the functor associated to \mathbf{p} will have to be applied to the modules associated to \mathbf{R}_i . It returns a fresh reference that represents the resulting module. Contrary to fragment , it can only be applied to module references.
- $\mathbf{p}.\text{Dyn}$ returns a reference that represents the client part of a server module \mathbf{p} . This is used for ELIOM_ϵ mixed structure that have both a server and a client part.

The first argument of fragment , fragment_m , bind and bind_m can also be a reference path $\text{Dyn}.f$, where Dyn is the locally bound Dyn field inside a module. This allows us to isolate some bound references inside a fresh module reference. This is useful for functors, as we will now demonstrate in [Example 13](#).

In this example, we again add integers² on the client while controlling the values and the control flow on the server. We want to define server modules that contain server values but also trigger some evaluation on the client, in a similar way to fragments. The first step is to define a module X on the server and to bind a corresponding module \mathbf{X} on the client. Similarly to the interpreted semantics presented in [Section 3.6](#), we add a Dyn field to the server module that points to the client module. Plain structures such as X are fairly straightforward, as we only need to declare each part statically and add the needed reference. bind_m allows to declare modules globally.

We then declare the functor F on the server and bind the functor \mathbf{F} on the client. The server-side functor contains a call to a fragment defined in the client-side functor. The difficulty here is that we should take care of differentiating between fragment closures produced by different functor applications. For this purpose, we use a similar technique than the one presented in [Section 3.6.3](#), which is to prefix the fragment closure identifier \mathbf{f} with the reference of the client-side module. This reference is available on the server side as the Dyn field and is generated by a call to fragment_m . When F is applied to X on the server, we generate a fresh reference \mathbf{R} and add $\{\mathbf{R}_1 \mapsto \mathbf{F} X_0\}$ to the execution queue. When $\text{exec}()$ is called, we introduce the additional binding $\{\text{Dyn} \mapsto \mathbf{R}_1\}$ in the environment and apply \mathbf{F} to X_0 , which will register the $\mathbf{R}_1.f$ fragment closure. Since it is the result of this specific functor application, the closure $\mathbf{R}_1.f$ will always add 4 to its argument. The rest of the execution proceed as shown in the previous section: we call a new fragment, which triggers the client-side addition $2 + 4$ and use an injection to pass the results around.

5.1.5 Type system rules. The ML_s and ML_c typing rules are presented in [Figures 35](#) and [36](#) as a small extension over the ML typing rules presented in [Appendix A.2](#). Note that the typing rules for the new primitives are weakly typed and are certainly not sound with respect to serialization and

²But better! or at very least, more obfuscated.

$$\begin{array}{c}
\text{FRAGMENT} \\
\frac{\forall i, \Gamma \triangleright_{\text{ML}_s} e_i : \text{serial}}{\Gamma \triangleright_{\text{ML}_s} \text{fragment } \mathbf{p} \ e_1 \ \dots \ e_n : \text{frag}} \\
\\
\text{INJECTION} \\
\frac{\Gamma \triangleright_{\text{ML}_s} e : \text{serial}}{\Gamma \triangleright_{\text{ML}_s} \text{injection } \mathbf{v} \ e : \varepsilon} \\
\\
\text{END} \\
\frac{}{\Gamma \triangleright_{\text{ML}_s} \text{end } () : \varepsilon} \\
\\
\text{FRAGMENT}_m \\
\frac{\forall i, \Gamma \triangleright_{\text{ML}_s} p_i : \mathcal{M}_i}{\Gamma \triangleright_{\text{ML}_s} \text{module Dyn} = \text{fragment}_m \ \mathbf{p} \ p_1.\text{Dyn} \ \dots \ p_n.\text{Dyn} : \varepsilon}
\end{array}$$

Fig. 35. Typing rules for ML_s

$$\begin{array}{c}
\text{BIND} \\
\frac{\Gamma \triangleright_{\text{ML}_c} e : \tau}{\Gamma \triangleright_{\text{ML}_c} \text{bind } \mathbf{p} = e : \varepsilon} \\
\\
\text{BIND}_m \\
\frac{\Gamma \triangleright_{\text{ML}_c} M : \mathcal{M}}{\Gamma \triangleright_{\text{ML}_c} \text{bind}_m \ \mathbf{p} = M : \varepsilon} \\
\\
\text{REFERENCE} \\
\frac{}{\Gamma \triangleright_{\text{ML}_c} \mathbf{x} : \text{serial}} \\
\\
\text{EXEC} \\
\frac{}{\Gamma \triangleright_{\text{ML}_c} \text{exec } () : \varepsilon}
\end{array}$$

Fig. 36. Typing rules for ML_c

deserialization. Given arbitrary ML_s and ML_c programs, there is no guarantee that (de)serialization will not fail at runtime. This is on purpose. Indeed, all these guarantees are provided by ELIOM itself. ML_s and ML_c are target languages that are very liberal by design, so that all patterns permitted by ELIOM are expressible with them. Furthermore, from an implementation perspective, ML_s and ML_c are simply OCAML libraries and do not rely on further compiler support. Note that Dyn fields are not reflected in signatures. The fragment FRAGMENT_m rule does not enforce that the Dyn field is present in all the arguments. This is enforced by construction during compilation.

5.1.6 Semantics rules. We define two reduction relations as extensions of the ML reduction rules (see [Appendix A.3](#)). The $\Rightarrow_{\text{ML}_s}$ reduction for ML_s server programs is presented in [Figure 37](#). The $\Rightarrow_{\text{ML}_c}$ reduction for ML_c client programs is presented in [Figure 38](#). Let us consider a server structure S_s and a client structure S_c . A paired execution of the two structures is presented below:

$$S_s \xRightarrow{\rho_s}_{\text{ML}_s} V_s, \xi, \zeta, \theta_s \qquad S_c, \xi \xRightarrow{\rho_c \mid \gamma \cup \zeta \rightarrow \gamma'}_{\text{ML}_c} V_c, \xi', \theta_c$$

Let us now detail these executions rules. As with the ML reduction, ρ_s and ρ_c are the local environments of values while θ_s and θ_c are the traces for server and client executions respectively. V_s and

<pre> module X = struct module Dyn = X0; let a = 2 end; module F(Y : M_s) = struct module Dyn = fragment_m F (Y.Dyn); let b = fragment Dyn.f (int^s Y.a); end; module Z = F(X); end (); injection v1 (frag^s Z.b); (a) Server program </pre>	<pre> bind_m X0 = struct let c = 4 end; bind_m F(Y : M_c) = struct bind Dyn.f = lambda.((int^c a) + Y.c); end; exec () let return = (frag^c v1); (b) Client program </pre>	<pre> {R1 ↦ F X0}; {r2 ↦ R1.f 2}; end (c) List of fragments v1 ↦ r2 (d) Mapping of injections </pre>
--	--	---

Example 13. Client-server program using module fragments

V_c are the returned values. Similarly to the interpreted semantics for ELIOM_ϵ , the client reduction uses global environments noted γ .

As we saw in the previous examples, server executions emits two sets of information during execution: a queue of fragments and a map of injections. Mapping of injection is a traditional (global) environment where bindings are noted $\{v \mapsto \dots\}$. The queue of fragments is noted ξ and contains end tokens end and fragment calls $\{r \mapsto f \ v_1 \dots v_n\}$. Concatenation of fragment queues is noted $\#$. We now see the various rules in more details.

Injections. Injection bindings are collected on the server through the INJECTION rule. When creating a new injection binding, we inject the server-side value using the injection of value operator, noted $\Downarrow v$ and presented in [Figure 31](#). This models the serialization of values before transmission from server to client by ensuring that only base values and references are injected. Other kinds of values should be handled using converters explicitly.

The injection environment ζ forms a valid client-side global environment. When executing the client-side program, we simply assume that ζ is included in the initial global environment γ .

Fragments and functors. On the server, fragments and functors calls are added to the queue through the FRAGMENT and FRAGMENT_m rules. In both rules, the reference of the associated closure or functor is provided, along with a list of arguments. A fresh reference symbolizing the fragment is generated and the call is added to the queue ξ . Note that in the case of regular fragments, the arguments are expressions which can themselves contains fragment calls. The module rule FRAGMENT_m is similar, the main difference being that it only accept module references as arguments of the call.

Fragment closures and functors are bound on the client through the BIND and BIND_m rules, which simply binds a reference to a value or a module value in the global environment γ . Since bind accepts references of the form Dyn.f , it must first resolves Dyn to the actual reference. This is done through the DYN rule.

Segmented execution. In ML_s and ML_c programs, the execution of fragments is segmented through the use of the `exec ()/end ()` instructions. On the server, end instructions are handled through the END rule, which simply adds an end token to the execution queue ξ . On the client, we use the EXEC rule, associated to the FRAG and FRAG_m rules. When `exec` is called, the EXEC rule triggers the execution of the segment of the queue until the next end token. Each token $\{r \mapsto f \ v_1 \dots v_n\}$ is executed with the FRAG rule as the function call $(f \ v_1 \dots v_n)$. The result of this function call is bound to r in the global environment γ . Similarly, functor calls are executed using the FRAG_m rule. Note that for functors we also introduce the Dyn field in the local environment, which allows local bind definitions. Once all the tokens have been executed in the considered fragment queue, we resume the usual execution.

5.2 Slicing

In [Section 5.1](#), we presented the two target languages ML_s and ML_c . We now present the compilation process transforming *one* ELIOM_ϵ program into two distinct ML_s and ML_c programs. Before giving a more formal description in [Section 5.3](#), we present the compilation process through three examples of increasing complexity.

Injections and fragments. [Example 14](#) presents an ELIOM_ϵ program containing only simple declarations involving fragments and injections without modules. The ELIOM_ϵ program is presented on the left, while the compiled ML_s and ML_c programs are presented on the right. In this example, a first fragment is created. It only contains an integer and is bound to a . A second fragment that

$$\begin{array}{c}
\text{INJECTION} \\
\frac{e \xRightarrow{\rho}_{\text{ML}_s} v, \xi, \zeta, \theta \quad S \xRightarrow{\rho}_{\text{ML}_s} V, \xi', \zeta', \theta'}{\text{injection } \mathbf{x} \ e; S \xRightarrow{\rho}_{\text{ML}_s} V, \xi + \xi', (\zeta \cup \{\mathbf{x} \mapsto \downarrow v\} \cup \zeta'), \theta @ \theta'} \\
\\
\text{FRAGMENT} \\
\frac{\mathbf{p} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{p}' \quad \forall i, e_i \xRightarrow{\rho}_{\text{ML}_s} v_i, \xi_i, \zeta_i, \theta_i \quad \mathbf{r} \text{ fresh}}{\text{fragment } \mathbf{p} \ e_1 \dots e_n \xRightarrow{\rho}_{\text{ML}_s} \mathbf{r}, (\xi_1 + \dots + \xi_n + \{\mathbf{r} \mapsto \mathbf{p}' \downarrow v_1 \dots \downarrow v_n\}), \cup_i \zeta_i, @_i \theta_i} \\
\\
\text{FRAGMENT}_m \\
\frac{\mathbf{p} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{p}' \quad \forall i, p_i.\text{Dyn} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{R}_i, \varepsilon, \varepsilon, \langle \rangle \quad \mathbf{R} \text{ fresh}}{\text{fragment}_m \mathbf{p} \ p_1.\text{Dyn} \dots p_n.\text{Dyn} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{R}, \{\mathbf{R} \mapsto \mathbf{p}' \mathbf{R}_1 \dots \mathbf{R}_n\}, \{\}, \langle \rangle} \\
\\
\text{END} \qquad \qquad \qquad \text{DYN} \qquad \qquad \qquad \text{SERVERCODE} \\
\frac{S \xRightarrow{\rho}_{\text{ML}_s} V, \xi, \zeta, \theta}{\text{end } (); S \xRightarrow{\rho}_{\text{ML}_s} V, \text{end} + \xi, \zeta, \theta} \qquad \frac{\rho(\text{Dyn}) = \mathbf{R}}{\text{Dyn.f} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{R.r}} \qquad \begin{array}{l} \text{Inherit the rules} \\ \text{from ML} \end{array}
\end{array}$$

Fig. 37. Semantics rules for $\text{ML}_s - S \xRightarrow{\rho}_{\text{ML}_s} V, \xi, \zeta, \theta$

$$\begin{array}{c}
\text{BIND} \\
\frac{\mathbf{p} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{p}' \quad e, \xi \xRightarrow{\rho | \gamma \rightarrow \gamma'}_{\text{ML}_c} v, \xi', \theta \quad S, \xi' \xRightarrow{\rho | \gamma' + \{\mathbf{p}' \mapsto v\} \rightarrow \gamma''}_{\text{ML}_c} V, \xi'', \theta'}{(\text{bind } \mathbf{p} = e; S), \xi \xRightarrow{\rho | \gamma \rightarrow \gamma''}_{\text{ML}_c} V, \xi'', \theta @ \theta'} \\
\\
\text{BIND}_m \\
\frac{\mathbf{p} \xRightarrow{\rho}_{\text{ML}_s} \mathbf{p}' \quad M, \xi \xRightarrow{\rho | \gamma \rightarrow \gamma'}_{\text{ML}_c} V_p, \xi', \theta \quad S, \xi' \xRightarrow{\rho | \gamma' + \{\mathbf{p}' \mapsto V_p\} \rightarrow \gamma''}_{\text{ML}_c} V, \xi'', \theta'}{(\text{bind}_m \mathbf{p} = M; S), \xi \xRightarrow{\rho | \gamma \rightarrow \gamma''}_{\text{ML}_c} V, \xi'', \theta @ \theta'} \\
\\
\text{REFERENCE} \qquad \qquad \qquad \text{EXEC} \\
\frac{\gamma(\mathbf{r}) = v}{\mathbf{r}, \xi \xRightarrow{\rho | \gamma \rightarrow \gamma'}_{\text{ML}_c} v, \xi, \langle \rangle} \qquad \frac{\xi \xRightarrow{\rho | \gamma \rightarrow \gamma'}_{\text{ML}_c} \{\}, \theta \quad S, \xi' \xRightarrow{\rho | \gamma' \rightarrow \gamma''}_{\text{ML}_c} V, \xi'', \theta'}{(\text{exec } (); S), \xi + \text{end} + \xi' \xRightarrow{\rho | \gamma \rightarrow \gamma''}_{\text{ML}_c} V, \xi'', \theta @ \theta'} \\
\\
\text{FRAG} \qquad \qquad \qquad \text{DYN} \\
\frac{(\mathbf{f} \ v_1 \dots v_n) \xRightarrow{\rho | \gamma \rightarrow \gamma'}_{\text{ML}_c} v, \theta \quad \xi \xRightarrow{\rho | \gamma' + \{\mathbf{r} \mapsto v\} \rightarrow \gamma''}_{\text{ML}_c} \{\}, \theta'}{\{\mathbf{r} \mapsto \mathbf{f} \ v_1 \dots v_n\} + \xi \xRightarrow{\rho | \gamma \rightarrow \gamma''}_{\text{ML}_c} \{\}, \theta @ \theta'} \qquad \frac{\rho(\text{Dyn}) = \mathbf{R}}{\text{Dyn.f} \xRightarrow{\rho}_{\text{ML}_c} \mathbf{R.r}} \\
\\
\text{FRAG}_m \qquad \qquad \qquad \text{CLIENTCODE} \\
\frac{\mathbf{F}(\mathbf{R}_1) \dots (\mathbf{R}_n) \xRightarrow{\rho \cup \{\text{Dyn} \mapsto \mathbf{R}\} | \gamma \rightarrow \gamma'}_{\text{ML}_c} V, \theta \quad \xi \xRightarrow{\rho | \gamma' + \{\mathbf{R} \mapsto V\} \rightarrow \gamma''}_{\text{ML}_c} \{\}, \theta'}{\{\mathbf{R} \mapsto \mathbf{F} \ \mathbf{R}_1 \dots \mathbf{R}_n\} + \xi \xRightarrow{\rho | \gamma \rightarrow \gamma''}_{\text{ML}_c} \{\}, \theta @ \theta'} \qquad \begin{array}{l} \text{Inherit the} \\ \text{ML rules} \end{array}
\end{array}$$

Fig. 38. Semantics rules for $\text{ML}_c - S, \xi \xRightarrow{\rho | \gamma \rightarrow \gamma'}_{\text{ML}_c} V, \xi', \theta$

uses a is created and bound on the server to b . Finally, b is used on the client via an injection. The program returns 4.

For each fragment, we emit a bind declaration on the client. The client expression contained in the fragment is abstracted and transformed in a closure that is bound to a fresh reference. The number of arguments of the closure corresponds to the number of injections inside the fragment. Similarly to the interpreted semantics, we use the client part of the converter on the client. In this case, $\{\{ 1 \}\}$ is turned into $\lambda().1$ and $\{\{ \text{frag}\%a + 1 \}\}$ is turned into $\lambda v.((\text{frag}^c v) + 1)$. On the server, each fragment is replaced by a call to the primitive fragment. The arguments of the call are the identifier of the closure and all the injections that are contained in the fragment. The fragment primitive, which was presented in [Section 5.1.3](#), registers that the closure declared on the client should be executed later on. Since all the arguments of fragment should be of type `serial`, we apply the client and server parts of the converters at the appropriate places. The `exec` and `end` primitives synchronize the execution so that the order of side effects is preserved. When `exec` is encountered, it executes queued fragment up to an end token which was pushed by an end primitive. We place an `exec/end` pair at each server section. This is enough to ensure that client code inside server fragment and client code in regular client declaration is executed in the expected order.

Note that injections, which occur outside of fragments, and escaped values, which occur inside fragments, are compiled in a very different way. Injections have the useful property that the use site and number of injections is completely static: we can collect all the injections on the server, independently of the client control flow and send them to the client. This is the property that allows us to avoid communications from the client to the server.

ELIOM _ε	ML _s	ML _c
$\text{let}_s a = \{\{ 1 \}\};$	$\text{let } a = \text{fragment } f_0 ();$ $\text{end } ();$	$\text{bind } f_0 = \lambda().1;$ $\text{exec } ();$
$\text{let}_s b = \{\{ \text{frag}\%a + 1 \}\};$	$\text{let } b = \text{fragment } f_1 (\text{frag}^s a);$ $\text{end } ();$	$\text{bind } f_1 = \lambda v.((\text{frag}^c v) + 1);$ $\text{exec } ();$
$\text{let}_c \text{return} = \text{frag}\%b + 2;$	$\text{injection } x (\text{frag}^s b);$	$\text{let } \text{return} = (\text{frag}^c x) + 2;$

Example 14. Compilation of expressions

Sections and modules. We now present an example with client and server modules in [Example 15](#). The lines of the various programs have been laid out in a way that should highlight their correspondence.

We declare a server module X containing a client fragment, a client functor F containing an injection, a client functor application Z and finally the client return value, with another injection. The compilation of the server module X proceeds in the following way: on the server, we emit a module X similar to the original declaration but where fragments have been replaced by a call to the fragment primitive. On the client, we only emit the call to `bind`, without any of the server-side code structure. Compilation for the rest of the code proceeds in a similar manner.

This compilation scheme is correct thanks to the following insight: In client and server modules or functors, the special instructions for fragments and injection can be freely lifted to the outer scope. Indeed, the fragment closure bound in f_0 can only reference client elements. Since the server X can only introduce server-side variables, the body of the fragment closure is independent from the server-side code. A similar remark can be made about the client functor F : the functor argument must be on the client, hence it cannot introduce new server binding. The server identifier that is

injected must have been introduced outside of functor and the injection can be lifted outside the functor.

Using this remark, the structure of the ML_s and ML_c programs is fairly straightforward: Code on the appropriate side has the same shape as in the original $ELIOM_\epsilon$ program and code on the other side only contains calls to the appropriate primitives.

$ELIOM_\epsilon$	ML_s	ML_c
<pre> module_s X = struct let_s a = {{ 2 }} let_s b = 4 end; module_c F(Y : M) = struct let_c a = Y.b + int%X.b end; module_c Z = F(struct let_c b = 2 end); let_c return = frag%X.a + Z.a; </pre>	<pre> module X = struct let a = fragment f₀ () let b = 4 end; end (); injection x₀ (int^s X.b); injection x₁ (frag^s X.a); </pre>	<pre> bind f₀ = λ().2; exec (); module F(Y : M) = struct let a = Y.b + (int^c x₀) end; module Z = F(struct let b = 2 end); let return = (frag^c x₁) + Z.a; </pre>

Example 15. Compilation of client and server modules and functors

Mixed modules. Finally, [Example 16](#) presents the compilation of mixed modules. In this example, we create a mixed structure X containing a server declaration and a client declaration with an injection. We define a functor F that takes a module containing a client integer and use it both inside a client fragment, and inside a client declaration. We then apply F to X and use an injection to compute the final result of the program.

The compilation of the mixed module X is similar to the procedure for programs: we compile each declaration and use the injection primitive as needed. Additionally, we add a `Dyn` field on the server-side version of the module. The content of the `Dyn` field is determined statically for simple structures (here, it is X_0). The client-side version of the module is first bound to X_0 using the bind_m primitive. We then declare X as a simple alias. This alias ensures that X is also usable in client sections transparently.

For functors, the process is similar. One additional complexity is that the `Dyn` field should be dynamically generated. For this purpose, we add a call to the fragment_m primitive. Each call to fragment_m generates a new, fresh identifier. We also prefix each call to fragment by the `Dyn` field. On the client, we emit two different functors. The first one is called F and contains only the client declarations to be used inside the rest of the client code. It is used for client-side usage of mixed functors. An example with the interpreted semantics was presented in [Section 3.6.3](#). The other one is bound to a new reference (here F_1) and contains both client declaration, along with calls to the `bind` and `exec` primitives. This function is used to perform client side effects: when the server version of F is applied, a call to F_1 is registered and will be executed when the client reaches the associated `exec` call (here, the last one).

5.3 Slicing rules

Given an $ELIOM_\epsilon$ module M (resp. module type \mathcal{M} , structure S , ...) and a location ι that is either client or server, we note $\langle M \rangle_\iota$ the result of the compilation of M to the location ι . The result of $\langle M \rangle_s$ is a module of ML_s and the result of $\langle M \rangle_c$ is a module in ML_c .

Let us defines a few notation. As before, we use $e[a \mapsto b]$ to denote the substitution of a by b in e . $e[a_i \mapsto b_i]_i$ denotes the repeated substitution of a_i by b_i in e . We note $\text{FRAGS}(e)$ (resp. $\text{INJS}(e)$)

the fragments (resp. injections) present in the expression e . We note $(e_i)_i$ the sequence of elements e_i . For ease of presentation, we use D_ζ (resp. \mathcal{D}_ζ) for definitions (resp. declarations) located on location ζ . In order to simplify our presentation of mixed functors, both in the slicing rules and in the simulation proofs, we consider a sliceability constraint which dictates which programs can be sliced.

Definition 1 (Sliceability). A program is said sliceable if mixed structures are only defined at top level, or directly inside a toplevel mixed functor.

We give an example of this constraint in [Example 17](#). The program presented on the left is not sliceable, since it contains a structure which is nested inside a structure in a functor. The semantically equivalent program on the right is sliceable, since structures are not nested. This restriction can be relaxed by using a transformation similar to lambda-lifting on mixed functors. In the rest of this section, we assume programs are sliceable and well typed.

We now describe how to slice the various constructions of our language. The slicing rules for modules and expressions are defined in [Figure 39](#). The slicing rules for structures and declarations are presented in [Figure 40](#).

Base structure and signature components are left untouched. Indeed, according to [Proposition 4](#), base elements are valid ML_ε elements. We do not need to modify them in any way. Signature components that are not valid on the target location are simply omitted. Signature components that are valid on the target have their type expressions translated. The translation of a type expression to the client is the identity: indeed, there are no new $ELIOM_\varepsilon$ type constructs that are valid on the client. Server types, on the other hand, can contain pieces of client types inside fragments $\{\tau_c\}$ and inside converters $\tau_s \rightsquigarrow \tau_c$. Fragments in ML_s are represented by a primitive type, `fragment`, without parameters. The type of converters is represented by the type of their server part, which is $\tau_s \rightarrow \text{serial}$. Module and module type expressions are traversed recursively. Functors and functor applications have each part sliced. Mixed functors are turned into normal functors.

Slicing of structure components inserts additional primitives that were described in [Section 5.1](#). In client structure components, we need to handle injections. We associate each injection to a

$ELIOM_\varepsilon$	ML_s	ML_c
<pre> module_m X = struct let_s a = 2 let_c b = 4 + int%a end; </pre>	<pre> module X = struct module Dyn = X₀; let a = 2; end (); injection x₀ (int^s a); end; </pre>	<pre> bind_m X₀ = struct exec (); let b = 4 + (int^c x₀) end; module X = X₀; </pre>
<pre> module_m F(Y : M) = struct let_s c = {{ Y.b }} let_c d = 2 * Y.b end; </pre>	<pre> module F(Y : M) = struct module Dyn = fragment_m F₁ (Y.Dyn); let c = fragment Dyn.f₀ (); end (); end; </pre>	<pre> bind_m F₁(Y : M) = struct bind_m Dyn.f₀ = λ().(Y.b); exec (); let d = 2 * Y.b end; module F(Y : M) = struct let d = 2 * Y.b end; </pre>
<pre> module_m Z = F(X); </pre>	<pre> module Z = F(X); end (); </pre>	<pre> module Z = F(X); exec () </pre>
<pre> let_c return = frag%Z.c + Z.d; </pre>	<pre> injection x₁ (frag^s Z.c); </pre>	<pre> let return = (frag^c x₁) + Z.d; </pre>

Example 16. Compilation of a mixed functor

<pre> module_m F(X : M) = struct module_m Y = struct ... end end </pre> <p>(a) An unsliceable functor</p>	<pre> module_m Y'(X : M) = struct ... end module_m F(X : M) = struct module_m Y = Y'(X) end </pre> <p>(b) A sliceable functor</p>
---	---

Example 17. The sliceability constraint

Type expressions

$$\langle \{\tau\} \rangle_s = \text{frag} \qquad \langle \tau_s \rightsquigarrow \tau_c \rangle_s = \langle \tau_s \rangle_s \rightarrow \text{serial}$$
Signature components

$$\langle \mathcal{D}_b; \mathcal{S} \rangle_t = \mathcal{D}_b; \langle \mathcal{S} \rangle_t \qquad \langle \mathcal{D}_\zeta; \mathcal{S} \rangle_t = \langle \mathcal{S} \rangle_t \quad \text{when } \zeta \neq t$$

$$= \langle \mathcal{D}_\zeta \rangle_t; \langle \mathcal{S} \rangle_t \quad \text{when } \zeta > t$$
Declarations and Definitions

$$\langle \text{type}_t t_i \rangle_t = \text{type } t_i \qquad \langle \text{type}_t t_i = \tau \rangle_t = \text{type } t_i = \langle \tau \rangle_t$$

$$\langle \text{val}_t x_i : \tau \rangle_t = \text{val } x_i : \langle \tau \rangle_t \qquad \langle \text{module}_\zeta X_i : \mathcal{M} \rangle_t = \text{module } X_i : \langle \mathcal{M} \rangle_t$$

$$\langle \text{let}_t x_i = e \rangle_t = \text{let } x_i = e \qquad \langle \text{module}_t X_i = M \rangle_t = \text{module } X_i = \langle M \rangle_t$$
Module Expressions

$$\langle \text{struct } S \text{ end} \rangle_t = \text{struct } \langle S \rangle_t \text{ end}$$

$$\langle M(M') \rangle_t = \langle M \rangle_t (\langle M' \rangle_t)$$

$$\langle \text{functor}(X_i : \mathcal{M})M \rangle_t = \text{functor}(X_i : \langle \mathcal{M} \rangle_t) \langle M \rangle_t$$

$$\langle \text{functor}_m(X_i : \mathcal{M})M \rangle_t = \text{functor}(X_i : \langle \mathcal{M} \rangle_t) \langle M \rangle_t$$
Module Type Expressions

$$\langle \text{sig } S \text{ end} \rangle_t = \text{sig } \langle S \rangle_t \text{ end}$$

$$\langle \text{functor}(X_i : \mathcal{M})M' \rangle_t = \text{functor}(X_i : \langle \mathcal{M} \rangle_t) \langle M' \rangle_t$$

$$\langle \text{functor}_m(X_i : \mathcal{M})M' \rangle_t = \text{functor}(X_i : \langle \mathcal{M} \rangle_t) \langle M' \rangle_t$$
Fig. 39. Slicing - $\langle \cdot \rangle_t$

new fresh reference noted x . In ML_s , we use the injection primitive to register the fact that the given server value should be associated to a given reference. In ML_c , we replace each injection by its associated reference. This substitution is applied both inside expressions and structures. Note that for each injection $f\%x$, we use the encoding part f^s and decoding part f^c for the server and client code, respectively. For server structure components, we apply a similar process to handle fragments. For each fragment, we introduce a reference noted f . In ML_s , we replace each fragment by a call to fragment with argument the associated reference and each escaped value inside the fragment (with the encoding part of the converters). We also add, after the translated component, a call to end which indicates that the execution of the component is finished. In ML_c , we use the bind primitives to associate to each reference a closure where all the escaped values are abstracted.

$$\begin{aligned}
\langle D_b; S \rangle_i &= D_b; \langle S \rangle_i & \langle D_m; S \rangle_i &= \langle D_m \rangle_i; \langle S \rangle_i \\
\langle D_c; S \rangle_s &= (\text{injection } \mathbf{x}_i (f_i^s \mathbf{x}_i);)_i \langle S \rangle_s \\
\langle D_c; S \rangle_c &= D_c \left[\overline{f_i \% \mathbf{x}_i} \mapsto \overline{f_i^c \mathbf{x}_i} \right]; & \text{Where } & \begin{cases} \overline{f_i \% \mathbf{x}_i} = \text{INJS}(D_c) \\ \overline{\mathbf{x}_i} \text{ is a list of fresh variables.} \end{cases} \\
\langle D_s; S \rangle_s &= \langle D_s \rangle_s \left[\{ \{ e_i \} \}_{f_i} \mapsto \text{fragment } f_i \overline{f_{i,j}^s a_{i,j}} \right]; \\
& \text{end } (); \\
& \langle S \rangle_s \\
\langle D_s; S \rangle_c &= \left(\text{bind Dyn. } f_i = \lambda \overline{\mathbf{x}_{i,j}}. e_i [f_{i,j} \% a_{i,j} \mapsto (f_{i,j}^c \mathbf{x}_{i,j})]; \right)_i \\
& \text{exec } (); \\
& \langle S \rangle_c \\
& \text{Where } \begin{cases} \{ \{ e_i \} \}_{f_i} = \text{FRAGS}(D_s) \\ \forall i, \overline{f_{i,j} \% a_{i,j}} = \text{INJS}(e_i) \\ \forall i, \overline{\mathbf{x}_{i,j}} \text{ are fresh variables} \end{cases} \\
\left\langle \begin{array}{l} \text{module}_m F_i(\overline{X_{i_k} : \langle \mathcal{M}_k \rangle_s}) = \text{struct} \\ S \\ \text{end}_F \end{array} \right\rangle_s &= \begin{array}{l} \text{module } F_i(\overline{X_{i_k} : \langle \mathcal{M}_k \rangle_s}) = \text{struct} \\ \text{module Dyn} = \text{fragment}_m F(\overline{X_{i_k} \cdot \text{Dyn}}); \\ \langle S \rangle_s \\ \text{end} \end{array} \\
\left\langle \begin{array}{l} \text{module}_m F_i(\overline{X_{i_k} : \langle \mathcal{M}_k \rangle_c}) = \text{struct} \\ S \\ \text{end}_F \end{array} \right\rangle_c &= \begin{array}{l} \text{bind}_m F(\overline{X_{i_k} : \langle \mathcal{M}_k \rangle_c}) = \text{struct } \langle S \rangle_c \text{ end}; \\ \text{module } F_i(\overline{X_{i_k} : \langle \mathcal{M}_k \rangle_c}) = \text{struct } S |_c \text{ end}; \end{array} \\
\langle \text{module}_m X_i = \text{struct } S \text{ end}_X \rangle_s &= \begin{array}{l} \text{module } X_i = \text{struct} \\ \text{module Dyn} = X; \\ \langle S \rangle_s \\ \text{end}_F \end{array} \\
\langle \text{module}_m X_i = \text{struct } S \text{ end}_X \rangle_c &= \begin{array}{l} \text{bind}_m X = \text{struct } \langle S \rangle_c \text{ end}; \\ \text{module } X_i = X; \end{array} \\
\langle \text{module } X_i = M \rangle_s &= \begin{array}{l} \text{module } X_i = \langle M \rangle_s; \\ \text{end } (); \end{array} \\
\langle \text{module } X_i = M \rangle_c &= \begin{array}{l} \text{module } X_i = \langle M \rangle_c; \\ \text{exec } (); \end{array}
\end{aligned}$$

Fig. 40. Slicing of declarations – $\langle D \rangle_i$

We also introduce the decoding part of each converter for escaped values. We then call `exec`, which executes all the pending fragments until the next `end ()`. This allows to synchronize interleaved side effects between fragments and client components.

Given the constraint of sliceability, a mixed module is either a multi-argument functor returning a structure, or it does not contain any structure at all. For each structure, we use the reference annotated on structures, as described in Section 3.6.2. Mixed modules without structures can simply be sliced by leaving the module expression unchanged. Mixed module types are also straightforward to slice. Mixed structures (with an arbitrary number of arguments) need special care. In ML_s ,

we add a Dyn field to the structure. The value of this field is the result of a call to the primitive fragment_m with arguments F and all the Dyn fields of the arguments of the functor. In ML_c , we create two structures for each mixed structure. One is simply a client functor where all the server parts have been removed. Note here that we don't use the slicing operation. The resulting structure does not contain any call to bind and exec . We also create another structure that uses the regular slicing operation. This structure is associated to F with the bind_m primitive

5.4 Typing preservation

One desirable property is that the introduction of new elements in the language and the compilation operation does not compromise the guarantees provided by the host language. To ensure this, we show that slicing a well typed ELIOM_ε program provides two well typed ML_c and ML_s programs.

We only consider typing environments Γ containing the primitive types frag and serial *i.e.*, $(\text{type}_s \text{ frag}) \in \Gamma$ and $(\text{type}_b \text{ serial}) \in \Gamma$. We also extend the slicing operation to typing environments. Slicing a typing environment is equivalent to slicing a signature with additional rules for converters. Converters, in ELIOM_ε , are not completely first class: they are only usable in injections and not manipulable in the expression language. As such, they must be directly provided by the environment. We add the two following slicing rules that ensures that converters are properly present in the sliced environment:

$$\begin{aligned} \langle \text{val } f : \tau_s \rightsquigarrow \tau_c \rangle_s &= \text{val } f^s : \langle \tau_s \rangle_s \rightsquigarrow \text{serial} \\ \langle \text{val } f : \tau_s \rightsquigarrow \tau_c \rangle_c &= \text{val } f^c : \text{serial} \rightsquigarrow \tau_c \end{aligned}$$

Theorem 3 (Compilation preserves typing). Let us consider M and \mathcal{M} such that $\Gamma \blacktriangleright_m M : \mathcal{M}$. Then $\langle \Gamma \rangle_s \blacktriangleright \langle M \rangle_s : \langle \mathcal{M} \rangle_s$ and $\langle \Gamma \rangle_c \blacktriangleright \langle M \rangle_c : \langle \mathcal{M} \rangle_c$

PROOF. We proceed by induction over the proof tree of $\Gamma \blacktriangleright_m M : \mathcal{M}$. The only difficult cases are client and server structure components and mixed structures. For brevity, we only detail the case of client structure components with one injection.

Let us consider D_c such that $\Gamma \blacktriangleright_m (D_c; S) : \mathcal{S}$ and $\text{INJS}(D_c) = f\%x$. We note x the fresh reference. By definition of the typing relation on ELIOM_ε , there exists Γ' and τ_c, τ_s such that $\Gamma \subset \Gamma'$, $\Gamma' \triangleright_s f : \tau_s \rightsquigarrow \tau_c$ and $\Gamma' \triangleright_s x : \tau_s$. We observe that there cannot be any server bindings in D_c , Hence we can assume $\Gamma' = \Gamma$. This is illustrated on the proof tree below.

$$\frac{\frac{\Gamma \triangleright_s f : \tau_s \rightsquigarrow \tau_c \quad \overline{\Gamma \triangleright_s x : \tau_s}}{\Gamma \triangleright_c f\%x : \tau_c} \quad \vdots}{\Gamma \blacktriangleright_m (D_c; S) : \mathcal{S}}$$

By definition of slicing on typing environments, $(\text{val } f^s : \langle \tau_s \rangle_s \rightarrow \text{serial}) \in \langle \Gamma \rangle_s$ and $(\text{val } f^c : \text{serial} \rightarrow \tau_c) \in \langle \Gamma \rangle_c$. By definition of ML_c and ML_s typing rules, we have $\langle \Gamma \rangle_s \triangleright_{\text{ML}_s} (f^s x) : \text{serial}$ and $\langle \Gamma \rangle_c \triangleright_{\text{ML}_c} (f^c x) : \tau_c$.

We easily have that $\langle \Gamma \rangle_s \blacktriangleright_{\text{ML}_s} \text{injection } x (f^s x) : \varepsilon$, as seen on the proof tree below.

$$\frac{\frac{(\text{val } f^s : \langle \tau_s \rangle_s \rightarrow \text{serial}) \in \langle \Gamma \rangle_s}{\langle \Gamma \rangle_s \triangleright_{\text{ML}_s} f^c : \langle \tau_s \rangle_s \rightsquigarrow \text{serial}} \quad \overline{\langle \Gamma \rangle_s \triangleright_{\text{ML}_s} x : \langle \tau_s \rangle_s}}{\langle \Gamma \rangle_s \blacktriangleright_{\text{ML}_s} \text{injection } x (f^s x) : \varepsilon}$$

By induction hypothesis on Γ , $(\text{val}_c x_j : \tau_c) \triangleright_m d_c [f\%x \mapsto x_j] : \varepsilon$ where v_j is fresh, we have $\langle \Gamma \rangle_c, (\text{val } x_j : \tau_c) \triangleright_{\text{ML}_c} d_c [f\%x \mapsto x_j] : \varepsilon$. We can then replace the proof tree of v_j by the one of

($f^c \mathbf{x}$). We simply need to ensure that the environments coincide. This is the case since f^c cannot be introduced by new bindings. We can then remove the binding of v_j from the environment, since it is unused. We obtain that $\langle \Gamma \rangle_c \triangleright_{ML_c} d_c[f\%x \mapsto (f^c \mathbf{x})] : \varepsilon$ which allows us to conclude.

$$\frac{\frac{(\text{val } f^c : \text{serial} \rightarrow \tau_c) \in \langle \Gamma \rangle_c}{\langle \Gamma \rangle_c \triangleright_{ML_c} f^c : \text{serial} \rightsquigarrow \tau_c} \quad \frac{}{\langle \Gamma \rangle_c \triangleright_{ML_c} \mathbf{x} : \text{serial}}}{\frac{\langle \Gamma \rangle_c \triangleright_{ML_c} (f^c \mathbf{x}) : \tau_c}{\vdots}}{\langle \Gamma \rangle_c \triangleright_{ML_c} D_c[f\%x \mapsto (f^c \mathbf{x})] ; S : S}$$

□

5.5 Semantics preservation

We now state that the compilation process preserves the semantics of $ELIOM_\varepsilon$ programs. In order to do that, we show that, given an $ELIOM_\varepsilon$ program P , the trace of its execution is the same as the concatenation of the traces of $\langle P \rangle_s$ and $\langle P \rangle_c$.

First, let us put some constraints on the constants of the $ELIOM_\varepsilon$, ML_s and ML_c language:

Hypothesis 1 (Well-behaved converters). Converters are said to be well-behaved if for each constant c in $Const$ such that $TypeOf(c) = \tau_s \rightsquigarrow \tau_c$, then $c^s \in Const_s$ and $c^c \in Const_c$.

We now assume that converters in $ELIOM_\varepsilon$, ML_s and ML_c are *well-behaved*. We can then state the following theorem.

Theorem 4 (Compilation preserves semantics). Given sets of constants where converters are well-behaved, given an $ELIOM_\varepsilon$ program P respecting the slicability hypothesis and such that $P \xrightarrow{\{\}} v, \theta$ then

$$\langle P \rangle_s \xrightarrow{\{\}}_{ML_s} (), \xi, \zeta, \theta_s \quad \langle P \rangle_c, \xi \xrightarrow{\{\} | \zeta \rightarrow \gamma}_{ML_c} v, \xi', \theta_c \quad \theta = \theta_s @ \theta_c$$

PROOF. The complete proof is given in [Appendix B](#). The proof proceed in the following way: First, we simplify the problem by applying a code transformation that hoist injections to the top level. We then proceed with a proof by simulation for client code ([Lemma 1](#)), server code ([Lemma 2](#)) and finally mixed code ([Lemma 3](#)). □

This theorem not only show that the return value is the same, but also that the trace is identical. This means that side effects happen in the same order and in the same location as specified by the interpreted semantics. While our simplified calculus does not have side effects, OCAML (and thus ELIOM) do, making such guarantee essential.

6 STATE OF THE ART AND COMPARISON

ELIOM takes inspiration from many sources. The two main influences are, naturally, the extremely diverse ecosystem of web programming languages and frameworks, which we explore in [Section 6.1](#), and the long lineage of ML programming languages. One of the important contributions of ELIOM is the use of a programming model similar to languages for distributed systems ([Section 6.2](#)) while using an execution model inspired by staged meta-programming ([Section 6.3](#)).

6.1 Web programming

Various directions have been explored to simplify Web development and to adapt it to current needs. ELIOM places itself in one of these directions, which is to use the same language on the server and the client. Several unified client-server languages have been proposed. They can be

split in two categories depending on their usage of JAVASCRIPT. JAVASCRIPT can either be used on the server, with NODE.JS, or as a compilation target, for example with GOOGLE WEB TOOLKIT for Java or EMSCRIPTEEN for C. The approach of compiling to JAVASCRIPT was also used to develop new client languages aiming to address the shortcomings of JAVASCRIPT. Some of them are new languages, such as HAXE, ELM or DART. Others are only JAVASCRIPT extensions, such as TYPESCRIPT or COFFEESCRIPT.³

However, these proposals only address the fact that JAVASCRIPT is an inadequate language for Web programming. They do not address the fact that the model of Web programming itself – server and client aspects of web applications are split in two distinct programs with untyped communication – raises usability, correctness and efficiency issues. A first attempt at tackling these concerns is to specify the communication between client and server. Such examples includes SOAP⁴ (mostly used for RPCs) and REST⁵ (for Web APIs). A more recent attempt is the GraphQL [GraphQL 2016] query language which attempts to describe, with a type system, the communications between the client and server parts of the application. These proposals are very powerful and convenient ways to check and document Web-based APIs. However, while making the contract between the client and the server more explicit, they further separate web applications into distinct tiers.

Tierless languages attempt to go in the opposite direction: by removing tiers and allowing web applications to be expressed in one single program, they make the development process easier and restore the possibility of encapsulation and abstraction without compromising correctness. In the remainder of this section, we attempt to give a fairly exhaustive taxonomy of tierless programming languages. We first give a high-level overview of the various trade-offs involved, then we give a detailed description of each language.

6.1.1 Code and data location. In ELIOM, code and data locations are specified through syntactic annotations. Other approaches for determining locations have been proposed. The first approach is to infer locations based on known elements through a control flow analysis (Stip.js, OPA, UR/WEB): database access is on the server, dynamic DOM interaction is done on the client, etc. Another approach is to extend the type system with locations information (LINKS, ML5). Locations can then be determined by relying on simple type inference and checking.

These various approaches present a different set of compromises:

- We believe that the semantics of a language should be easy to predict by looking at the code, which is why ELIOM uses syntactic annotations to specify locations. This fits well within the OCAML language, which is specifically designed to have a predictable behavior. On the other end of the spectrum, languages with inferred location sacrifice predictability for a very light-weight syntax which provides very little disruption over the rest of the program. Typed based approaches sit somewhere in the middle: locations are not visible in the code but are still accessible through types. Such approaches benefit greatly from IDEs allowing exploring inferred types interactively.
- Naturally, explicit approaches are usually more expressive than implicit approaches. Specifying locations manually gives programmers greater control over the performance of their applications. Furthermore, it allows to express mixed *data structures*, *i.e.*, data structures that contain both server and client parts as presented in Section 2.3. Such idioms are difficult to express when code locations are inferred. We demonstrate this with an example in the UR/WEB description.

³A fairly exhaustive list of languages compiling to JAVASCRIPT can be found in <https://github.com/jashkenas/coffeescript/wiki/List-of-languages-that-compile-to-JS>

⁴<https://en.wikipedia.org/wiki/SOAP>

⁵https://en.wikipedia.org/wiki/Representational_state_transfer

- Type-directed approaches to infer code location is an extremely elegant approach. It can be employed either in an algebraic effect setting (LINKS) or as modal logic annotations (ML5). By its type-directed nature, error messages can be expressed in term of the source language and it should lend itself naturally to separate compilation (although this has not yet been achieved). However, such novel type systems significantly extend traditional general purpose type systems (the ML one, in this case) to the point where it seems difficult to retrofit them on an existing languages. One lead would be to provide a tight integration using a form of Foreign Function Interface. Such integration has yet to be proposed.
- ELIOM, as a language based on OCAML, is an effectful language. Marrying inference of locations and a side-effecting semantics is delicate. The Stip.js library [Philips et al. 2014] attempts to solve this by automatically providing replication and eventual consistency on shared references. This might cause many more communications that necessary if not done carefully. We believe that such decision is better left in the hands of the programmer.

6.1.2 Slicing. Once code location has been determined, the tierless program must be sliced in (at least) two components. In ELIOM, slicing is done statically at compile time in a modular manner: each module is sliced independently. Another common approach is to use a static whole-program slicing transformation (UR/WEB, Stip.js). This is most common for languages where code location is inferred, simply due to the fact that such inference is often non-modular. This allows precise analysis of location that can benefit from useful code transformations such as CPS transformation [Philips et al. 2016], inlining and defunctionalization. However, this can make it difficult for the users to know where each piece of code is executed and hinder error messages,. It also prevents any form of separate compilation.

Finally, slicing can be done at runtime simply by generating client JAVASCRIPT code “on the fly” during server execution (LINKS, HOP, PHP). Such solution has several advantages: it is easier to implement and provides a very flexible programming style by allowing programmers to compose the client program in arbitrary ways. The downside is that it provides less guarantees to the users. Furthermore, it prevents generating and optimizing a single JAVASCRIPT file in advance, which is beneficial for caching and execution purposes.

Separate and incremental compilation. Most current mainstream compiled language support some form of incremental compilation. Indeed, incremental compilation avoids recompiling files of which no dependency has changed. This accelerates the feedback loop between development and testing greatly and allow very fast recompilation times. In the case of statically typed languages, it also allows immediate checking of the modified file thus providing developers very fast iteration cycles. The easiest way to implement incremental compilation is through separate compilation, where each file can be compiled completely independently. Furthermore, separate compilation is compatible with link-time optimization and thus does not prevent generation of heavily optimized code, as demonstrated by nearly every C compiler. As a consequence, we consider languages that do not support incremental compilation completely unusable for practical usages.

6.1.3 Communications. ELIOM uses asymmetric communication between client and server (see Section 2.2.5): everything needed to execute the client code is sent during the initial communication that also sends the Web page. It also exposes a convenient API for symmetric communications using RPC (Section 2.3.1) and broadcasts, which must be used manually.

We thus distinguish several kind of communications. First, manual communications are exposed through normal APIs and are performed explicitly by programmers. Of course, the convenience and safety of such functions vary a lot depending on the framework. Then, we consider automatic

communications, that are inserted automatically by the language at appropriate points, as determined by code locations and slicing. We can further decompose automatic communications further in two categories. In static asymmetric communications, information is sent from the server to the client automatically, when sending the page. In dynamic symmetric communications, information is sent back and forth between the client and the server dynamically through some form of channel (AJAX, websockets, Comet, ...).

While symmetric communications are very expressive, they impose a significant efficiency overhead: a permanent connection must be established and each communication imposes a round trip between client and server. Furthermore, such communication channel must be very reliable. On the other hand, asymmetric communications are virtually free: data is sent with the web page (and is usually much smaller). Only a thin instrumentation is needed. Of course, the various communication methods can be mixed in arbitrary manner. ELIOM, for example, uses both automatic asymmetric and manual communications.

Offline usage. Many web applications are also used on Mobile phones, where connection is intermittent at best. As such, we must consider the case where the web application produced by a tierless language is used offline. In this context, asymmetric communication offer a significant advantage: given the initially transmitted information by the server, the client program can run perfectly fine without connection. This guarantee, however, does not extend to dynamic manual communications done by the use of RPCs and channels. Philips et al. [2014] explore this question for symmetric communications through their R5 requirement.

6.1.4 Type systems. Type safety in the context of tierless languages can encompass several notions. The first notion is the traditional distinction between weakly and strongly typed languages. In the interest of avoiding a troll war among the jury, we will not comment further. A more interesting question is whether communication errors between client and server are caught by the typechecker. This is, surprisingly, not the case of UR/WEB since location inference and slicing is done very late in the compilation process, far after type checking. One consequence of this is that slicing errors are fairly difficult to understand [Chlipala 2015b, page 10].⁶ While the ELIOM formalization is type safe, the ELIOM implementation is not, due to the use of wrapping and Marshall, which will fail at runtime on functional values.

Another remark is the distinction between client and server universes.⁷ ELIOM has separate type universes for client and server types (see Section 3.3.3). Most tierless languages do not provide such distinction, notably for the purpose of convenience. Distributed systems such as ACUTE, however, do make such distinction to provide a solution for API versionning and dynamic reloading of code. In this case, there are numerous distinct type universes.

Module systems. The notion of module system varies significantly depending on the language. In ELIOM we consider an ML-style module system composed of a small typed language with structures and functors. We believe modules are essential for building medium to large sized programs: this has been demonstrated for general purpose languages but also holds for web programming languages, as demonstrated by the size of large modern websites (the web frontend of facebook alone is over 9 millions lines of code). Even JAVASCRIPT recently obtained a module system in ES6. In the context of tierless languages, an interesting question is the interaction between locations and modules. In particular, can modules contain elements of different locations and, for statically typed languages, are locations reflected in signatures?

⁶ “However, the approach we adopted instead, with ad-hoc static analysis on whole programs at compile time, leads to error messages that confuse even experienced Ur/Web programmers.”

⁷Or, more philosophically: Is your favorite language platonist or nominalist ?

Types and documentation. Type systems are indisputably very useful for correctness purposes, but they also serve significant documentation purposes. Indeed, given a function, its type signature provides many properties. In traditional languages, this can range from very loose (arguments and return types) to very precise (with dependent types and parametricity [Wadler 1989]). In the context of tierless languages, important questions we might want to consider are “Where can I call this function?” and “Where should this argument come from?”. The various languages exposes this information in different ways: ELIOM does not expose location in the types, but it is present in the signature. ML5 exposes this information directly in the types. UR/WEB and LINKS do not expose that information at all.

6.1.5 Details on some specific approaches. We now provide an in-depth comparison with the most relevant approaches. A summary in Figure 41 classifies each approach according to the main distinctive features described in the previous paragraphs. Each language or framework is also described below.

	Locations	Slicing	Communications	Type safe	Host language
ELIOM	Syntactic	Modular	Asymmetric	✓	OCAML
LINKS	Type-based*	Dynamic*	Symmetric	✓	-
UR/WEB	Inferred	Global	(A)symmetric~	✓*	-
Haste	Type-based	Modular	Symmetric	✓	HASKELL
HOP	Syntactic	Dynamic*	(A)symmetric~	×	JAVASCRIPT*
Meteor.js	Syntactic	Dynamic	Manual	×	JAVASCRIPT
Stip.js	Inferred	Global	Symmetric*	×	JAVASCRIPT
ML5	Type-based	Global*	Symmetric	✓	-
Acute	Syntactic	Modular	Distributed	✓	OCAML

Fig. 41. Summary of the various tierless languages

See previous sections for a description of each headline. A star * indicates that details are available in the description of the associated language. A tilde ~ indicates that we are unsure, either because the information was not specified, or because we simply missed it.

UR/WEB [Chlipala 2015a,b] is a new statically typed language especially designed for Web programming. It features a rich ML-like type and module system and a fairly original execution model where programs only execute as part of a web-server request and do not have any state (the language is completely pure). While similar in scope to ELIOM, it follows a very different approach: Location inference and slicing are done through a whole-program transformation operated on a fairly low level representation. Notably, this transformation relies on inlining and removal of high-order functions (which are not supported by the runtime). The advantages of this approach are twofold: It makes UR/WEB applications extremely fast (in particular because it doesn’t use a GC: memory is trashed after each request) and it requires very little syntactic overheads, allowing programs to be written in a very elegant manner.

The downsides, however, are fairly significant. UR/WEB’s approach is incompatible with any form of separate compilation. Many constructs are hard-coded into the language, such as RPCs and reactive signals and it does not seem possible to implement them as libraries. The language is clearly not general and has a limited expressivity, in particular when trying to use mixed data-structures (see Section 2.3). For example, Example 18. presents the server function `button_list` which takes a list of labels and client functions and generates a list of buttons. We show the ELIOM implementation and a tentative UR/WEB implementation. The UR/WEB version typechecks but

slicing fails. We are unable to write a working version and do not believe it to be possible: indeed, in the ELIOM version we use a client fragment to build the list l as a mixed data-structure. This annotation is essential to make the desired semantics explicit. Some examples are expressible only using reactive signals, which present a very different semantics.

```

1 let client handler _ = alert "clicked!"
2 let server l =
3   [ ("Click!", [%client handler]) ]
4
5 let server button_list lst =
6   ul (List.map (fun (name, action) ->
7     li [button
8       ~button_type : Button
9       ~a:[a_onclick action]
10      [pdata name]])
11     lst)
12
13 let main () =
14   body (button_list l)

```

(a) ELIOM version

```

1 fun main () : transaction page =
2 let
3   fun handler _ = alert "clicked!"
4   val l = Cons ("Click!", handler), Nil)
5
6 fun button_list lst =
7   case lst of
8     Nil => <xml/>
9   | Cons ((name, action), r) =>
10    <xml>
11      <button value={name}
12        onclick={action}/>
13      {button_list r}
14    </xml>
15 in
16   return <xml>
17     <body>{button_list l}</body>
18   </xml>
19 end

```

(b) Tentative UR/WEB version. Typechecks but does not compile.

Example 18. Programs building a list of buttons from a list of client side actions

HOP [Serrano et al. 2006] is a dialect of Scheme for programming Web applications. Its successor, HOP.js [Serrano and Prunet 2016], takes the same concepts and brings them to JAVASCRIPT. The implementation of HOP.js is very complete and allow them to run both the JAVASCRIPT and the scheme dialect while leveraging the complete node.js ecosystem. HOP uses very similar language constructions to the one provided by ELIOM: \sim -expressions are fragments and $\$$ -expressions are injections. All functions seem to be **shared** by default. Communications are asymmetric when possible and use channels otherwise. However, contrary to ELIOM, slicing is done dynamically during server execution [Loitsch and Serrano 2007]. In the tradition of Scheme, HOP only uses a minimal type system for optimizations and does not have a notion of location. In particular HOP does not provide static type checking and does not statically enforce the separation of client and server universes (such as preventing the use of database code inside the client). The semantics of HOP has been formalized [Boudol et al. 2012; Serrano and Queinnec 2010] and does present similarities to the interpreted ELIOM semantics (Section 3.6). HOP is however significantly more dynamic than ELIOM: it allows dynamic communication patterns through the use of channels and allows nested fragments in the style of Lisp quotations which allows to generate client code inside client code.

For dynamically-typed inclined programmers, HOP currently presents the most convincing approach to tierless Web programming. In particular given its solid implementation, great flexibility and support for the JAVASCRIPT ecosystem.

LINKS [Cooper et al. 2006] is an experimental functional language for client-server Web programming with a syntax close to JAVASCRIPT and an ML-like type system. Its type system is

extended with a notion of *effects*, allowing a clean integration of database queries in the language [Lindley and Cheney 2012]. In Example 19, we highlight two notable points of LINKS: the function `adults` takes as argument a list `l` and returns the name of all the person over 18. This function has no effect and can thus run on the client, the server, but can also be transformed into SQL to run in a database query. On the other hand, the `print` function has an effect called “wild” which indicates it can’t be run inside a query. Effects are also used to provide type-safe channel-based concurrency.

LINKS also allows to annotate functions by indicating on which location they should run. Those annotations, however, are not reflected in the type system. Communications are symmetric and completely dynamic through the use of AJAX. Client-server slicing is dynamic (although some progress has been made towards static *query* slicing [Cheney et al. 2014]) and can introduce “code motion”, which can moves closures from the server to the client. This can be extremely problematic in practice, both from an efficiency and a security point of view. The current implementation of LINKS is interpreted but a compilation scheme leveraging the Multicore-OCAML efforts has been recently added.

Although LINKS is very seducing, the current implementation presents many shortcomings given its statically typed nature: slicing is dynamic and produces fairly large JAVASCRIPT code and the type system does not really track client-server locations.

```

1 links> fun adults(l) { for (x <- l) where (x.age >= 18) [(name = x.name)] } ;;
2 adults = fun : ((age:Int, name:a|_) -> [(name:a)])
3
4 links> print ;;
5 print : (String) {wild}-> ()

```

Example 19. Small pieces of LINKS code

METEOR.JS [Meteor.js 2017] is a framework where both the client and the server sides of an application are written in JAVASCRIPT. It has no built-in mechanism for sections and fragments but relies on conditional `if` statements on the `Meteor.isClient` and `Meteor.isServer` constants. It does not perform any slicing. This means that there are no static guarantees over the respective execution of server and client code. Besides, it provides no facilities for client-server communication such as fragments and injections. Compared to ELIOM, this solution only provides coarse-grained composition.

STIP.JS [Philips et al. 2014] allows to slice tierless JAVASCRIPT programs with a minimal amount of annotations. It emits METEOR.JS programs with explicit communications. Annotations are optionally provided through the use of comments, which means that STIP.JS are actually perfectly valid JAVASCRIPT programs. Location inference and slicing are whole-program static transformations. Communications are symmetric, through the use of fairly elaborate consistency and replication mechanisms for shared references. This approach allows the programmer to write code with very little annotations. As opposed to UR/WEB, manual annotations are possible, which might allow to express delicate patterns such as mixed data-structures and prevents security issues.

6.2 Distributed programming

Tierless languages in general are very inspired by distributed programming languages. The main difference being that distributed programs contain an arbitrary number of locations while tierless web programs only have two: client and server. Communications are generally symmetric

and dynamic, due to the multi-headed aspect of distributed systems. There are of course numerous programming languages dedicated to distributed programming. We present here two relevant approaches that put greater emphasis on the typing and tierless aspects.

Haste [Ekblad 2017] is a HASKELL EDSL for distributed programming. This DSL allows to express complex orchestrations of multiple nodes and external components (for example databases and IoT components), with handling of distinct type universes when necessary. Instead of using syntactic annotations, locations are determined through typing. This approach works particularly well in the context of HASKELL, thanks to the advanced type system and the syntactic support for monads and functors. Multiple binaries are produced from one program. Slicing relies on type information and dead code elimination, as provided by the GHC compiler. Explicit slicing markers similar to ELIOM's section annotations are the subject of future work. Communications are dynamic and symmetric through the use of websockets. One notable feature of this DSL is that it offers a client-centric view: The control flow is decided by the client which pilots the other nodes. This is the opposite of ELIOM where the server can assemble pieces of client code through fragments. This work also inherits the HASKELL and GHC features in term of modules, data abstraction and separate-compilation. A module language has been developed for HASKELL by Kilpatrick et al. [2014].

An earlier version, HASTE.APP [Ekblad and Claessen 2014], was limited to only one client and one server and used a monadic approach to structure tierless programs.

ML5 [Murphy VII et al. 2007] is an ML language that introduces new constructs for type-safe communication between distributed actors through the use of location annotations inside the types called "modal types". It is geared towards a situation where all actors have similar capabilities. It uses dynamic communication, which makes the execution model very different from ELIOM. ML5 provides a very rich type system that allows to precisely export the capabilities of the various locations. For example, it is possible to talk about addresses on distant locations and pass them around arbitrary. ELIOM only supports such feature through the use of fragments, for client code.

Unfortunately, ML5 does not have a module system. However, we believe that ML5's modal types can play a role similar to ELIOM's location annotations on declarations, including location polymorphism. ML5 uses a global transformation for slicing. Given the rich typing information present in ML5's types, it should lend itself fairly well to a modular slicing approach, but this has not been done.

ACUTE [Sewell et al. 2007] is an extension of OCAML for distributed programming. It provides typesafe serialization and deserialization and also allows arbitrary loading of modules at runtime. Like ELIOM, it provides a full-blown module system. However, it takes an opposite stance on the execution model: each actor runs independent programs and communications are completely dynamic.

Handling of multiple type universes is done by providing a description of the type with each message and by versioning APIs. In particular, great care is taken to provide type safe serialization by also transmitting the type of messages alongside each message. This gives ACUTE very interesting capabilities, such as reloading only part of the distributed system in a type-safe way.

6.3 Staged meta-programming

An important insight regarding ELIOM is that, while it is a tierless programming language and tries to disguise itself as a distributed programming language, ELIOM corresponds exactly to a staged meta-programming language. ELIOM simply provides only two stages: stage 0 is the server, stage 1 is the client. ELIOM's client fragments are the equivalent of stage quotations.

Most approaches to partial evaluation are done implicitly (not unlike tierless languages with implicit locations). We take inspiration from several approaches that combine staged meta-programming with explicit stages annotations that are reflected in the type system. We only look at the most relevant approaches but a longer description of the history of staged meta-programming approaches can be found in Taha [1999, Chapter 7].

METAOCAML [Kiselyov 2014] is an extension of OCAML for meta programming. It introduces a quotation annotation for staged expressions, whose execution is delayed. Quotations and antiquotations corresponds exactly to fragments and injections. The main difference is that METAOCAML is much more dynamic: quoted code does not have to be completely closed when produced and well-scopedness is checked dynamically, just before running the quoted code. This allows very dynamic behaviors such as automatic insertion of let-bindings [Kiselyov 2015] and dynamically determining staged stream pipelines [Kiselyov et al. 2017]. One difference is the choice of universes: ELIOM has two universes, client and server, which are distinct. METAOCAML has a single type universe but a series of scopes, for each stage, included in one another.

METAOCAML itself provides no support for modules and only leverages the OCAML module system. Staging annotations are only on expressions, not on declarations.

Modular macros [Nicole 2016; Yallop and White 2015] are another extension of OCAML. It uses staging to implement macros. It provides both a quotation-based expression language along with staging annotations on declarations. It also aims to support modules and functors. The slicing can be seen as dynamic (since code is executed at compile time to produce pieces of programs). In particular, this allows to lift most of the restriction imposed on multi-stage functors. They also use a notion similar to converters, except that the serialization format here is simply the OCAML AST.

The main difference compared to ELIOM is how the asymmetry between stage 0 and stage 1 is treated. Only one type universe is used and there is no notion of slicing that would allow a distant execution.

Feltman et al. [2016] presents a slicing technique for a two-staged simply typed lambda calculus. Their technique is similar to the one used in ELIOM. They distinguish their language in three parts: $1\mathcal{G}$, which corresponds to base code; $1\mathcal{M}$, which corresponds to server code; and $2\mathcal{M}$, which corresponds to client code. They also provide a proof of equivalence between the dynamic semantics and the slicing techniques. This proof has been mechanized in Twelf. While their work is done in a more general settings, they do not specify how to transfer rich data types across stages (which is solved in ELIOM using converters). They also do not propose a module system.

7 CONCLUSION

We presented ELIOM, a Modular Tierless Web programming language that supports static typing, an efficient execution model, abstraction and modularity. We showed its design through numerous examples and formalized it. In our formalization, we first gave a simple interpreted semantics, that is suitable to explain to programmers. We then gave a more efficient compilation scheme and showed that it preserves both typing and semantics. The compilation scheme follows precisely the current implementation of ELIOM.

Many design choices of ELIOM were inspired by practical concerns. Indeed, in order to be useful, a language must have an ecosystem. The simplest way to have an ecosystem is to reuse the one of an existing language, in our case the OCAML one. Since separate compilation and modularity are indispensable for any non-trivial programming projects, we needed to make modularity, abstraction and the tierless annotations interact. ELIOM is thus inspired by many different elements such as ML

languages, tierless web programming and staged meta-programming. In particular, ELIOM aims to combine a semantics inspired by staged meta-programming and an ML-style module system.

We believe that such combination can be used in a wider context than simply Web programming. Indeed, OCSIGEN has also been used to create Mobile applications [Besport 2017]. We would like to extend our language to be usable to any kind of client-server applications that requires dynamicity, type safety and modularity. We also consider extending our model to explicitly support multiple distinct clients. Finally, we aim to further explore the programming patterns made possible by ELIOM, for instance multi-tiers functional reactive programming.

REFERENCES

- Nada Amin and Tiark Rompf. 2017. Type soundness proofs with definitional interpreters. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 666–679. <https://doi.org/10.1145/3009837>
- Brian E. Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. 2008. Engineering formal metatheory. In *Proceeding of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, George C. Necula and Philip Wadler (Eds.). ACM, 3–15. <http://arthur.chargueraud.org/research/2007/binders/>
- Vincent Balat. 2014. Rethinking Traditional Web Interaction: Theory and Implementation. *International Journal on Advances in Internet Technology* (2014). http://www.ariajournals.org/internet_technology/
- Besport 2017. *BeSport*. <http://www.besport.com/>.
- Didier Le Botlan and Didier Rémy. 2003. ML^F: raising ML to the power of system F. In *ICFP*. ACM, 27–38.
- Gérard Boudol, Zhengqin Luo, Tamara Rezk, and Manuel Serrano. 2012. Reasoning about Web Applications: An Operational Semantics for HOP. *ACM Trans. Program. Lang. Syst.* 34, 2 (2012), 10.
- James Cheney, Sam Lindley, Gabriel Radanne, and Philip Wadler. 2014. Effective quotation: relating approaches to language-integrated query. In *Proceedings of the ACM SIGPLAN 2014 workshop on Partial evaluation and program manipulation, PEPM 2014, January 20-21, 2014, San Diego, California, USA*, Wei-Ngan Chin and Jurriaan Hage (Eds.). ACM, 15–26. <https://doi.org/10.1145/2543728.2543738>
- Adam Chlipala. 2015a. An Optimizing Compiler for a Purely Functional Web-Application Language. In *ICFP*.
- Adam Chlipala. 2015b. Ur/Web: A Simple Model for Programming the Web. In *Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '15)*. ACM, New York, NY, USA, 153–165. <https://doi.org/10.1145/2676726.2677004>
- Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. 2006. Links: Web Programming Without Tiers. In *FMCO*. 266–296.
- Anton Eklblad. 2017. A meta-EDSL for distributed web applications. In *Proceedings of the 10th ACM SIGPLAN International Symposium on Haskell, Oxford, United Kingdom, September 7-8, 2017*, Iavor S. Diatchki (Ed.). ACM, 75–85. <https://doi.org/10.1145/3122955.3122969>
- Anton Eklblad and Koen Claessen. 2014. A Seamless, Client-centric Programming Model for Type Safe Web Applications. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell (Haskell '14)*. ACM, New York, NY, USA, 79–89. <https://doi.org/10.1145/2633357.2633367>
- Nicolas Feltman, Carlo Angiuli, Umut A. Acar, and Kayvon Fatahalian. 2016. Automatically Splitting a Two-Stage Lambda Calculus. See [Thiemann 2016], 255–281. https://doi.org/10.1007/978-3-662-49498-1_11
- J. Garrigue. 2009. A Certified Interpreter for ML with Structural Polymorphism. (2009).
- Gencore 2017. *New York University Gencore*. <http://gencore.bio.nyu.edu/>.
- GraphQL 2016. GraphQL. (2016). <http://graphql.org/>
- Hashids 2017. *Hashids*. <http://hashids.org/>.
- Scott Kilpatrick, Derek Dreyer, Simon L. Peyton Jones, and Simon Marlow. 2014. Backpack: retrofitting Haskell with interfaces. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, Suresh Jagannathan and Peter Sewell (Eds.). ACM, 19–32. <https://doi.org/10.1145/2535838.2535884>
- Oleg Kiselyov. 2014. The Design and Implementation of BER MetaOCaml - System Description. In *Functional and Logic Programming - 12th International Symposium, FLOPS 2014, Kanazawa, Japan, June 4-6, 2014. Proceedings (Lecture Notes in Computer Science)*, Michael Codish and Eijiro Sumii (Eds.), Vol. 8475. Springer, 86–102. https://doi.org/10.1007/978-3-319-07151-0_6
- Oleg Kiselyov. 2015. Generating Code with Polymorphic let: A Ballad of Value Restriction, Copying and Sharing. In *Proceedings ML Family / OCaml Users and Developers workshops, ML Family/OCaml 2015, Vancouver, Canada, 3rd & 4th September 2015. (EPTCS)*, Jeremy Yallop and Damien Doligez (Eds.), Vol. 241. 1–22. <https://doi.org/10.4204/EPTCS.241.1>

- Oleg Kiselyov, Aggelos Biboudis, Nick Palladinis, and Yannis Smaragdakis. 2017. Stream fusion, to completeness. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 285–299. <http://dl.acm.org/citation.cfm?id=3009880>
- Daniel K. Lee, Karl Crary, and Robert Harper. 2007. Towards a mechanized metatheory of standard ML. In *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007*, Martin Hofmann and Matthias Felleisen (Eds.). ACM, 173–184. <https://doi.org/10.1145/1190216.1190245>
- Xavier Leroy. 1994. Manifest Types, Modules, and Separate Compilation. In *Conference Record of POPL'94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon, USA, January 17-21, 1994*, Hans-Juergen Boehm, Bernard Lang, and Daniel M. Yellin (Eds.). ACM Press, 109–122. <https://doi.org/10.1145/174675.176926>
- Xavier Leroy. 1995. Applicative Functors and Fully Transparent Higher-Order Modules. In *Conference Record of POPL'95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*, Ron K. Cytron and Peter Lee (Eds.). ACM Press, 142–153. <https://doi.org/10.1145/199448.199476>
- Xavier Leroy. 1996. A Syntactic Theory of Type Generativity and Sharing. *J. Funct. Program.* 6, 5 (1996), 667–698. <https://doi.org/10.1017/S0956796800001933>
- Sam Lindley and James Cheney. 2012. Row-based effect types for database integration. In *Proceedings of TLDI 2012: The Seventh ACM SIGPLAN Workshop on Types in Languages Design and Implementation, Philadelphia, PA, USA, Saturday, January 28, 2012*, Benjamin C. Pierce (Ed.). ACM, 91–102. <https://doi.org/10.1145/2103786.2103798>
- Florian Loitsch and Manuel Serrano. 2007. Hop Client-Side Compilation. In *Proceedings of the Eighth Symposium on Trends in Functional Programming, TFP 2007, New York City, New York, USA, April 2-4, 2007. (Trends in Functional Programming)*, Marco T. Morazán (Ed.), Vol. 8. Intellect, 141–158.
- Meteor.js 2017. *Meteor.js*. <http://meteor.com>.
- Tom Murphy VII, Karl Crary, and Robert Harper. 2007. Type-Safe Distributed Programming with ML5. In *TGC (Lecture Notes in Computer Science)*, Gilles Barthe and Cédric Fournet (Eds.), Vol. 4912. Springer, 108–123.
- Olivier Nicole. 2016. Bringing typed, modular macros to OCaml. (2016). https://oliviernicole.github.io/about_macros.html
- Ocsigen Toolkit 2017. *Ocsigen Toolkit*. <http://ocsigen.org/ocsigen-toolkit/>.
- Scott Owens. 2008. A Sound Semantics for OCamlLight. In *Programming Languages and Systems, 17th European Symposium on Programming, ESOP 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings (Lecture Notes in Computer Science)*, Sophia Drossopoulou (Ed.), Vol. 4960. Springer, 1–15. https://doi.org/10.1007/978-3-540-78739-6_1
- Scott Owens, Magnus O. Myreen, Ramana Kumar, and Yong Kiam Tan. 2016. Functional Big-Step Semantics, See [Thiemann 2016], 589–615. https://doi.org/10.1007/978-3-662-49498-1_23
- Laure Philips, Coen De Roover, Tom Van Cutsem, and Wolfgang De Meuter. 2014. Towards Tierless Web Development Without Tierless Languages. In *Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software (Onward! 2014)*. ACM, New York, NY, USA, 69–81. <https://doi.org/10.1145/2661136.2661146>
- Laure Philips, Joeri De Koster, Wolfgang De Meuter, and Coen De Roover. 2016. Dependence-driven delimited CPS transformation for JavaScript. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences, GPCE 2016, Amsterdam, The Netherlands, October 31 - November 1, 2016*, Bernd Fischer and Ina Schaefer (Eds.). ACM, 59–69. <https://doi.org/10.1145/2993236.2993243>
- Pumgrana 2017. *Pumgrana*. <http://www.pumgrana.com/>.
- Gabriel Radanne, Vasilis Papavasileiou, Jérôme Vouillon, and Vincent Balat. 2016a. Eliom: tierless Web programming from the ground up. In *IFL 2016, Leuven, Belgium, August 31 - September 2, 2016*, Tom Schrijvers (Ed.). ACM, 8:1–8:12. <https://doi.org/10.1145/3064899.3064901>
- Gabriel Radanne and Jérôme Vouillon. 2018. Tierless Web programming in the large. In *WWW 2018, Lyon, Paris*.
- Gabriel Radanne, Jérôme Vouillon, and Vincent Balat. 2016b. Eliom: A Core ML Language for Tierless Web Programming. In *APLAS 2016, Hanoi, Vietnam, November 21-23, 2016, Proceedings (Lecture Notes in Computer Science)*, Atsushi Igarashi (Ed.), Vol. 10017. 377–397. https://doi.org/10.1007/978-3-319-47958-3_20
- Andreas Rossberg, Claudio V. Russo, and Derek Dreyer. 2014. F-ing modules. *J. Funct. Program.* 24, 5 (2014), 529–607. <https://doi.org/10.1017/S0956796814000264>
- Manuel Serrano, Erick Gallesio, and Florian Loitsch. 2006. Hop: a language for programming the Web 2.0. In *OOPSLA Companion*. 975–985.
- Manuel Serrano and Vincent Prunet. 2016. A glimpse of Hopjs. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 180–192. <https://doi.org/10.1145/2951913.2951916>
- Manuel Serrano and Christian Queinnec. 2010. A multi-tier semantics for Hop. *Higher-Order and Symbolic Computation* 23, 4 (2010), 409–431.

- Peter Sewell, James J. Leifer, Keith Wansbrough, Francesco Zappa Nardelli, Mair Allen-Williams, Pierre Habouzit, and Viktor Vafeiadis. 2007. Acute: High-level programming language design for distributed computation. *J. Funct. Program.* 17, 4-5 (2007), 547–612. <https://doi.org/10.1017/S0956796807006442>
- Walid Taha. 1999. *Multi-stage Programming: Its Theory and Applications*. Ph.D. Dissertation. Oregon Graduate Institute of Science and Technology.
- Peter Thiemann (Ed.). 2016. *Programming Languages and Systems - 25th European Symposium on Programming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*. Lecture Notes in Computer Science, Vol. 9632. Springer. <https://doi.org/10.1007/978-3-662-49498-1>
- Mads Tofte. 1988. *Operational Semantics and Polymorphic Type Inference*. Ph.D. Dissertation. University of Edinburgh.
- Tutorial 2017. *Ocsigen Tutorial*. <https://ocsigen.org/tuto/manual>.
- TyXML 2017. *TyXML*. <http://ocsigen.org/tyxml/>.
- Jérôme Vouillon. 2008. Lwt: a cooperative thread library. In *Proceedings of the ACM Workshop on ML, 2008, Victoria, BC, Canada, September 21, 2008*, Eijiro Sumii (Ed.). ACM, 3–12. <https://doi.org/10.1145/1411304.1411307>
- Jérôme Vouillon and Vincent Balat. 2014. From bytecode to JavaScript: the Js_of_ocaml compiler. *Software: Practice and Experience* 44, 8 (2014), 951–972. <https://doi.org/10.1002/spe.2187>
- Philip Wadler. 1989. Theorems for Free!. In *Proceedings of the fourth international conference on Functional programming languages and computer architecture, FPCA 1989, London, UK, September 11-13, 1989*, Joseph E. Stoy (Ed.). ACM, 347–359. <https://doi.org/10.1145/99370.99404>
- Andrew K Wright and Matthias Felleisen. 1994. A syntactic approach to type soundness. *Information and computation* 115, 1 (1994), 38–94.
- Jeremy Yallop and Leo White. 2015. Modular macros. *OCaml Workshop* (2015). <http://www.lpw25.net/ocaml2015-abs1.pdf>

A THE ML CALCULUS

This section describes our version of the ML calculus, which contains the minimal amount of ingredients that allows us to describe the ELIOM extensions: a core calculus with polymorphism, let bindings and parametrized datatypes in the style of Wright and Felleisen [1994], accompanied by a fully featured module system with separate compilation and applicative functors in the style of Leroy [1995]. We first present the syntax of the language in Appendix A.1 and its type system in Appendix A.2. We then present the semantics of this language in Appendix A.3. The description in this appendix is self-contained, and thus restate the basic description found in Section 3.1.

A.1 Syntax

Expressions

$e ::= c$	(Constant)
$ x_i p.x$	(Variables)
$ Y$	(Fixpoint)
$ (e e)$	(Application)
$ \lambda x.e$	(Function)
$ \text{let } x = e \text{ in } e$	(Let binding)
$c \in \text{Const}$	(Constants)

Path

$p ::= X_i | p.X | p_1(p_2)$

Type Schemes

$\sigma ::= \forall(\alpha)^*. \tau$

Type Expressions

$\tau ::= \alpha$	(Type variables)
$ \tau \rightarrow \tau$	(Function types)
$ (\tau^*)_{t_i} (\tau^*)p.t$	(Type constructors)

(a) The expression language

Module Expressions

$M ::= X_i p.X$	(Variables)
$ (M : M)$	(Type constraint)
$ M_1(M_2)$	(Functor application)
$ \text{functor}(X_i : M)M$	(Functor)
$ \text{struct } S \text{ end}$	(Structure)

Module types

$\mathcal{M} ::= \text{sig } S \text{ end}$	(Signature)
$ \text{functor}(X_i : \mathcal{M}_1)\mathcal{M}_2$	(Functor)

Signature body

$S ::= \varepsilon | \mathcal{D}; S$

Signature components

$\mathcal{D} ::= \text{val } x_i : \tau$	(Values)
$ \text{type } t_i = \tau$	(Types)
$ \text{type } t_i$	(Abstract types)
$ \text{module } X_i : \mathcal{M}$	(Modules)

Structure body

$S ::= \varepsilon | \mathcal{D}; S$

Structure components

$D ::= \text{let } x_i = e$	(Values)
$ \text{type } t_i = \tau$	(Types)
$ \text{module } X_i = M$	(Modules)

Environments

$\Gamma ::= S$

Programs

$P ::= \text{prog } S \text{ end}$

(b) The module language

Fig. 42. ML grammar

Let us first define some notations and meta-syntactic variables. As a general rule, the expression language is in lowercase (e) and the module language is in uppercase (M). Module types are in calligraphic letters (\mathcal{M}). More precisely: x are variables, p are module paths, X are module variables, τ are type expressions and t are type constructors. x_i , X_i and t_i are identifiers (for values,

modules and types). Identifiers (such as x_i) have a name part (x) and a stamp part (i) that distinguish identifiers with the same name. Only the name part of identifiers is exposed in module signature. α -conversion should keep the name intact and change only the stamp, which allow to preserve module signatures. Sequences are noted with a star; for example τ^* is a sequence of type expressions. Indexed sequences are noted (τ_i) , with an implicit range. Substitution of a by b in e is noted $e[a \mapsto b]$. Repeated substitution of each a_i by the corresponding b_i is noted $e[a_i \mapsto b_i]_i$. The syntax is presented in [Figure 42](#).

Expressions. The expression language is a fairly simple extension of the lambda calculus with a fixpoint combinator Y and let bindings $\text{let } x = e_1 \text{ in } e_2$. The language is parametrized by a set of constants *Const*. Variables can be qualified by a module path p . Paths can be either module identifiers such as X_i , a submodule access such as $X_i.Y$, or a path application such as $X_i(Y_j.Z)$. Note that, as said earlier, that fields of modules are only called by their name, without stamp.

Types. Types are composed of type variables α , function types $\tau_1 \rightarrow \tau_2$ and parametrized type constructors $(\tau_1, \tau_2, \dots, \tau_k)t_i$. Type constructors can have an arbitrary number of parameters, including zero. Type constructors can be qualified by a module path p . Type schemes, noted σ , are type expressions that are universally quantified by a list of type variables. Type schemes can also have free variables. For example: $\forall \alpha. (\alpha, \beta)t_i$.

Modules. The module language is quite similar to a simple lambda calculus: Functors are functions over module (except that arguments are annotated with their types). Module application is noted $M_1(M_2)$. Modules can also be constrained by a module type: $(M : \mathcal{M})$. Finally, a module can be a structure which contains a list of value, types or module definitions: `struct let $x_i = 2$ end`. Programs are lists of definitions.

Module types. Module types can be either the type of a functor or a signature, which contains a list of value, types and module descriptions. Type descriptions can expose their definition or can be left abstract. Typing environments are simply module signatures. We note them Γ for convenience.

A.2 Type system

We now present the ML type system. For ease of presentation, we proceed in two steps: we will first forget that the module language exists, and present a self-contained type system for the expression language. We then extend the typing relation to handle modules.

A.2.1 The expression language. We introduce the following judgments:

$\Gamma \triangleright e : \tau$ The expression e has type τ in the environment Γ . See [Figure 44](#).

$\Gamma \triangleright \tau_1 \approx \tau_2$ Types τ_1 and τ_2 are equivalent in environment Γ . See [Figure 45](#).

$\Gamma \vDash \tau$ The type τ is well formed in the environment Γ . See [Figure 43](#).

We note $\text{TypeOf}(c)$ the type scheme of a given constant c . The instantiation relation is noted $\sigma > \tau$ for a type scheme σ and a type τ . The converse operation which closes a type according to an environment is noted $\text{Close}(\Gamma, \tau)$. We use $(D) \in \Gamma$ to test if a given type or value is declared in the environment Γ . Note that for types, $(\text{type } (\alpha_i)t) \in \Gamma$ holds also if t is not abstract in Γ .

Polymorphism. One of the main benefit of programming language of the ML family is the ability to easily define and use functions that operate on values of various types. For example, the `map` function can applies to all lists, regardless of the type of their content. Indeed, the type of `map` is polymorphic:

$$\text{map} : \forall \alpha \beta. (\alpha \rightarrow \beta) \rightarrow (\alpha)\text{list} \rightarrow (\beta)\text{list}$$

From a type checking point of view, this is possible thanks to two operations: instantiation and abstraction. Instantiation takes a type scheme, which is a type where some variables have been universally quantified, and replace all the quantified type variables by some type. It is used when looking up a variable (rule VAR) or typechecking a constant (rule CONST). For example, the type of map can be instantiated to the following type.

$$(\text{int} \rightarrow \text{bool}) \rightarrow (\text{int})\text{list} \rightarrow (\text{bool})\text{list}$$

Once instantiated, the map function can be applied on a list with concretes types. Naturally, we also need the converse operation: constructing a type scheme given a type containing some type variables. Closing a type depends on the current typing environments, we only abstract type variables that have not been introduced by previous binders. $\text{Close}(\Gamma, \tau)$ returns the type scheme $\forall \alpha_1 \dots \alpha_n. \tau$ where the α_i are free variables of τ that are not present in Γ . While it is possible to apply the closing operation at any step of a typing derivation, it is only useful at the introduction point of type variables, in let bindings (rule LETIN). In the following example, we derive a polymorphic type for a function that constructs a pair with an element from the environment. We first use the close operation to obtain a type scheme for f . Note that since α is present in the environment, it is not universally quantified. We then use the instance operation to apply f to an integer constant.

$$\frac{\begin{array}{c} \vdots \\ \hline (\text{val } a : \alpha; \text{val } b : \beta) \triangleright (a, b) : \alpha * \beta \end{array}}{(\text{val } a : \alpha) \triangleright \lambda b. (a, b) : \beta \rightarrow \alpha * \beta} \quad \frac{\forall \beta. \beta \rightarrow \alpha * \beta \triangleright \text{int} \rightarrow \alpha * \text{int} \quad \text{Const}(3) \triangleright \text{int}}{\Gamma \triangleright f : \text{int} \rightarrow \alpha * \text{int} \quad \Gamma \triangleright 3 : \text{int}}}{(\text{val } a : \alpha; \text{val } f : \forall \beta. \beta \rightarrow \alpha * \beta) \triangleright f \ 3 : (\alpha * \text{int})} \quad \frac{}{(\text{val } a : \alpha) \triangleright \text{let } f = \lambda b. (a, b) \text{ in } f \ 3 : (\alpha * \text{int})}$$

Parametric datatypes. Parametric polymorphism introduces type variables in type expressions. In the presence of type definitions, it is natural to expect the ability to write type definitions which can contain type variables. This leads us to parametric datatypes: datatypes which are parametrized by a set of variables. $(\alpha)\text{list}$ is of course an example of such datatype. Note that care must be taken when deciding the equivalence of types. If the type is not abstract, *i.e.*, its definition is available, we can always unfold the definition, as shown in rule DEFTYPEEQ. However, when considering an abstract type, we cannot unfold the type definition. Instead, we check that head symbols are compatible and that parameters are equivalent pairwise⁸. This is done in rule ABSTYPEEQ.

$$\frac{\text{TYPEVAL} \quad (\text{type } (\alpha_i)t) \in \Gamma \quad \forall i, \Gamma \models \tau_i}{\Gamma \models (\tau_i)t} \quad \frac{\text{ARROWVAL} \quad \Gamma \models \tau_1 \quad \Gamma \models \tau_2}{\Gamma \models \tau_1 \rightarrow \tau_2} \quad \frac{\text{VARVAL}}{\Gamma \models \alpha}$$

Fig. 43. Type validity rules – $\Gamma \models \tau$

A.2.2 *The module language.* We introduce the following judgments:

- $\Gamma \triangleright M : \mathcal{M}$ The module M is of type \mathcal{M} in Γ . See Figure 46.
- $\Gamma \triangleright \mathcal{M} <: \mathcal{M}'$ The module type \mathcal{M} is a subtype of \mathcal{M}' in Γ . See Figure 47.
- $\Gamma \models \mathcal{M}$ The module type \mathcal{M} is well-formed in Γ . See Figure 48.

The typing rules for OCAML-style modules are quite complex. In particular, the inner details of the rules are not well known, even by OCAML programmers. Before presenting the typing rules

⁸This is similar to the handling of free symbols in the unification literature.

$$\begin{array}{c}
\text{VAR} \\
\frac{(\text{val } x : \sigma) \in \Gamma \quad \sigma > \tau}{\Gamma \triangleright x : \tau} \\
\\
\text{LAMBDA} \\
\frac{\Gamma; (\text{val } x : \tau_1) \triangleright e : \tau_2}{\Gamma \triangleright \lambda x. e : \tau_1 \rightarrow \tau_2} \\
\\
\text{CONST} \\
\frac{\text{TypeOf}(c) > \tau}{\Gamma \triangleright c : \tau} \\
\\
\text{LETIN} \\
\frac{\Gamma \triangleright e_1 : \tau_1 \quad \Gamma; (\text{val } x : \text{Close}(\tau_1, \Gamma)) \triangleright e_2 : \tau_2}{\Gamma \triangleright \text{let } x = e_1 \text{ in } e_2 : \tau_2} \\
\\
\text{EQUIV} \\
\frac{\Gamma \triangleright e : \tau_1 \quad \Gamma \triangleright \tau_1 \approx \tau_2}{\Gamma \triangleright e : \tau_2} \\
\\
\text{APP} \\
\frac{\Gamma \triangleright e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \triangleright e_2 : \tau_1}{\Gamma \triangleright (e_1 e_2) : \tau_2} \\
\\
\text{Y} \\
\frac{\text{Y}}{\Gamma \triangleright \text{Y} : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2} \\
\\
\text{Close}(\tau, \Gamma) = \forall \alpha_0 \dots \alpha_n. \tau \text{ with } \{\alpha_0, \dots, \alpha_n\} = \text{FreeTypeVar}(\tau) \setminus \text{FreeTypeVar}(\Gamma)
\end{array}$$

Fig. 44. ML expression typing rules – $\Gamma \triangleright e : \tau$

$$\begin{array}{c}
\text{REFLEQ} \\
\frac{}{\Gamma \triangleright \tau \approx \tau} \\
\\
\text{TRANSEQ} \\
\frac{\Gamma \triangleright \tau_1 \approx \tau_2 \quad \Gamma \triangleright \tau_2 \approx \tau_3}{\Gamma \triangleright \tau_1 \approx \tau_3} \\
\\
\text{COMMEQ} \\
\frac{\Gamma \triangleright \tau_2 \approx \tau_1}{\Gamma \triangleright \tau_1 \approx \tau_2} \\
\\
\text{FUNEQ} \\
\frac{\Gamma \triangleright \tau_1 \approx \tau'_1 \quad \Gamma \triangleright \tau_2 \approx \tau'_2}{\Gamma \triangleright \tau_1 \rightarrow \tau_2 \approx \tau'_1 \rightarrow \tau'_2} \\
\\
\text{DEFTYPEEQ} \\
\frac{(\text{type } (\alpha_i) t = \tau) \in \Gamma}{\Gamma \triangleright (\tau_i) t \approx \tau[\alpha_i \mapsto \tau_i]_i} \\
\\
\text{ABSTYPEEQ} \\
\frac{(\text{type } (\alpha_i) t) \in \Gamma \quad \forall i, \Gamma \triangleright \tau_i \approx \tau'_i}{\Gamma \triangleright (\tau_i) t \approx (\tau'_i) t}
\end{array}$$

Fig. 45. Type equivalence rules – $\Gamma \triangleright \tau \approx \tau'$

in details, we will attempt to give insight on why some features are present in the languages and what are their advantages. For this purpose, we present two examples illustrating the need for applicative functors and strengthening, respectively. We assume that readers are familiar with simpler usages of ML modules.

Applicative Functors. Let us consider the following scenario: we are given a module G implementing a graph data-structure and would like to implement a simple graph algorithm which takes a vertex and returns all the accessible vertices. We would like the returned module to contain a function of type $G.\text{graph} \rightarrow G.\text{vertex} \rightarrow \text{set_of_vertices}$. How to implement `set_of_vertices`? An easy but inefficient way would be to use lists. A better way is to use proper sets (implemented with balanced binary tree, for example). In OCAML, this is provided in the standard library by the functor `Set.Make`, presented in [Section 2.1](#), which takes a module implementing comparison functions for the given type. We would obtain a signature similar to the one below.

```

module Access(G : Graph) : sig
  module VerticesSet : sig ... end
  val run : G.graph → G.vertex → VerticesSet.set
end

```

However, this means we need to expose a complete module implementing set of vertices that is independent from any other set module. This prevents modularity, since any usage of our new

function must use this specific set implementation. Furthermore, this make the signature bigger than strictly necessary. What we really want to expose is that the return type comes from an application of *Set.Make*. Fortunately, we can do so by using the following signature.

```
module Access(G : Graph) : sig
  val run : G.graph → G.vertex → Set.Make(G.Vertex).set
end
```

Here, we export the fact that the set type must be the result of a functor application on a module that is compatible with *G.Vertex*. The type system guarantees that any such functor application will produces types that are equivalent. In particular, if multiple libraries uses the *Access* functor, their sets will be of the same types, which make composition of libraries easier. This behavior of functors is usually called *applicative*.

Strengthening. Let us now consider the program presented in [Example 20](#). We assume the existence of two modules, presented in [Example 20a](#). The module *Showable* exposes the abstract type *t*, along with a *show* function that turns it into a string. The module *Elt* exposes a type *t* equal to *Showable.t* and a value that inhabits this type. The program is presented in [Example 20b](#). We define a functor *F* taking two arguments *E* and *S* whose signature are similar to *Elt* and *Showable*, respectively. The main difference is that *E* comes first and *S.t* is defined as an alias of *E.t*. The functor uses the *show* function on the element in *E* to create a string. It is natural to expect the functor application *F(Elt)(Showable)* to type check, since *Elt.t = Showable.t*. We must, however, check for module inclusion. While *Elt* is clearly included in the signature of the argument *E*, the same is not clear for *Showable*. We first need to enrich its type signature with additional type equalities. We give *Showable* the type sig type *t = Showable.t ... end*. It makes sense to enrich the signature in such a manner since *Showable* is already in the environment. Given this enriched signature, we can now deduce that $\blacktriangleright(\text{type } t = \text{Showable.t}) <: (\text{type } t = E.t)$ since $E.t = \text{Elt.t} = \text{Showable.t}$.

```
module Showable : sig
  type t
  val show : t → string
end
module Elt : sig
  type t = Showable.t
  val v : elt
end
```

(a) Typing environment

```
module F
  (E : sig type t val v : t end)
  (S : sig type t = E.t val show : t → string end)
= struct
  let s = (S.show E.v)
end
module X = F(Elt)(Showable)
```

(b) Application of multi-argument functor using manifests

Example 20. Program using functors and manifest types

The operation that consists in enriching type signatures of module identifiers with new equalities by using elements in the environment is called *strengthening* [[Leroy 1994](#)].

A.2.3 Typing rules. In the previous two examples, we showcased some delicate interactions between functor, type equalities and modularity in the context of an ML module system. We now see in details how the rules presented in [Figures 46, 47 and 48](#) produce these behaviors.

Qualified access. Unqualified module variables are typechecked in a similar manner than regular variables in the expression language, with the *MODVAR* typing rule. Qualified access (of the form *X.a*), both for the core and the module language, need more work. As with the expression language,

the typing environment is simply a list of declaration. In particular, typing environments do not store paths. This means that in order to prove $\Gamma \triangleright p.a : \tau$, we must first verify that the module p typechecks in $\Gamma : \Gamma \triangleright p : \mathcal{M}$. We then need to verify that the module type \mathcal{M} contains a declaration ($\text{val } a : \tau$). This is done in the `QUALMODVAR` rule for the module language. The rules for the expression language are given in [Figure 49](#).

Let us now try to apply these rules to the module X with the following module type. X contains a type t and a value a of that type. We note that module type \mathcal{X} .

$$X : \text{sig type } t; \text{ val } a : t \text{ end}$$

We wish to typecheck $X.a$. One expected type for this expression is $X.t$. However, the binding of v in \mathcal{X} gives the type t , with no mention of X . We need to prefix the type variable t by the access path X . This is done in the rule `QUALMODVAR` by the substitution $\mathcal{M}[n_i \mapsto p.n \mid n_i \in \text{BV}(\mathcal{S}_1)]$ which prefixes all the bound variables of \mathcal{S}_1 , noted $\text{BV}(\mathcal{S}_1)$, by the path p . Note here that we substitute only by the names declared *before* the variable a . Indeed, a variable or a type can only reference names declared previously in ML. To prove that $X.a$ has the type $X.t$, we can write the following type derivation.

$$\text{QUALVAR} \frac{\text{MODVAR} \frac{(\text{module } X : \mathcal{X}) \in (\text{module } X : \mathcal{X})}{(\text{module } X : \mathcal{X}) \triangleright X : \text{sig type } t; \text{ val } a : t \text{ end}}}{(\text{module } X : \mathcal{X}) \triangleright X.a : X.t} \text{ WITH } X.t = t[t \mapsto p.t]$$

Strengthening. The strengthening operation, noted \mathcal{M}/p , is defined in [Figure 50](#) and is used in the `STRENGTH` rule. It takes a module type \mathcal{M} and a path p and returns a module type \mathcal{M}' where all the type declarations, abstract or not, have been replaced by type aliases pointing to the path p . These type aliases are usually called “manifest types”. This operator relies on the following idea: if p is of type \mathcal{M} , then p is available in the environment. In order to expose as many type equalities as possible, it suffices to give p a type where all the type definition point to definitions available in the environment. This way, we preserve type equalities even for abstract types. This also mean that type equalities can be deduced by only looking at the path and the module type. In particular, we do not need to look at the implementation of p , which is important for the purpose of separate compilation.

Applicative functors. Let us consider a functor F with the following type. It takes a module containing a single type t and return a module containing an abstract type t' and a conversion function.

$$F : \text{functor}(X : \text{sig type } t \text{ end})(\text{sig type } t'; \text{ val } \text{make} : X.t \rightarrow t' \text{ end})$$

If we consider two modules X_1 and X_2 , does $X_1 = X_2$ imply $F(X_1).t = F(X_2).t$? If that is the case, we say that functors are *applicative*. Otherwise, they are *generative*⁹. Here, we consider the applicative behavior of functors. This is implemented with the last strengthening rule which ensures that the body of functors is also strengthened. For example, if \mathcal{M} is the type of the functor above, \mathcal{M}/F is the following module type:

$$\text{functor}(X : \text{sig type } t \text{ end})(\text{sig type } t' = F(X).t; \text{ val } \text{make} : X.t \rightarrow t' \text{ end})$$

This justifies the presence of application inside paths. Otherwise, such type manifests inside functors could not be represented. A more type-theoretic description of generative and applicative functors can be found in [Leroy \[1996\]](#).

⁹SML only supports generative functors. OCAML originally only supported applicative functors, but also supports the generative behavior since version 4.03.

Separate compilation. Separate compilation is an important properties of programming languages. In fact, almost all so-called “mainstream” languages support it. We can distinguish two aspects of this property: separate typechecking and separate code generation. In both cases, it means that in order to process the file (either to type check it or to transform it into another representation), we only need to look at the type of its dependencies, not their implementation.

It turns out that the ML module system with manifest types lends itself very well to separate typechecking [Leroy 1994]. Indeed, let us consider a program as a list of modules. Each module represents a compilation unit (*i.e.*, a file). Since module bindings in the typing environment only contains module types, and not the actual module, typechecking a file only needs the module type of the previous files, which ensure that we can typecheck each file separately, as long as all its dependencies have been typechecked before. This is expressed more formally in [Theorem 5](#).

Theorem 5 (Separate Typechecking). Given a list of module declarations that form a typed program, there exists an order such that each module can be typechecked with only knowledge of the type of the previous modules.

More formally, given a list of n declarations D_i and a signature \mathcal{S} such that

$$\blacktriangleright (D_1; \dots; D_n) : \mathcal{S}$$

then there exists n definitions \mathcal{D}_i and a permutation π such that

$$\forall i < n, \mathcal{D}_1; \dots; \mathcal{D}_i \blacktriangleright \mathcal{D}_{i+1} : \mathcal{D}_{i+1} \quad \blacktriangleright \mathcal{D}_{\pi(1)}; \dots; \mathcal{D}_{\pi(n)} < : \mathcal{S}$$

PROOF. It is always possible to reorder declarations in a signature using the SUBSTRUCT rule. This means we can choose the appropriate permutation of definitions that matches the order of declarations. The rest follows by definition of the typing relation. \square

A.3 Semantics

We now define the semantics of our ML language. We use a rule-based big step semantics with traces. Traces allows us to reason about execution order in a way that is compatible with modules, as we will see in [Appendix A.3.1](#).

We note v for values in the expression language and V for values in the module language. Values are defined in [Figure 52](#). Values in the expression language can be either constants or lambdas. Module values are either structures, which are list of bindings of values, or functors. We note ρ the execution environment. Execution environments are a list of value bindings. Note here that the execution environment is not mutable, since reference cells are not in the language. We note the concatenation of environment $+$. Environment access is noted $\rho(x) = v$ where x has value v in ρ . The same notation is also used for structures. Traces are lists of messages. For now, we consider messages that are values and are emitted with a `print` operation. The empty trace is noted $\langle \rangle$. Concatenation of traces is noted $@$.

Given an expression e (resp. a module m), an execution environment ρ , a value v (resp. V) and a trace θ ,

$$e \xRightarrow{\rho} v, \theta$$

means that e reduces to v in ρ and prints θ . The reduction rules are given in [Figure 53](#). The rules for the expression language are fairly traditional. Variables and paths must be resolved using the VAR and QUALVAR rules. Applications are done in two steps: first, we reduce both the function and the argument with the APP rule, then we apply the appropriate reduction rule for the application: BETA for lambda expressions, Y for fixpoints and DELTA for constants. The δ operation gives meaning to application of a constant to a value. $\delta(c, v) = v', \theta$ means that c applied to v returns v' and emits

the trace θ . Let bindings are treated in a similar manner than lambda expressions: the left hand side is executed, added to the environment, then the right hand side is executed.

The module language has similar rules for identifiers and application. In this case, the BETA and APP rule have been combined in MODBETA. Additional rules for declarations are also present. Type declarations are ignored (TYPEDECL). Values and module declarations (VALDECL and MODDECL) are treated similarly to let bindings: the body of the binding is executed, added to the environment and then the rest of the structure is executed.

A.3.1 Traces and Printing. Traces allow us to visualize the execution order of programs. In particular, if we prove that code transformation preserves traces, it ensures that the execution order is preserved. Traces allow us to reason about execution without introducing references and other side-effecting operations in our language, which would make the presentation significantly more complex.

One example of operation using traces is the print constant. Typing and semantics of print are provided in [Figure 51](#). print accepts any value, prints it, and returns it. From a typing point of view, print has the same type as the identity: a polymorphic function which returns its input. We make use of the fact that the CONST typing rule also uses the instantiation for type schemes. The semantics of print is provided via the DELTA rule: it returns its argument directly but also emits a trace containing the given argument.

We now present an example using print. We assume the existence of the type int, a set of constant corresponding to the integers and an associated operation +. We wish to type and execute the expression e defined as `let x = (print 3) in (print (x + 1))`

Let us first show that e is of type int. The type derivation is provided in [Example 21](#). The typing derivation is fairly direct: we use the CONST rule to type print as $\text{int} \rightarrow \text{int}$ and apply it to integers with the rule APP. We can now look at the execution of e , which returns 4 with a trace $\langle 3; 4 \rangle$. The execution derivation is shown in [Example 22](#). The first step is to decompose the let-binding. We first reduce `(print 3)`, which can be directly done with the DELTA rule. This gives us 3 with a trace $\langle 3 \rangle$. We then reduce `(print (x + 1))` in the environment where x is associated to 3. Before resolving the application of print with the DELTA rule, we need to reduce its argument with the APP rule. We obtain 4 with a trace $\langle 4 \rangle$. We return the result of the right hand side of the left and the concatenation of both traces by usage of the LETIN rule, which gives us 4 with a trace $\langle 3; 4 \rangle$.

A.3.2 Modules. We now present an example of reduction involving modules. Our example program P is presented in [Example 23a](#). It consists of two declarations: a module declaration X which contains a single declaration a , and the return value of the program, which is equal to $X.a$. It is fairly easy to see that the program P return a value of type int, hence we focus on the execution of P , which is presented in [Example 23b](#). The derivation is slightly simplified for clarity. In particular, rules such as EMPTYSTRUCT are elided. The first step is to apply the PROGRAM and MODULEDECL rules in order to execute the content of each declaration. The declaration of X , on the left side, can be reduced by first applying the STRUCT rule in order to extract the content of the module structure, then VALDECL, to reduce the declaration of a . These reductions give us the structure value $\{a \mapsto 3\}$. We now execute the declaration of return. According to the MODULEDECL rule, we must do so in a new environment containing X : $\{X \mapsto \{a \mapsto 3\}\}$. In order to reduce $X.a$, we must use the QUALMODVAR rule, which reduces qualified variables. This means we first reduce X , which according to the environment gives us $\{a \mapsto 3\}$, noted V . We then look up a in V , which returns 3. To return, we first compose the resulting structure value from both declaration: $\{X \mapsto \{a \mapsto 3\}\} + \{\text{return} \mapsto 3\}$. We then lookup return in this structure, which gives us 3.

$$\begin{array}{c}
\text{CONST} \frac{\text{TypeOf}(\text{print}) \triangleright \text{int} \rightarrow \text{int} \quad \text{TypeOf}(3) = \text{int}}{\triangleright \text{print} : \text{int} \rightarrow \text{int}} \quad \frac{\quad}{\triangleright 3 : \text{int}} \quad \frac{\quad}{\triangleright \text{print} : \text{int} \rightarrow \text{int}} \quad \frac{\quad}{\triangleright x + 1 : \text{int}} \\
\text{APP} \frac{\quad}{\triangleright (\text{print } 3) : \text{int}} \quad \frac{\quad}{(\text{val } x : \text{int}) \triangleright (\text{print } (x + 1)) : \text{int}} \\
\text{LETIN} \frac{\quad}{\triangleright \text{let } x = (\text{print } 3) \text{ in } (\text{print } (x + 1)) : \text{int}}
\end{array}$$

Example 21. Typing derivation for $e \triangleright e : \text{int}$

$$\begin{array}{c}
\text{DELTA} \frac{\delta(\text{print}, 3) = 3, \langle 3 \rangle \quad \frac{\quad}{\text{print} \Rightarrow \text{print}, \langle \rangle} \quad \frac{\quad}{x + 1 \xrightarrow{\{x \mapsto 3\}} 4, \langle \rangle} \quad \frac{\delta(\text{print}, 4) = 4, \langle 4 \rangle}{(\text{print } 4) \Rightarrow 4, \langle 4 \rangle}}{\frac{\quad}{(\text{print } 3) \Rightarrow 3, \langle 3 \rangle} \quad \frac{\quad}{(\text{print } (x + 1)) \xrightarrow{\{x \mapsto 3\}} 4, \langle 4 \rangle}}{\text{let } x = (\text{print } 3) \text{ in } (\text{print } (x + 1)) \Rightarrow 4, \langle 3; 4 \rangle}} \\
\text{APP} \quad \text{DELTA} \quad \text{APP} \\
\text{LET} \quad \text{LET}
\end{array}$$

Example 22. Execution derivation for $e \rightarrow e \Rightarrow 4, \langle 3; 4 \rangle$

```

prog
  module X = struct let a = 3 end
  let return = X.a
end

```

(a) The program P

$$\begin{array}{c}
\text{VALDECL} \frac{3 \Rightarrow 3, \langle \rangle}{\text{let } a = 3 \Rightarrow \{a \mapsto 3\}, \langle \rangle} \quad \text{MODVAR} \frac{\rho(X) = V}{X \xrightarrow{\rho} V \equiv \{a \mapsto 3\}, \langle \rangle} \quad \text{QUALMODVAR} \frac{V(a) = 3}{X.a \xrightarrow{\{X \mapsto \{a \mapsto 3\}\}} 3, \langle \rangle} \\
\text{STRUCT} \frac{\left(\begin{array}{l} \text{struct} \\ \text{let } a = 3 \\ \text{end} \end{array} \right) \Rightarrow \{a \mapsto 3\}, \langle \rangle}{\left(\begin{array}{l} \text{module } X = \text{struct let } a = 3 \text{ end} \\ \text{let return} = X.a \end{array} \right) \Rightarrow \{X \mapsto \{a \mapsto 3\}\} + \{\text{return} \mapsto 3\}, \langle \rangle} \\
\text{MODULEDECL} \quad \text{MODULEDECL} \\
\text{PROGRAM} \frac{\quad}{P \Rightarrow 3, \langle \rangle} \quad \text{PROGRAM}
\end{array}$$

(b) Execution of P

Example 23. Example of execution with modules

A.3.3 Notes on Soundness. Soundness properties, which correspond to the often misquoted “Well typed programs cannot go wrong.”, have been proven for many variants of the ML language. Unfortunately, stating and proving the soundness property for big step semantics and ML modules requires a fairly large amount of machinery which we do not attempt to provide here. Instead, we give pointers to various relevant work containing such proofs.

Soundness for a small step semantics of our expression language is provided by [Wright and Felleisen \[1994\]](#). At a larger scale, [Owens \[2008\]](#) proves the soundness of a small step semantics for a very large portion of the OCAML expression language using the Locally Nameless Coq framework

[Aydemir et al. 2008]. Soundness of a big step semantics has been proved and mechanized for several richer languages [Amin and Rompf 2017; Garrigue 2009; Lee et al. 2007; Owens et al. 2016; Tofte 1988].

Unfortunately, as far as we are aware, soundness of Leroy’s module language with higher order applicative functors has not been proved directly and is a fairly delicate subject. The most recent work of interest is [Rossberg et al. 2014], which presents an elaboration scheme from ML modules, including applicative OCAML-style modules, into System F_ω . Soundness then relies on soundness of the elaboration (provided in the article) and soundness of System F_ω . In this work, the applicative/generative behavior of functors is decided depending on its purity, which is much more precise than what is done in OCAML.

$$\begin{array}{c}
\text{MODVAR} \\
\frac{(\text{module } X_i : \mathcal{M}) \in \Gamma}{\Gamma \triangleright X_i : \mathcal{M}} \\
\\
\text{QUALMODVAR} \\
\frac{\Gamma \triangleright p : (\text{sig } \mathcal{S}_1; \text{module } X_i : \mathcal{M}; \mathcal{S}_2 \text{ end})}{\Gamma \triangleright p.X : \mathcal{M}[n_i \mapsto p.n \mid n_i \in \text{BV}(\mathcal{S}_1)]} \\
\\
\text{STRENGTH} \\
\frac{\Gamma \triangleright p : \mathcal{M}}{\Gamma \triangleright p : \mathcal{M}/p} \\
\\
\frac{\Gamma \triangleright M : \mathcal{M}' \quad \Gamma \triangleright \mathcal{M}' <: \mathcal{M}}{\Gamma \triangleright M : \mathcal{M}} \quad \frac{\Gamma \triangleright M_1 : \text{functor}(X_i : \mathcal{M})\mathcal{M}' \quad \Gamma \triangleright M_2 : \mathcal{M}}{\Gamma \triangleright M_1(M_2) : \mathcal{M}'[X_i \mapsto M_2]} \\
\\
\frac{\Gamma \vDash \mathcal{M} \quad X_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{module } X_i : \mathcal{M}) \triangleright M : \mathcal{M}'}{\Gamma \triangleright \text{functor}(X_i : \mathcal{M})M : \text{functor}(X_i : \mathcal{M})\mathcal{M}'} \quad \frac{\Gamma \vDash \mathcal{M} \quad \Gamma \triangleright M : \mathcal{M}}{\Gamma \triangleright (M : \mathcal{M}) : \mathcal{M}} \\
\\
\frac{\Gamma \triangleright e : \tau \quad x_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{val } x_i : \text{Close}(\tau, \Gamma)) \triangleright S : \mathcal{S}}{\Gamma \triangleright (\text{let } x_i = e; s) : (\text{val } x_i : \tau; \mathcal{S})} \\
\\
\frac{\Gamma \vDash \tau \quad t_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{type } t_i = \tau) \triangleright S : \mathcal{S}}{\Gamma \triangleright (\text{type } t_i = \tau; s) : (\text{type } t_i = \tau; \mathcal{S})} \\
\\
\frac{\Gamma \triangleright M : \mathcal{M} \quad X_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{module } X_i : \mathcal{M}) \triangleright S : \mathcal{S}}{\Gamma \triangleright (\text{module } X_i = M; s) : (\text{module } X_i : \mathcal{M}; \mathcal{S})} \quad \frac{\Gamma \triangleright S : \mathcal{S}}{\Gamma \triangleright \text{struct } S \text{ end} : \text{sig } \mathcal{S} \text{ end}} \\
\\
\frac{}{\Gamma \triangleright \varepsilon : \varepsilon}
\end{array}$$

Fig. 46. Module typing rules – $\Gamma \triangleright m : \mathcal{M}$

$$\begin{array}{c}
\text{SUBSTRUCT} \\
\frac{\pi : [1; m] \rightarrow [1; n] \quad \forall i \in [1; m], \Gamma; \mathcal{D}_1; \dots; \mathcal{D}_n \triangleright \mathcal{D}_{\pi(i)} <: \mathcal{D}'_i}{\Gamma \triangleright (\text{sig } \mathcal{D}_1; \dots; \mathcal{D}_n \text{ end}) <: (\text{sig } \mathcal{D}'_1; \dots; \mathcal{D}'_n \text{ end})} \\
\\
\frac{\Gamma \triangleright \tau_1 \approx \tau_2}{\Gamma \triangleright (\text{val } x_i : \tau_1) <: (\text{val } x_i : \tau_2)} \quad \frac{\Gamma \triangleright \mathcal{M}_1 <: \mathcal{M}_2}{\Gamma \triangleright (\text{module } X_i : \mathcal{M}_1) <: (\text{module } X_i : \mathcal{M}_2)} \\
\\
\frac{\Gamma \triangleright \mathcal{M}'_a <: \mathcal{M}_a \quad \Gamma, (\text{module } X : \mathcal{M}'_a) \triangleright \mathcal{M}_r <: \mathcal{M}'_r}{\Gamma \triangleright \text{functor}(X : \mathcal{M}_a)\mathcal{M}_r <: \text{functor}(X : \mathcal{M}'_a)\mathcal{M}'_r} \quad \frac{\Gamma \triangleright \tau_1 \approx \tau_2}{\Gamma \triangleright (\text{type } t_i = \tau_1) <: (\text{type } t_i = \tau_2)} \\
\\
\frac{}{\Gamma \triangleright (\text{type } t_i) <: (\text{type } t_i)} \quad \frac{\Gamma \triangleright t_i \approx \tau}{\Gamma \triangleright (\text{type } t_i) <: (\text{type } t_i = \tau)} \quad \frac{}{\Gamma \triangleright (\text{type } t_i = \tau_1) <: (\text{type } t_i)}
\end{array}$$

Fig. 47. Module subtyping rules – $\Gamma \triangleright \mathcal{M} <: \mathcal{M}'$

$$\begin{array}{c}
\frac{\Gamma \vDash \mathcal{S}}{\Gamma \vDash \text{sig } \mathcal{S} \text{ end}} \quad \frac{}{\Gamma \vDash \varepsilon} \quad \frac{\Gamma \vDash \mathcal{M}_a \quad x_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{module } X_i : \mathcal{M}_a) \vDash \mathcal{M}_r}{\Gamma \vDash \text{functor}(X_i : \mathcal{M}_a) \mathcal{M}_r} \\
\frac{t_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{type } t_i) \vDash \mathcal{S}}{\Gamma \vDash \text{type } t_i; \mathcal{S}} \quad \frac{\Gamma \vDash \mathcal{M} \quad x_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{module } X_i : \mathcal{M}) \vDash \mathcal{S}}{\Gamma \vDash \text{module } X_i : \mathcal{M}; \mathcal{S}} \\
\frac{\Gamma \vDash \tau \quad t_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{type } t_i = \tau) \vDash \mathcal{S}}{\Gamma \vDash \text{type } t_i = \tau; \mathcal{S}} \quad \frac{\Gamma \vDash \tau \quad x_i \notin \text{BV}(\Gamma) \quad \Gamma; (\text{val } x_i : \tau) \vDash \mathcal{S}}{\Gamma \vDash \text{val } x_i : \tau; \mathcal{S}}
\end{array}$$

Fig. 48. Module type validity rules – $\Gamma \vDash \mathcal{M}$

$$\begin{array}{c}
\frac{\text{QUALVAR} \quad \Gamma \triangleright p : (\text{sig } \mathcal{S}_1; \text{val } x_i : \tau; \mathcal{S}_2 \text{ end})}{\Gamma \triangleright p.x : \tau[n_i \mapsto p.n \mid n_i \in \text{BV}(\mathcal{S}_1)]} \quad \frac{\text{QUALDEFTYPEEQ} \quad \Gamma \triangleright p : (\text{sig } \mathcal{S}_1; \text{type } t_i = \tau; \mathcal{S}_2 \text{ end})}{\Gamma \triangleright (\tau_i)p.t \approx \tau[n_i \mapsto p.n \mid n_i \in \text{BV}(\mathcal{S}_1)][\alpha_i \mapsto \tau_i]_i} \\
\frac{\text{QUALABSTYPEEQ} \quad \Gamma \triangleright p : (\text{sig } \mathcal{S}_1; \text{type } t_i; \mathcal{S}_2 \text{ end}) \quad \forall i, \Gamma \triangleright \tau_i \approx \tau'_i}{\Gamma \triangleright (\tau_i)p.t \approx (\tau'_i)p.t}
\end{array}$$

Fig. 49. Additional typing rules for the expression language

$$\begin{array}{c}
\varepsilon/p = \varepsilon \\
(\text{sig } \mathcal{S} \text{ end})/p = \text{sig } \mathcal{S}/p \text{ end} \\
(\text{module } X_i = \mathcal{M}; \mathcal{S})/p = \text{module } X_i = \mathcal{M}/p; \mathcal{S}/p \\
(\text{type } t_i = \tau; \mathcal{S})/p = \text{type } t_i = (\alpha^*)p.t; \mathcal{S}/p \\
(\text{type } t_i; \mathcal{S})/p = \text{type } t_i = (\alpha^*)p.t; \mathcal{S}/p \\
(\text{val } x_i : \tau; \mathcal{S})/p = \text{val } x_i : \tau; \mathcal{S}/p \\
(\text{functor}(X_i : \mathcal{M})\mathcal{M}')/p = \text{functor}(X_i : \mathcal{M})(\mathcal{M}'/p(X_i))
\end{array}$$

Fig. 50. Module strengthening operation – \mathcal{M}/p

$$\begin{array}{c}
\text{PRINTTY} \quad \text{PRINTEXEC} \\
\text{TypeOf}(\text{print}) = \forall \alpha. (\alpha \rightarrow \alpha) \quad \delta(\text{print}, v) = v, \langle v \rangle
\end{array}$$

Fig. 51. Typing and execution rules for print

Expressions	Modules	Bindings
$v ::= c$ (Constant)	$V ::= \{ V_b^* \}$ (Structure)	$V_b ::= \{ x_i \mapsto v \}$ (Values)
$ \lambda x. \rho. e$ (Function)	$ \text{functor}(\rho)(X_i : \mathcal{M})M$	$ \{ X_i \mapsto V \}$ (Modules)

Fig. 52. ML values

$\frac{\text{VAR} \quad \rho(x) = v}{x \xRightarrow{\rho} v, \langle \rangle}$	$\frac{\text{QUALVAR} \quad p \xRightarrow{\rho} V, \theta \quad V(x) = v}{p.x \xRightarrow{\rho} v, \theta}$	$\frac{\text{CONSTANT}}{c \xRightarrow{\rho} c, \langle \rangle}$	$\frac{\text{CLOSURE}}{\lambda x. e \xRightarrow{\rho} \lambda x. \rho. e, \langle \rangle}$
$\frac{\text{LETIN} \quad e' \xRightarrow{\rho} v', \theta \quad e \xRightarrow{\rho + \{x_i \mapsto v'\}} v, \theta'}{(\text{let } x = e' \text{ in } e) \xRightarrow{\rho} v, \theta @ \theta'}$	$\frac{\text{APP} \quad e \xRightarrow{\rho} v, \theta \quad e' \xRightarrow{\rho} v', \theta' \quad (v \ v') \xRightarrow{\rho} v'', \theta''}{(e \ e') \xRightarrow{\rho} v'', \theta @ \theta' @ \theta''}$		
$\frac{\text{BETA} \quad e \xRightarrow{\rho' + \{x_i \mapsto v\}} v', \theta}{(\lambda x. \rho'. e \ v) \xRightarrow{\rho} v', \theta}$	$\frac{\text{Y} \quad (v \ \lambda x. (\text{Y } v \ x)) \xRightarrow{\rho} v', \theta}{(\text{Y } v) \xRightarrow{\rho} v', \theta}$	$\frac{\text{DELTA} \quad \delta(c, v) = v', \theta}{(c \ v) \xRightarrow{\rho} v', \theta}$	
$\frac{\text{MODVAR} \quad \rho(X) = V}{X \xRightarrow{\rho} V, \langle \rangle}$	$\frac{\text{QUALMODVAR} \quad p \xRightarrow{\rho} V', \theta \quad V'(X) = V}{p.X \xRightarrow{\rho} V, \theta}$	$\frac{\text{STRUCT} \quad S \xRightarrow{\rho} V_s, \theta}{(\text{struct } S \ \text{end}) \xRightarrow{\rho} V_s, \theta}$	$\frac{\text{EMPTYSTRUCT}}{\varepsilon \xRightarrow{\rho} \{ \}, \langle \rangle}$
$\frac{\text{MODCLOSURE}}{\text{functor}(X : \mathcal{M})M \xRightarrow{\rho} \text{functor}(\rho)(X : \mathcal{M})M, \langle \rangle}$	$\frac{\text{MODCONSTR} \quad M \xRightarrow{\rho} V, \theta}{(M : \mathcal{M}) \xRightarrow{\rho} V, \theta}$		
$\frac{\text{MODBETA} \quad M \xRightarrow{\rho} \text{functor}(\rho')(X : \mathcal{M})M_f, \theta \quad M' \xRightarrow{\rho} V', \theta' \quad M_f \xRightarrow{\rho' + \{x_i \mapsto V'\}} V'', \theta''}{M(M') \xRightarrow{\rho} V'', \theta @ \theta' @ \theta''}$			
$\frac{\text{TYPEDECL} \quad S \xRightarrow{\rho} V_s, \theta}{(\text{type } t_i = \tau; S) \xRightarrow{\rho} V_s, \theta}$	$\frac{\text{MODULEDECL} \quad M \xRightarrow{\rho} V, \theta \quad S \xRightarrow{\rho + \{X_i \mapsto V\}} V_s, \theta'}{(\text{module } X_i = M; S) \xRightarrow{\rho} \{X_i \mapsto V\} + V_s, \theta @ \theta'}$		
$\frac{\text{VALDECL} \quad e \xRightarrow{\rho} v, \theta \quad S \xRightarrow{\rho + \{x_i \mapsto v\}} V_s, \theta'}{(\text{let } x_i = e; S) \xRightarrow{\rho} \{x_i \mapsto v\} + V_s, \theta'}$	$\frac{\text{PROGRAM} \quad S \xRightarrow{\rho} V_s, \theta}{\text{prog } S \ \text{end} \xRightarrow{\rho} V_s(\text{return}), \theta}$		

Fig. 53. Big step semantics – $e \xRightarrow{\rho} v, \theta$

B SEMANTICS PRESERVATION

We now prove [Theorem 4](#) which states that given an ELIOM_ε program P , the trace of its execution is the same as the concatenation of the traces of $\langle P \rangle_s$ and $\langle P \rangle_c$. More formally:

Theorem 6 (Compilation preserves semantics). Given sets of constants where converters are well-behaved, given an ELIOM_ε program P respecting the slicability hypothesis and such that $P \xrightarrow{\{\}} v, \theta$ then

$$\langle P \rangle_s \xrightarrow{\{\}}_{\text{ML}_s} () , \xi, \zeta, \theta_s \quad \langle P \rangle_c, \xi \xrightarrow{\{\} | \zeta \rightarrow \gamma}_{\text{ML}_c} v, \xi', \theta_c \quad \theta = \theta_s @ \theta_c$$

B.1 Hoisting

In [Section 5.2](#), we mentioned that a useful property of injections and fragments is that they can be partially lifted outside sections. This property can be used to simplify the simulation proofs. We consider the code transformation that hoists the content of injections out of fragments, client declarations and mixed functors in a way that preserve semantics. This transformation can be decomposed in two parts.

Injections. We decompose injections inside fragments and client declarations into simpler components. For example, the ELIOM_ε piece of code presented in [Example 24a](#) is decomposed in [Example 24b](#) by moving out the application of the converter and leaving only a call to the `serial` converter. All injections using a converter than is not `serial` nor `frag` can be decomposed in such a way.

Since injections can only be used on variables or constants and that no server bindings can be introduced inside a fragment, scoping is preserved. Furthermore, by definition of converters and their client and server components, this transformation preserves typing. It also preserves the dynamic semantics as long as the order of hoisting correspond to the order of evaluation. This can be seen by inspecting the reduction relation for server code under client contexts $\Rightarrow_{c/s}$. Finally, it trivially preserves the semantics of the compiled program since it corresponds exactly to how converters are decomposed during compilation.

<pre>let a = 1 + 2 in {{ 3 + int%a }}</pre> <p>(a) Fragment with injections</p>	<pre>let a = 1 + 2 in let a' = (int^s a) in {{ 3 + (int^c serial%a') }}</pre> <p>(b) Fragment with hoisted injections</p>
---	---

Example 24. Hoisting on fragments

This allows us to assume that reduction of server code in client context only uses variable lookup and never leads to any evaluation. In particular, this will avoid having to deal with the case of fragments being executed inside the reduction of another fragment (to see why this could happen, consider the case of a converter of type $\forall \alpha_c. (\text{unit} \rightarrow \{\alpha_c\}) \rightsquigarrow \alpha_c$).

In the rest of this section, we assume that reductions of the ELIOM_ε rule `FRAGMENT` are always of the following shape:

$$\frac{e \xrightarrow{\rho_c}_{c/s} \bar{e}, \varepsilon, \langle \rangle}{\{\{ e \}\} \xrightarrow{\rho_s}_s \mathbf{r}, (\text{bind } \mathbf{r} = \bar{e}), \langle \rangle} \quad \text{where } \bar{e} = e[f_i \% x_i \mapsto \rho_s(x_i)]_i \text{ and } f_i \in \{\text{serial}, \text{frag}\}$$

and that reductions of the ELIOM_ε rule CLIENTDECL are always of the following shape:

$$\frac{}{D_c \xrightarrow{\rho_s}_{c/s} \overline{D_c}, \varepsilon, \langle \rangle} \text{ where } \overline{D_c} = D_c[f_i \% x_i \mapsto \rho_s(x_i)]_i \text{ and } f_i \in \{\text{serial}, \text{frag}\}$$

Injections inside mixed modules. We also hoist injections completely out of mixed contexts to the outer englobing scope. For example in the functor presented in [Example 25a](#), we can lift the injection out of the functor, as show in [Example 25b](#). This is valid since injections can only reference content outside of the functor, by typing. Semantics is similarly preserved since injections inside functors are reduced immediately when encountering a functor, as per rule MODCLOSURE in [Figure 32](#).

This allows us to assume that the reduction of a mixed module will never lead to the reduction of an injection.

```

lets x = ...
modulem F(X : M) = struct
  letc y = f % x
end
(a) Mixed functor with injections

```

```

lets x = ...
letc y' = f % x
modulem F(X : M) = struct
  letc y = y'
end
(b) Mixed functor with hoisted injections

```

Example 25. Hoisting on mixed modules

B.2 Preliminaries

Let us start with some naming conventions. Identifiers with a hat, such as $\widehat{\gamma}$, are related to the compiled semantics. For example, while the server environment for the interpreted semantics is noted ρ_s , the environment for the execution of the target language ML_s is noted $\widehat{\rho}_s$. This naming convention is only for ease of reading and does not apply a formal relation between the objects with and without hats, unless indicated explicitly.

B.2.1 Remarks about global environments. Let us make some preliminary remarks about global environments in the ELIOM_ε client generated programs and in ML_c .

Given a global environment γ resulting of \Rightarrow_c , it contains only two kinds of references:

- Closure fragments, noted \mathbf{f} , which come from the execution of `bind env`. The associated value is always a environment (*i.e.*, a signature).
- Fragment values, noted \mathbf{r} , which come from the execution of `bind with`.

In the rest of this section, we consider that we can always decompose global environments γ in two parts: a fragment value environment γ_r containing all the references \mathbf{r} that were produced by `bind with` and a fragment closure environment γ_f containing only binding of the form $\{\mathbf{f} \mapsto \rho\}$ that were produced by `bind env`.

Similarly, given a global environment $\widehat{\gamma}$ used in ML_c . There are only three kind of references:

- Closure fragments, noted \mathbf{f} , which come from the slicing of syntactic fragments in the source program. The associated value is always a closure.
- Fragment values, noted \mathbf{r} come from the execution of fragments in the fragment queue.

- Injections, noted \mathbf{x} . The associated values must be serializable, and hence can only be references or constants in $Const_b$.

In the rest of this section, we consider that we can always decompose global compiled environment $\widehat{\gamma}$ into a fragment closure environment $\widehat{\gamma}_f$, a fragment value environment $\widehat{\gamma}_r$ and an injection environment ζ .

B.2.2 Client equivalence.

Definition 2 (Client values equivalence). Given v an ELIOM client value, v' an ML_c value and ζ an environment of references, v and v' are equivalent under ζ , noted $v \simeq_{\zeta}^c v'$, if and only if they are equals after substitution by ζ : $v[\zeta] = v'[\zeta]$.

We extend this notation to environments and traces.

Definition 3 (Global environment equivalence). We say that an $ELIOM_e$ global environment $\gamma = \gamma_f \cup \gamma_r$ and an ML_c global environment $\widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta$ are synchronized if and only if the following conditions hold.

- The reference environments are equivalent: $\gamma_r \simeq_{\zeta}^c \widehat{\gamma}_r$
- The domains of γ_f and $\widehat{\gamma}_f$ coincides, and:
 - For each \mathbf{f} in these environments such that $\gamma_f(\mathbf{f}) = \rho$ and that $\widehat{\gamma}_f(\mathbf{f}) = \lambda x_0 \dots x_n. \widehat{\rho}.e$, then the following property must hold.
We must have that $\rho \simeq_{\zeta}^c \widehat{\rho}$ and that for all $v_0, \dots, v_n, \widehat{v}_0, \dots, \widehat{v}_n$ such that for all $i, v_i \simeq_{\zeta}^c \widehat{v}_i$; then:

$$e[x_i \mapsto v_i]_i \xrightarrow{\rho | \gamma \rightarrow \gamma}_c v, \theta \implies (\lambda \bar{x}_i. \widehat{\rho}.e \widehat{v}_0 \dots \widehat{v}_n) \xrightarrow{|\widehat{\gamma} \rightarrow \widehat{\gamma}}_{ML_c} \widehat{v}, \widehat{\theta}$$

with $v \simeq_{\zeta}^c \widehat{v}$ and $\theta \simeq_{\zeta}^c \widehat{\theta}$

- For each \mathbf{F} in these environments such that $\gamma_f(\mathbf{F}) = \rho$ and that $\widehat{\gamma}_f(\mathbf{F}) = \text{functor}(\widehat{\rho})(Y_i : \mathcal{M}_i)_i \widehat{S}$, we have $\rho \simeq_{\zeta}^c \widehat{\rho}$.

Definition 4 (Fragment closure environment). We consider that $\widehat{\gamma}_f$ is a fragment closure environment for the $ELIOM_e$ server expression e_s , noted $FCE(\widehat{\gamma}_f, e_s)$, if for each $\{\mathbf{f} \mapsto \lambda \bar{x}_i. \widehat{\rho}.e'\}$ in $\widehat{\gamma}_f$, for each $\{\{e\}\}_f$ in $FRAGS(e_s)$ we have $e' = e[f_i \% x_i \mapsto x_i]_i$.

Definition 5 (Functor closure environment). We consider that $\widehat{\gamma}_f$ is a functor closure environment for the $ELIOM_e$ module expression M , noted $FCE(\widehat{\gamma}_f, M)$, if for each $\{\mathbf{F} \mapsto \text{functor}(\widehat{\rho})(Y_i : \mathcal{M}_i)_i \widehat{S}\}$ in $\widehat{\gamma}_f$, for each $(\text{struct } S \text{ end}_F)$ in M_s we have $\widehat{S} = \langle S \rangle_c$. Additionally, we require that $\widehat{\gamma}_f$ be a fragment closure environment for each expression contained in S .

In the rest of this section, we use the same notation for both properties. We extend this notation to server declarations, server values (by looking under closures) and server environments.

Lemma 1 (Reduction up to equivalence). Given $\rho, \widehat{\rho}, \gamma = \gamma_f \cup \gamma_r, \widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta, e$ and \widehat{e} such that:

$$\rho \simeq_{\zeta}^c \widehat{\rho} \qquad \gamma \simeq_{\zeta}^c \widehat{\gamma} \qquad e[\zeta] = \widehat{e}[\zeta] \qquad e \xrightarrow{\rho | \gamma \rightarrow \gamma}_c v, \theta$$

Then we have:

$$\widehat{e} \xrightarrow{\widehat{\rho} | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{ML_c} \widehat{v}, \widehat{\theta} \qquad v \simeq_{\zeta}^c \widehat{v} \qquad \theta \simeq_{\zeta}^c \widehat{\theta}$$

PROOF. The only difference between ELIOM_e client expressions and ML_c expressions are the presence of extra references for injections in ML_c . Indeed, syntactic injections have been removed either by the server execution or by compilation and `bind` constructs are only accessible at the module level. Since we assume that the original expression e and the compiled expression \widehat{e} are the same up to the injection environment ζ , we can trivially mimic the execution of e in \widehat{e} by induction. \square

B.2.3 Server equivalence.

Definition 6 (Server value equivalence). Given v an ELIOM server value, \widehat{v} an ML_s value. We say they are equivalent, noted $v \simeq^s \widehat{v}$ if and only if

$$v[\{\{ e_i \}\}_f \mapsto \text{fragment } \mathbf{f} \overline{x_{i,j}}]_i = \widehat{v} \quad \text{where } \overline{x_{i,j}} = \text{INJS}(e_i)$$

We extend this notation to environments and traces.

B.3 Server expressions and structures

We first look at server expressions and structures. By definition of the server reduction relation for ELIOM_e , the emitted program is a series of binds.

Lemma 2 (Server expressions are simulable). We consider an ELIOM_e server expression e ; the ELIOM_e environments ρ_s, ρ_c and $\gamma = \gamma_f \cup \gamma_r$; the target environment $\widehat{\rho}_s, \widehat{\rho}_c$ and $\widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta$.

If the expression e has valid server and client executions:

$$e \xrightarrow{\rho_s}_s v, \mu, \theta_s \quad \mu \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c \{\}, \theta_c$$

and the following invariants hold:

$$\widehat{\rho}_c \simeq_{\zeta}^c \rho_c \quad \widehat{\rho}_s \simeq^s \rho_s \quad \widehat{\gamma} \simeq^c \gamma \quad \text{FCE}(\widehat{\gamma}_f, e) \quad \text{FCE}(\widehat{\gamma}_f, \rho_s)$$

Then $\widehat{e} = e[\{\{ e_i \}\}_f \mapsto \text{fragment } \mathbf{f} \overline{x_{i,j}}]_i$ has an equivalent execution.

$$\frac{}{\widehat{e} \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{v}, \xi_{\bullet}, \{\}, \widehat{\theta}_s} \quad \frac{}{\text{exec } (\cdot), \xi_{\bullet} \cdot \text{end} \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \rightarrow \widehat{\gamma}'}_{\text{ML}_c} \varepsilon, [\cdot], \widehat{\theta}_c}$$

with the following invariants:

$$\widehat{\gamma}' \simeq^c \gamma' \quad \text{FCE}(\widehat{\gamma}'_f, v) \quad \widehat{v} \simeq^s v \quad \widehat{\theta}_s \simeq^s \theta_s \quad \widehat{\theta}_c \simeq_{\zeta}^c \theta_c$$

PROOF. We consider an expression e ; the ELIOM_e environments ρ_s, ρ_c and $\gamma = \gamma_f \cup \gamma_r$; the target environment $\widehat{\rho}_s, \widehat{\rho}_c$ and $\widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta$. such that

$$\widehat{\gamma} \simeq^c \gamma \quad \widehat{\rho}_c \simeq_{\zeta}^c \rho_c \quad \widehat{\rho}_s \simeq^s \rho_s \quad \text{FCE}(\widehat{\gamma}_f, e) \quad \text{FCE}(\widehat{\gamma}_f, \rho_s)$$

We will proceed by induction over the executions of e and μ . The only case of interest is when the server expression is a fragment.

- CASE $\{\{ e \}\}_f$.

We assume that the following executions hold:

$$\frac{\rho_s(x_i) = v_i}{\{\{ e \}\}_f \xrightarrow{\rho_s}_m \mathbf{r}, \text{bind } \mathbf{r} = \bar{e} \text{ with } \mathbf{f}, \langle \cdot \rangle} \quad \frac{\gamma(\mathbf{f}) = \rho \quad \bar{e} \xrightarrow{\rho | \gamma \rightarrow \gamma'}_c v_c, \theta_c}{(\text{bind } \mathbf{r} = \bar{e} \text{ with } \mathbf{f}) \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c \{\}, \theta_c}$$

where $\bar{e} = e[f_i \% x_i \mapsto \downarrow v_i]_i$ and $\gamma' = \gamma \cup \{\mathbf{r} \mapsto v_c\}$. We have \widehat{e} equal to `fragment f x1 . . . xn`.

We first consider the execution of \widehat{e} . We can easily construct the following execution.

$$\frac{\widehat{\rho}_s(x_i) = \widehat{v}_i}{\text{fragment } f \ x_1 \dots x_n \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \mathbf{r}, \{\mathbf{r} \mapsto f \ \downarrow \widehat{v}_1 \dots \downarrow \widehat{v}_n\}, \{\}, \langle \rangle}$$

By hypothesis, for each i , $v_i \simeq_{\widehat{\gamma}}^s \widehat{v}_i$. This gives us that $\downarrow v_i \simeq_{\widehat{\gamma}}^c \downarrow \widehat{v}_i$. We trivially have that $\mathbf{r} \simeq_{\widehat{\gamma}}^s \mathbf{r}$

Let us now look at the client execution. By client execution of μ , $\gamma(f) = \rho$. Since $\gamma \simeq_{\widehat{\zeta}}^c \widehat{\gamma}$, we have $\{f \mapsto \lambda x_0 \dots x_n. \widehat{\rho}.e'\} \in \widehat{\gamma}$ and $\rho \simeq_{\widehat{\zeta}}^c \widehat{\rho}$. Furthermore, since $FCE(\widehat{\gamma}_f, \{\{e\}\}_f)$, we know that that $e' = e[f_i \% x_i \mapsto x_i]_i$. We have by hypothesis that $\bar{e} \xrightarrow{\rho_c | \gamma \rightarrow \gamma}_c v_c, \theta_c$. Since $\bar{e} = e'[x_i \mapsto \downarrow v_i]$ and since for all i , $\downarrow v_i \simeq_{\widehat{\gamma}}^c \downarrow \widehat{v}_i$, we can use **Lemma 1** to build the following reduction:

$$\frac{\frac{\frac{e' \xrightarrow{\widehat{\rho}_c \cup \{x_i \mapsto \widehat{v}_i\}_i | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{v}, \widehat{\theta}_c}{\lambda x_1 \dots x_n. \widehat{\rho}.e' \ \downarrow \widehat{v}_1 \dots \downarrow \widehat{v}_n \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{v}, \widehat{\theta}_c}}{f \ \downarrow \widehat{v}_1 \dots \downarrow \widehat{v}_n \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{v}, \widehat{\theta}_c}}{\text{exec } (), \{\mathbf{r} \mapsto f \ \downarrow \widehat{v}_1 \dots \downarrow \widehat{v}_n\} \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \varepsilon, [], \widehat{\theta}_c}$$

Where $\widehat{\gamma}' = \widehat{\gamma} \cup \{\mathbf{r} \mapsto \widehat{v}\}$. By **Lemma 1**, we have that $v \simeq_{\widehat{\gamma}}^c \widehat{v}$ and $\theta_c \simeq_{\widehat{\gamma}}^c \widehat{\theta}_c$. The only part that is changed in γ' and $\widehat{\gamma}'$ is the fragment reference environment, hence we easily have that $\widehat{\gamma}' \simeq^c \gamma'$, which concludes. \blacksquare

• OTHER CASES.

In other cases, we first note that references manipulated inside server code can only fragment references \mathbf{r} . By hypothesis, the same references are considered before and after compilation. Since the fragment closure environment hypothesis ranges over all server expressions, including the one in closures, it is easy to preserve it during execution. The rest is a very simple induction. \blacksquare

Corollary 1 (Server module declarations are simulable). We consider an ELIOM_e server declaration D_s ; the ELIOM_e environments ρ_s, ρ_c and $\gamma = \gamma_f \cup \gamma_r$; the target environment $\widehat{\rho}_s, \widehat{\rho}_c$ and $\widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta$.

If the expression e has valid server and client executions:

$$D \xrightarrow{\rho_s}_s V, \mu, \theta_s \qquad \mu \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c \{\}, \theta_c$$

and the following invariants hold:

$$\widehat{\rho}_c \simeq_{\widehat{\zeta}}^c \rho_c \qquad \widehat{\rho}_s \simeq^s \rho_s \qquad \widehat{\gamma} \simeq^c \gamma \qquad FCE(\widehat{\gamma}_f, D) \qquad FCE(\widehat{\gamma}_f, \rho_s)$$

Then $\widehat{D} = D[\{\{e_i\}\}_f \mapsto \text{fragment } f \ \overline{x_{i,j}}]_i$ have an equivalent execution.

$$\frac{\widehat{D} \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{V}, \xi_{\bullet}, \{\}, \widehat{\theta}_s}{\text{exec } (), \xi_{\bullet} \text{ ++ end} \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \varepsilon, [], \widehat{\theta}_c}$$

with the following invariants:

$$\widehat{\gamma}' \simeq^c \gamma' \qquad FCE(\widehat{\gamma}'_f, V) \qquad \widehat{V} \simeq^s V \qquad \widehat{\theta}_s \simeq^s \theta_s \qquad \widehat{\theta}_c \simeq_{\widehat{\zeta}}^c \theta_c$$

B.4 Mixed structures

Lemma 3 (Structures are simulable). We consider a slicable structure S ; the ELIOM_ε environments ρ_s, ρ_c and $\gamma = \gamma_f \cup \gamma_r$; the target environment $\widehat{\rho}_s, \widehat{\rho}_c$ and $\widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta$.

If the structure has valid server and client executions:

$$S \xrightarrow{\rho_s}_m V_s, \mu, \theta_s \quad \mu \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c V_c, \theta_c$$

and the following invariants hold:

$$\widehat{\rho}_c \simeq_{\zeta}^c \rho_c \quad \widehat{\rho}_s \simeq^s \rho_s \quad \widehat{\gamma} \simeq^c \gamma \quad FCE(\widehat{\gamma}_f, S) \quad FCE(\widehat{\gamma}_f, \rho_s)$$

then for any ξ , the compiled structures have equivalent executions

$$\langle S \rangle_s \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{V}_s, \xi, \zeta, \widehat{\theta}_s \quad \langle S \rangle_c, \xi \cdot \text{++} \xi \xrightarrow{\widehat{\rho}_c | \zeta \cup \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{V}_c, \xi', \widehat{\theta}_c$$

with the following invariants:

$$\begin{array}{ccc} \widehat{\gamma}' \simeq^c \gamma' & \widehat{V}_s \simeq^s V_s & \widehat{\theta}_s \simeq^s \theta_s \\ FCE(\widehat{\gamma}'_f, V_s) & \widehat{V}_c \simeq_{\widehat{\gamma}'}^c V_c & \widehat{\theta}_c \simeq_{\widehat{\gamma}'}^c \theta_c \end{array}$$

PROOF. We consider a slicable structure S ; the ELIOM_ε environments ρ_s, ρ_c and $\gamma = \gamma_f \cup \gamma_r$; the target environment $\widehat{\rho}_s, \widehat{\rho}_c$ and $\widehat{\gamma} = \widehat{\gamma}_f \cup \widehat{\gamma}_r \cup \zeta$. such that

$$\widehat{\rho}_c \simeq_{\widehat{\gamma}}^c \rho_c \quad \widehat{\rho}_s \simeq^s \rho_s \quad \widehat{\gamma} \simeq^c \gamma \quad FCE(\widehat{\gamma}_f, S) \quad FCE(\widehat{\gamma}_f, \rho_s)$$

We will now proceed by induction over the execution of S .

- CASE $S = D_b; S'$ – BASE DECLARATION.

We assume that the following executions hold:

$$\frac{D_b \xrightarrow{\rho_s}_b V_s, \varepsilon, \theta_s \quad S' \xrightarrow{\rho_s + V_s}_m V'_s, \mu, \theta'_s}{D_b; S' \xrightarrow{\rho_s}_m V_s + V'_s, (D_b; \mu), \theta_s @ \theta'_s} \quad \frac{D_b \xrightarrow{\rho_c | \gamma \rightarrow \gamma}_c V_c, \theta_c \quad \mu \xrightarrow{\rho_c + V_c | \gamma \rightarrow \gamma'}_c V'_c, \theta'_c}{D_b; \mu \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c V_c + V'_c, \theta_c @ \theta'_c}$$

Let us consider the executions of D_b . By definition of base, it contains neither injections nor fragments. By **Proposition 5**, $D_b \xrightarrow{\rho}_b V_s, \varepsilon, \theta_s$ and $D_b \xrightarrow{\rho_c | \gamma \rightarrow \gamma}_c V_b, \theta_c$ both correspond to $D_b \xrightarrow{\rho_s}_b V_s, \varepsilon, \theta_s$ and $D_b \xrightarrow{\rho_c}_c V_c, \varepsilon, \theta_c$ respectively. By definition, base fragments can't be present, hence we also have $FCE(\widehat{\gamma}_f, V_s)$

Additionally, the compilation functions are the identity on base, which mean that $\langle D_b \rangle_s$ and $\langle D_b \rangle_c$ contains only ML constructs. The reduction relation over ML_s and ML_c coincide with the ML one on the ML fragment of the language. Hence, for any ξ , we have $\langle D_b \rangle_s \xrightarrow{\rho_s}_{\text{ML}_s} \widehat{V}_s, [], \{ \}, \widehat{\theta}_s$ and $\langle D_b \rangle_c, \xi \xrightarrow{\rho_c | \widehat{\gamma} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{V}_c, \xi, \widehat{\theta}_c$ with

$$\widehat{V}_s \simeq^s V_s \quad \widehat{V}_c \simeq_{\widehat{\gamma}}^c V_c \quad \widehat{\theta}_s \simeq^s \theta_s \quad \widehat{\theta}_c \simeq_{\widehat{\gamma}}^c \theta_c$$

Let us consider the execution of S' and μ . We easily have the following properties:

$$\widehat{\rho}_c + \widehat{V}_c \simeq_{\zeta}^c \rho_c + V_c \quad \widehat{\rho}_s + \widehat{V}_s \simeq^s \rho_s + V_s \quad \widehat{\gamma} \simeq^c \gamma \quad FCE(\widehat{\gamma}_f, S) \quad FCE(\widehat{\gamma}_f, \rho_s + V_s)$$

hence, by induction on the execution of S' and μ' , we have $\langle S' \rangle_s \xrightarrow{\widehat{\rho}_s + \widehat{V}_s}_{\text{ML}_s} \widehat{V}'_s, \xi_{\bullet}, \zeta_{\bullet}, \widehat{\theta}'_s$ and $\langle S' \rangle_c, \xi_{\bullet} + \xi \xrightarrow{\widehat{\rho}_c + \widehat{V}_c}$ for any ξ , with

$$\begin{array}{ccc} \widehat{\gamma}' \simeq^c \gamma' & \widehat{V}'_s \simeq^s V'_s & \widehat{\theta}_s \simeq^s \theta_s \\ \text{FCE}(\widehat{\gamma}'_f, V'_s) & \widehat{V}'_c \simeq^c V'_c & \widehat{\theta}'_c \simeq^c \theta'_c \end{array}$$

We can then build the following derivations:

$$\frac{\langle D_b \rangle_s \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{V}_s, [], \{ \}, \theta_s \quad \langle S' \rangle_s \xrightarrow{\widehat{\rho}_s + \widehat{V}_s}_{\text{ML}_s} \widehat{V}'_s, \xi_{\bullet}, \zeta_{\bullet}, \widehat{\theta}'_s}{\langle D_b; S' \rangle_s \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{V}_s + \widehat{V}'_s, \xi_{\bullet}, \zeta_{\bullet}, \widehat{\theta}_s @ \widehat{\theta}'_s}$$

$$\frac{\langle D_b \rangle_c, \xi_{\bullet} + \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \cup \xi_{\bullet} \rightarrow \widehat{\gamma} \cup \xi_{\bullet}}_{\text{ML}_c} \widehat{V}_c, \xi_{\bullet} + \xi, \widehat{\theta}_c \quad \langle S' \rangle_c, \xi_{\bullet} + \xi \xrightarrow{\widehat{\rho}_c + \widehat{V}_c | \widehat{\gamma} \cup \xi_{\bullet} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{V}'_c, \xi_{\bullet}, \widehat{\theta}'_c}{\langle D_b; S \rangle_c, \xi_{\bullet} + \xi \xrightarrow{\rho_c | \widehat{\gamma} \cup \xi_{\bullet} \rightarrow \widehat{\gamma}}_{\text{ML}_c} \widehat{V}_c + \widehat{V}'_c, \xi_{\bullet}, \widehat{\theta}_c @ \widehat{\theta}'_c}$$

and the following invariants are easily verified:

$$\begin{array}{ccc} \widehat{\gamma}' \simeq^c \gamma' & \widehat{V}_s + \widehat{V}'_s \simeq^s V_s + V'_s & \widehat{\theta}_s @ \widehat{\theta}'_s \simeq^s \theta_s @ \theta'_s \\ \text{FCE}(\widehat{\gamma}'_f, V_s + V'_s) & \widehat{V}_c + \widehat{V}'_c \simeq^c V_c + V'_c & \widehat{\theta}_c @ \widehat{\theta}'_c \simeq^c \theta_c @ \theta'_c \end{array} \quad \blacksquare$$

• CASE $S = D_s; S'$ – SERVER DECLARATION.

We assume that the following executions hold:

$$\frac{D_s \xrightarrow{\rho_s}_s V_s, \mu, \theta_s \quad S' \xrightarrow{\rho_s + V_s}_m V'_s, \mu', \theta'_s}{D_s; S' \xrightarrow{\rho_s}_m V_s + V'_s, (\text{bind env } \mathbf{f}_i)_i; \mu; \mu', \theta_s @ \theta'_s}$$

$$\frac{(\text{bind env } \mathbf{f}_i)_i \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c \{ \}, \langle \rangle \quad \mu \xrightarrow{\rho_c | \gamma' \rightarrow \gamma''}_c V_c, \theta_c \quad \mu' \xrightarrow{\rho_c | \gamma'' \rightarrow \gamma'''}_c V'_c, \theta'_c}{(\text{bind env } \mathbf{f}_i)_i; \mu; \mu' \xrightarrow{\rho_c | \gamma \rightarrow \gamma'''}_c V_c + V'_c, \theta_c @ \theta'_c}$$

Let us note $\{ \{ e_i \} \}_{\mathbf{f}_i}$ the fragments syntactically present in D_s . let us note $\{ \{ e_j \} \}_{\mathbf{f}_j}$ the fragments executed during the reduction of D_s and \mathbf{r}_j the associated fresh variables.

We have the following compilations:

$$\begin{array}{l} \langle D_s; S' \rangle_s = \langle D_s \rangle_s [\{ \{ e_i \} \}_{\mathbf{f}_i} \mapsto \text{fragment } \mathbf{f}_i \overline{x_{i,k}}]_i; \text{end} (); \langle S' \rangle_s \\ \langle D_s; S' \rangle_c = (\text{bind } \mathbf{f}_i = \lambda \overline{x_{i,k}}. e_i [f_{i,k} \% x_{i,k} \mapsto x_{i,k}]); \text{exec} (); \langle S' \rangle_c \end{array}$$

After hoisting, converters can only be the `serial` or `frag`. Its server and client parts are the identity, hence we simply omit them. We also note that $\langle D_s \rangle_s$ differs with D_s only on type annotations and type declarations which are ignored by reduction relations. We note $\widehat{D}_s = \langle D_s \rangle_s [\{ \{ e_i \} \} \mapsto \text{fragment } \mathbf{f}_i \overline{x_{i,k}}$

Let us consider the reduction of $(\text{bind } \mathbf{f}_i = \lambda \overline{x_{i,k}}. \widehat{e}_i)_i$. Let us note $e'_i = e_i [f_{i,j} \% x_{i,j} \mapsto x_{i,j}]$. For any queue ξ , we have the following reduction:

$$\frac{\lambda x_1 \dots x_m. e'_i, \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}_i \rightarrow \widehat{\gamma}_i}_{\text{ML}_c} \lambda \overline{x_{i,k}}. \widehat{\rho}_c. e'_i, \xi, \langle \rangle}{\forall i, \text{bind } \mathbf{f}_i = \lambda \overline{x_{i,k}}. e'_i, \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}_i \rightarrow \widehat{\gamma}_{i+1}}_{\text{ML}_c} \{ \}, \xi, \langle \rangle}$$

$$(\text{bind } \mathbf{f}_i = \lambda \overline{x_{i,k}}. e'_i)_i, \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}_i \rightarrow \widehat{\gamma}_{n+1}}_{\text{ML}_c} \{ \}, \xi, \langle \rangle$$

where $\widehat{\gamma}_1 = \widehat{\gamma}$ and $\widehat{\gamma}_{i+1} = \widehat{\gamma}_i \cup \{\mathbf{f}_i \mapsto \lambda x_1 \dots x_m. \widehat{\rho}_c. e'_i\}$. Let γ_{bind} be $\{\mathbf{f}_i \mapsto \lambda x_1 \dots x_m. \widehat{\rho}_c. e'_i\}_i$. We note $\widehat{\gamma}' = \widehat{\gamma}_{n+1} = \widehat{\gamma} \cup \gamma_{bind}$ and $\widehat{\gamma}'_f = \widehat{\gamma}_f \cup \gamma_{bind}$.

Since $(bind \ env \ \mathbf{f}_i)_i \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c \{\}, \langle \rangle$, we have $\gamma' = \gamma \cup \{\mathbf{f}_i \mapsto \rho_c\}_i$. and $\rho_c \simeq^c \widehat{\rho}_c$, we have that $\gamma' \simeq^c \widehat{\gamma}'$. Furthermore, given one of the \mathbf{f}_i in γ_{bind} , each fragment annotated with this \mathbf{f}_i syntactically appear in D_s by uniqueness of the annotation function. This also holds inside functors, since each \mathbf{f}_i will be prefixed by a unique module reference. Hence $FCE(\widehat{\gamma}'_f, D_s)$ and $FCE(\widehat{\gamma}'_f, \rho_s)$.

We now have all the ingredients to uses **Corollary 1** on the execution of D_s and μ . This gives us the following reductions:

$$\overline{\widehat{D}_s \xrightarrow{\widehat{\rho}_s}_{ML_s} \widehat{V}_s, \xi_\bullet, \{\}, \widehat{\theta}_s} \quad \overline{exec \ () , \xi_\bullet ++ end \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}' \rightarrow \widehat{\gamma}''}_{ML_c} \varepsilon, [\], \widehat{\theta}_c}$$

with the following invariants:

$$\widehat{\gamma}'' \simeq^c \gamma'' \quad FCE(\widehat{\gamma}''_f, V) \quad \widehat{V} \simeq^s V \quad \widehat{\theta}_s \simeq^s \theta_s \quad \widehat{\theta}_c \simeq^c \theta_c$$

We remark that $\zeta'' = \zeta$ since no injection took place during a server section and that $\widehat{\gamma}''_f = \widehat{\gamma}'_f$, by definition of the reduction for $exec$.

We now consider the execution of S' . The following invariants holds:

$$\widehat{\rho}_c \simeq^c \rho_c \quad \widehat{\rho}_s + \widehat{V}_s \simeq^s \rho_s + V_s \quad \widehat{\gamma}'' \simeq^c \gamma'' \quad FCE(\widehat{\gamma}''_f, S') \quad FCE(\widehat{\gamma}''_f, \rho'_s + V_s)$$

By induction on the execution of S' and μ' , we have $\langle S' \rangle_s \xrightarrow{\widehat{\rho}_s + \widehat{V}_s}_{ML_s} \widehat{V}'_s, \xi'_\bullet, \zeta'_\bullet, \widehat{\theta}'_c$ and $\langle S' \rangle_c, \xi'_\bullet ++ \xi' \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}'' \cup \zeta'_\bullet \rightarrow \widehat{\gamma}''' }_{ML_c}$ where

$$\begin{array}{ccc} \widehat{\gamma}''' \simeq^c \gamma''' & \widehat{V}'_s \simeq^s V'_s & \widehat{\theta}'_s \simeq^s \theta'_s \\ FCE(\widehat{\gamma}'''_f, V'_s) & \widehat{V}'_c \simeq^c V'_c & \widehat{\theta}'_c \simeq^c \theta'_c \end{array}$$

Finally, we can construct the following executions:

$$\begin{array}{c} \overline{\widehat{D}_s \xrightarrow{\widehat{\rho}_s}_{ML_s} \widehat{V}_s, \xi_\bullet, \{\}, \widehat{\theta}_s} \quad \langle S' \rangle_s \xrightarrow{\rho_s + V_s}_{ML_s} \widehat{V}'_s, \xi'_\bullet, \zeta'_\bullet, \widehat{\theta}'_s} \\ \overline{\widehat{D}_s; end \ () ; \langle S' \rangle_s \xrightarrow{\widehat{\rho}_s}_{ML_s} \widehat{V}_s + \widehat{V}'_s, \xi_\bullet ++ end ++ \xi'_\bullet, \zeta'_\bullet, \widehat{\theta}_s @ \widehat{\theta}'_s} \\ \overline{\widehat{\mu}; exec \ () ; \langle S' \rangle_c, \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}'' \cup \zeta'_\bullet \rightarrow \widehat{\gamma}''' }_{ML_c} V'_c, \xi', \widehat{\theta}'_c} \\ \overline{\widehat{\mu}; exec \ () ; \langle S' \rangle_c, \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}'' \cup \zeta'_\bullet \rightarrow \widehat{\gamma}''' }_{ML_c} \widehat{V}'_c, \xi', \widehat{\theta}_c @ \widehat{\theta}'_c} \end{array}$$

where $\xi = \xi_\bullet ++ end ++ \xi'_\bullet ++ \xi'$ and $\widehat{\mu} = (bind \ \mathbf{f}_i = \lambda \bar{x}_i. e'_i)_i$. We verify the following invariants:

$$\begin{array}{ccc} \widehat{\gamma}''' \simeq^c \gamma''' & \widehat{V}_s + \widehat{V}'_s \simeq^s V_s + V'_s & \widehat{\theta}_s + \widehat{\theta}'_s \simeq^s \theta_s + \theta'_s \\ FCE(\widehat{\gamma}'''_f, V_s + V'_s) & \widehat{V}'_c \simeq^c V'_c & \widehat{\theta}_c + \widehat{\theta}'_c \simeq^c \theta_c + \theta'_c \end{array} \quad \blacksquare$$

• CASE $S = D_c; S'$ – CLIENT DECLARATION.

We assume that the following executions hold:

$$\frac{D_c \xrightarrow{\rho_s}_{c/s} \overline{D_c}, \varepsilon, \langle \rangle \quad S' \xrightarrow{\rho_s}_m V'_s, \mu', \theta'_s}{D_c; S' \xrightarrow{\rho_s}_m V'_s, (\overline{D_c}; \mu'), \theta'_s} \quad \frac{\overline{D_c} \xrightarrow{\rho_c | \gamma \rightarrow \gamma}_c V_c, \theta_c \quad \mu' \xrightarrow{\rho_c + V_c | \gamma \rightarrow \gamma'}_c V'_c, \theta'_c}{\overline{D_c}; \mu' \xrightarrow{\rho_c | \gamma \rightarrow \gamma'}_c V_c + V'_c, \theta_c @ \theta'_c}$$

Let us note $f_i \% x_i$ the injections in D_c and \mathbf{x}_i the associated fresh variables. Since hoisting has been applied, all the f_i are either serial or frag. Furthermore, no fragments are executed due to the execution of injections and $\overline{D_c} = D_c[f_i \% x_i \mapsto \rho_s(x_i)]_i$.

We have the following compilations:

$$\begin{aligned} \langle D_c; S \rangle_s &= (\text{injection } \mathbf{x}_i \ x_i;)_i \text{ end } (); \langle S \rangle_s \\ \langle D_c; S \rangle_c &= \text{exec } (); D_c[f_i \% x_i \mapsto \mathbf{x}_i]_i; \langle S \rangle_c \end{aligned}$$

In the rest of this proof, we note $\widehat{D}_c = D_c[f_i \% x_i \mapsto \mathbf{x}_i]_i$.

We consider the server reduction $D_c \Rightarrow_{c/s} \overline{D_c}$. We know that $\overline{D_c} = D_c[f_i \% x_i \mapsto \downarrow \rho_s(x_i)]_i$. Let us note $v_i = \rho_s(x_i)$. We can build the following ML_s reduction:

$$\frac{\widehat{\rho}_s(x_i) = \widehat{v}_i}{\forall i. \text{injection } \mathbf{x}_i \ x_i \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \varepsilon, [], \{\mathbf{x}_i \mapsto \downarrow \widehat{v}_i\}, \langle \rangle} (\text{injection } \mathbf{x}_i \ x_i;)_i; \text{end } () \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \varepsilon, \text{end}, \{\mathbf{x}_i \mapsto \downarrow \widehat{v}_i\}_i, \langle \rangle$$

Since $\widehat{\rho}_s \simeq^s \rho_s$, we also have that $\widehat{v}_i \simeq^s v_i$ and $\downarrow \widehat{v}_i \simeq^c \downarrow v_i$ for each i . We note $\zeta_\bullet = \{\mathbf{x}_i \mapsto \downarrow \widehat{v}_i\}_i$. By definition of the slicing relation, the \mathbf{x}_i are fresh, hence they are not bound in $\widehat{\gamma}$. We can thus construct the global environment $\widehat{\gamma}' = \widehat{\gamma} \cup \zeta_\bullet$. Since we only extend the part with injection references, we still have that $\gamma \simeq^c \widehat{\gamma}'$.

We now consider the client reduction $\overline{D_c} \xrightarrow{\rho_c | \gamma \rightarrow \gamma}_c V_c, \theta_c$. We know that $\overline{D_c}$ is equal to $D_c[f_i \% x_i \mapsto \downarrow v_i]_i$, hence the reduction tree contains for each i a reduction $\downarrow v_i \xrightarrow{|\gamma \rightarrow \gamma}_c \downarrow v_i, \langle \rangle$. To obtain a reduction of $\widehat{D}_c = D_c[f_i \% x_i \mapsto \mathbf{x}_i]_i$, we simply substitute each of these subreduction by one of the form $\mathbf{x}_i, \xi \xrightarrow{|\widehat{\gamma}' \rightarrow \widehat{\gamma}'}_{\text{ML}_c} \downarrow \widehat{v}_c, \xi, \langle \rangle$. for any queue ξ . By [Lemma 1](#), we can build the following reduction:

$$\widehat{D}_c, \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}' \rightarrow \widehat{\gamma}'}_{\text{ML}_c} \widehat{V}_c, \xi, \widehat{\theta}_c$$

where $\widehat{V}_c \simeq^c_{\xi'} V_c$ and $\widehat{\theta}_c \simeq^c_{\xi'} \theta_c$, for any queue ξ .

We now consider the execution of S' . We have the following properties:

$$\widehat{\rho}_c + \widehat{V}_c \simeq^c_{\xi'} \rho_c + V_c \quad \widehat{\rho}_s \simeq^s \rho_s \quad \widehat{\gamma}' \simeq^c \gamma \quad FCE(\widehat{\gamma}'_f, S') \quad FCE(\widehat{\gamma}'_f, \rho_s)$$

By induction on the execution of S' and μ' , we have $\langle S' \rangle_c \xrightarrow{\widehat{\rho}_c + \widehat{V}_c}_{\text{ML}_s} \widehat{V}'_s, \xi'', \zeta''_\bullet, \widehat{\theta}'_c$ and $\langle S' \rangle_s, \xi'_\bullet + \xi' \xrightarrow{\widehat{\rho}_s | \zeta''_\bullet \cup \widehat{\gamma}' \rightarrow \widehat{\gamma}'}$ where

$$\begin{aligned} \widehat{\gamma}'' \simeq^c \gamma' & & \widehat{V}'_s \simeq^s_{\xi''} V'_s & & \widehat{\theta}'_c \simeq^s_{\xi''} \theta'_s \\ FCE(\widehat{\gamma}''_f, V'_s) & & \widehat{V}'_c \simeq^c_{\xi''} V'_c & & \widehat{\theta}'_c \simeq^c_{\xi''} \theta'_c \end{aligned}$$

Finally, we can build the following derivations:

$$\begin{array}{c}
\frac{}{\widehat{V}_i, \text{injection } \mathbf{x}_i \ x_i \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \varepsilon, [], \zeta_\bullet, \langle \rangle} \quad \frac{}{\langle S' \rangle_s \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{V}'_s, \xi'_\bullet, \zeta'_\bullet, \widehat{\theta}'_s} \\
\frac{}{(\text{injection } \mathbf{x}_i \ (f_i^s \ x_i);)_i \ \text{end } (); \langle S' \rangle_s \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \widehat{V}'_s, \text{end} + \xi'_\bullet, \zeta_\bullet \cup \zeta'_\bullet, \widehat{\theta}'_s} \\
\frac{\widehat{D}_c, \xi'_\bullet + \xi' \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}' \cup \zeta'_\bullet \rightarrow \widehat{\gamma}' \cup \zeta'_\bullet}_{\text{ML}_c} \widehat{V}_c, \xi'_\bullet + \xi', \widehat{\theta}'_c \quad \langle S' \rangle_c, \xi'_\bullet + \xi' \xrightarrow{\widehat{\rho}_c + \widehat{V}_c | \widehat{\gamma}' \cup \zeta'_\bullet \rightarrow \widehat{\gamma}''}_{\text{ML}_c} \widehat{V}'_c, \xi', \widehat{\theta}'_c}{\widehat{D}_c; \langle S' \rangle_c, \xi'_\bullet + \xi' \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}' \cup \zeta'_\bullet \rightarrow \widehat{\gamma}''}_{\text{ML}_c} \widehat{V}_c + \widehat{V}'_c, \xi', \widehat{\theta}_c @ \widehat{\theta}'_c} \\
\frac{}{\text{exec } (); \widehat{D}_c; \langle S' \rangle_c, \text{end} + \xi'_\bullet + \xi' \xrightarrow{\widehat{\rho}_c | \widehat{\gamma}' \cup \zeta'_\bullet \cup \zeta'_\bullet \rightarrow \widehat{\gamma}''}_{\text{ML}_c} \widehat{V}_c + \widehat{V}'_c, \xi', \widehat{\theta}_c @ \widehat{\theta}'_c}
\end{array}$$

We verify the following invariants:

$$\begin{array}{ccc}
\widehat{\gamma}'' \simeq^c \gamma' & \widehat{V}_s \simeq^s V_s & \widehat{\theta}_s \simeq^s \theta_s \\
FCE(\widehat{\gamma}'_f, V_s + V'_s) & \widehat{V}_c + \widehat{V}'_c \simeq^c_{\widehat{\gamma}''} V_c + V'_c & \widehat{\theta}_c + \widehat{\theta}'_c \simeq^c_{\widehat{\gamma}''} \theta_c + \theta'_c \quad \blacksquare
\end{array}$$

• CASE $\text{module}_m X = M; S'$ – DECLARATION OF A MIXED MODULE.

We assume that the following executions hold:

$$\begin{array}{c}
\frac{M \xrightarrow{\rho_s}_m V_s, M^c, \mu, \theta_s}{\text{module}_m X = M \xrightarrow{\rho_s}_m \{X \mapsto V_s\}, \text{module } X = M^c; \mu, \theta_s \quad S' \xrightarrow{\rho_s + \{X \mapsto V_s\}}_m V'_s, \mu', \theta'_s} \\
\frac{}{\text{module}_m X = M; S' \xrightarrow{\rho_s}_m \{X \mapsto V_s\} + V'_s, (\mu; \text{module } X = M^c; \mu'), \theta_s @ \theta'_s} \\
\frac{M^c \xrightarrow{\rho_c | \gamma' \rightarrow \gamma''}_c V_c, \theta'_c}{\mu \xrightarrow{\rho_c | \gamma' \rightarrow \gamma'}_c \{ \}, \theta_c \quad \text{module } X = M^c \xrightarrow{\rho_c | \gamma' \rightarrow \gamma''}_c \{X \mapsto V_c\}, \theta'_c \quad \mu' \xrightarrow{\rho_c + \{X \mapsto V_c\} | \gamma'' \rightarrow \gamma''}_c V'_c, \theta'_c} \\
\frac{}{\mu; \text{module } X = M^c; \mu' \xrightarrow{\rho_c | \gamma' \rightarrow \gamma''}_c \{X \mapsto V_c\} + V'_c, \theta_c @ \theta'_c @ \theta'_c}
\end{array}$$

Let us assume that we can build the following reductions

$$\begin{array}{c}
\frac{}{\langle \text{module}_m X = M \rangle_s \xrightarrow{\widehat{\rho}_s}_{\text{ML}_s} \{X \mapsto \widehat{V}_s\}, \xi_\bullet, \zeta, \widehat{\theta}_s} \\
\frac{}{\langle \text{module}_m X = M \rangle_c, \xi_\bullet + \xi \xrightarrow{\widehat{\rho}_c | \zeta \cup \widehat{\gamma} \rightarrow \widehat{\gamma}'}_{\text{ML}_c} \{X \mapsto \widehat{V}_c\}, \xi, \widehat{\theta}_c}
\end{array}$$

for any ξ , and that the following invariants hold:

$$\begin{array}{ccc}
\widehat{\gamma}' \simeq^c \gamma'' & \widehat{V}_s \simeq^s V_s & \widehat{\theta}_s \simeq^s \theta_s \\
FCE(\widehat{\gamma}'_f, V_s) & \widehat{V}_c \simeq^c_{\widehat{\gamma}'} V_c & \widehat{\theta}_c \simeq^c_{\widehat{\gamma}'} \theta_c + \theta'_c
\end{array}$$

By induction on the execution of S' and μ' , we can build the following reduction: $\langle S' \rangle_s \xrightarrow{\widehat{\rho}_s + \{X \mapsto \widehat{V}_s\}}_{\text{ML}_s} V'_s, \xi'_\bullet, \zeta'_\bullet$

and $\langle S' \rangle_c, \xi'_\bullet + \xi \xrightarrow{\widehat{\rho}_c + \{X \mapsto \widehat{V}_c\} | \zeta' \cup \widehat{\gamma}' \rightarrow \widehat{\gamma}''}_{\text{ML}_c} V'_c, \xi, \theta'_c$, which allows us to conclude.

To build the compiled reduction, we will operate by case analysis over M .

- SUBCASE $M = \text{struct } S \text{ end}_X$ – DECLARATION OF A MIXED STRUCTURE.

We have $\mu = \text{bind } X = (\text{struct } \mu_0 \text{ end})$ and $M^c = X$ with the following reductions:

$$\frac{\frac{S \xrightarrow{\rho_s} {}_m V_s, \mu, \theta_s}{\text{struct } S \text{ end} \xrightarrow{\rho_s} {}_m V_s + \{\text{Dyn} \mapsto X\}, X, \mu, \theta_s}}{\frac{\mu_0 \xrightarrow{\rho_c | \gamma \rightarrow \gamma'} {}_c V_c, \theta_c \quad \frac{\gamma''(X) = V_c}{X \xrightarrow{\rho_c | \gamma'' \rightarrow \gamma''} {}_c V_c, \langle \rangle}}{(\text{bind } X = \text{struct } \mu_0 \text{ end}; \text{ module } X = X) \xrightarrow{\rho_c | \gamma \rightarrow \gamma''} {}_c, \theta_c}}$$

where $\gamma'' = \gamma' \cup \{X \mapsto V_c\}$.

We have the following compilations:

$$\begin{aligned} \langle \text{module}_m X = \text{struct } S \text{ end}_X \rangle_s &= \begin{array}{l} \text{module } X = \text{struct} \\ \text{module } \text{Dyn} = X; \\ \langle S \rangle_s \\ \text{end} \end{array} \\ \langle \text{module}_m X = \text{struct } S \text{ end}_X \rangle_c &= \begin{array}{l} \text{bind}_m X = \text{struct } \langle S \rangle_c \text{ end}; \\ \text{module } X = X; \end{array} \end{aligned}$$

By induction on the execution of S and μ , we have $\langle S \rangle_s \xrightarrow{\widehat{\rho}_s} {}_{ML_s} \widehat{V}_s, \xi_\bullet, \zeta_\bullet, \widehat{\theta}_s$ and $\langle S \rangle_c, \xi_\bullet + \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \cup \xi_\bullet \rightarrow \widehat{\gamma}'} {}_{ML_c} \widehat{V}_c, \xi, \widehat{\theta}_c$ with the following invariants:

$$\begin{array}{lll} \widehat{\gamma}' \simeq^c \gamma' & \widehat{V}_s \simeq^s V_s & \widehat{\theta}_s \simeq^s \theta_s \\ FCE(\widehat{\gamma}'_f, V_s) & \widehat{V}_c \simeq^c_{\widehat{\gamma}'} V_c & \widehat{\theta}_c \simeq^c_{\widehat{\gamma}'} \theta_c \end{array}$$

We can then build the following executions:

$$\begin{array}{l} \frac{\langle S \rangle_s \xrightarrow{\widehat{\rho}_s} {}_{ML_s} \widehat{V}_s, \xi_\bullet, \zeta_\bullet, \widehat{\theta}_s}{\langle \text{module}_m X = \text{struct } S \text{ end}_X \rangle_s \xrightarrow{\widehat{\rho}_s} {}_{ML_s} \{X \mapsto \{\text{Dyn} \mapsto X\} + \widehat{V}_s\}, \xi_\bullet, \zeta_\bullet, \widehat{\theta}_s} \\ \frac{\langle S \rangle_c, \xi_\bullet + \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \cup \xi_\bullet \rightarrow \widehat{\gamma}'} {}_{ML_c} \widehat{V}_c, \xi, \widehat{\theta}_c}{\text{module } X = \text{struct } \langle S \rangle_c \text{ end}, \xi_\bullet + \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \cup \xi_\bullet \rightarrow \widehat{\gamma}'} {}_{ML_c} \{X \mapsto \widehat{V}_c\}, \xi, \widehat{\theta}_c} \\ \langle \text{module}_m X = \text{struct } S \text{ end} \rangle_c, \xi_\bullet + \xi \xrightarrow{\widehat{\rho}_c | \widehat{\gamma} \cup \xi_\bullet \rightarrow \widehat{\gamma}'} {}_{ML_c} \{X \mapsto \widehat{V}_c\}, \xi, \widehat{\theta}_c \end{array}$$

Where $\widehat{\gamma}'' = \widehat{\gamma}' \cup \{X \mapsto \widehat{V}_c\}$. We verify the following invariants:

$$\begin{array}{lll} \widehat{\gamma}'' \simeq^c \gamma' & \widehat{V}_s + \{\text{Dyn} \mapsto X\} \simeq^s V_s + \{\text{Dyn} \mapsto X\} & \widehat{\theta}_s \simeq^s \theta_s \\ FCE(\widehat{\gamma}''_f, V_s) & \widehat{V}_c \simeq^c_{\widehat{\gamma}'} V_c & \widehat{\theta}_c \simeq^c_{\widehat{\gamma}'} \theta_c \end{array}$$

which concludes. ■

- SUBCASE $M = \text{functor}(X_i : M_i); \text{struct } S \text{ end}_F$ – DECLARATION OF A MIXED FUNCTOR.

In this case, we have the client program $\mu = \text{bind env F}$ and the module expression $M^c = \text{functor}(X_i : \mathcal{M}_i)_i \text{struct } S|_c \text{ end}$. The following reductions hold:

$$\frac{\left(\begin{array}{l} \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S \\ \text{end}_F \end{array} \right)}{\left(\begin{array}{l} \text{bind env F} \\ \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S|_c \\ \text{end} \end{array} \right), \langle \rangle} \xRightarrow{\rho_s} \{F \mapsto V_s\}, \left(\begin{array}{l} \text{bind env F} \\ \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S|_c \\ \text{end} \end{array} \right), \langle \rangle$$

$$\frac{\left(\begin{array}{l} \text{bind env F} \\ \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S|_c \\ \text{end} \end{array} \right)}{\left(\begin{array}{l} \text{bind env F} \\ \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S|_c \\ \text{end} \end{array} \right)} \xRightarrow{\rho_c | \gamma \rightarrow \gamma'} \{F \mapsto V_c\}, \langle \rangle$$

Where $\gamma' = \gamma \cup \{F \mapsto \rho_c\}$ and the following values:

$$V_s = \text{functor}(\rho_s)(X_i : \mathcal{M}_i)_i \text{struct } S \text{ end}$$

$$V_c = \text{functor}(\rho_c)(X_i : \mathcal{M}_i)_i \text{struct } S|_c \text{ end}$$

We recall that by hoisting, the body of the functors contains no injection, hence we don't need to evaluate server code in the client part.

We have the following compilations:

$$\left\langle \begin{array}{l} \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S \\ \text{end}_F \end{array} \right\rangle_s = \begin{array}{l} \text{module } F(X_i : \langle \mathcal{M}_i \rangle_s)_i = \text{struct} \\ \text{module Dyn} = \text{fragment}_m F(X_i.\text{Dyn})_i; \\ \langle S \rangle_s \\ \text{end} \end{array}$$

$$\left\langle \begin{array}{l} \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S \\ \text{end}_F \end{array} \right\rangle_c = \begin{array}{l} \text{bind}_m F(X_i : \langle \mathcal{M}_i \rangle_c)_i = \text{struct } \langle S \rangle_c \text{ end}; \\ \text{module } F(X_i : \langle \mathcal{M}_i \rangle_c)_i = \text{struct } S|_c \text{ end}; \end{array}$$

We trivially have the following execution:

$$\left\langle \begin{array}{l} \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S \\ \text{end}_F \end{array} \right\rangle_s \xRightarrow{\widehat{\rho}_s} \{F \mapsto \widehat{V}_s\}, \xi, \bullet, \{\}, \langle \rangle$$

$$\left\langle \begin{array}{l} \text{module}_m F(X_i : \mathcal{M}_i)_i = \text{struct} \\ S \\ \text{end}_F \end{array} \right\rangle_c, \xi \xRightarrow{\widehat{\rho}_c | \widehat{\gamma} \rightarrow \widehat{\gamma}'} \{F \mapsto \widehat{V}_c\}, \xi, \langle \rangle$$

Where $\gamma' = \gamma \cup \{F \mapsto \widehat{V}_F\}$ with the following values:

$$\widehat{V}_s = \text{functor}(\widehat{\rho}_s)(X_i : \mathcal{M}_i)_i \text{struct } \langle S \rangle_s \text{ end}$$

$$\widehat{V}_c = \text{functor}(\widehat{\rho}_c)(X_i : \mathcal{M}_i)_i \text{struct } S|_c \text{ end}$$

$$\widehat{V}_F = \text{functor}(\widehat{\rho}_c)(X_i : \mathcal{M}_i)_i \text{struct } \langle S \rangle_c \text{ end}$$

We now need to show that the invariants still hold. We easily have that $\widehat{V}_c \simeq_c^c V_c$. By definition of equivalence over mixed functors, we have $\widehat{V}_s \simeq^s V_s$. Indeed, the body of the functor in \widehat{V}_s is the server compilation of the body of the mixed functor V_s and the captured environments corresponds. Finally, we have that the body of \widehat{V}_F is the client compilation of the body of the mixed functors and that the capture environment corresponds to $\gamma'(F)$. Thus we get that $FCE(\widehat{\gamma}'_f, V_s)$. By definition of the annotation function, the reference F could not have appeared on a previously executed structure, hence we still have that $FCE(\widehat{\gamma}'_f, \widehat{\rho}_s)$.

Hence, all the following invariants are respected, which concludes.

$$\begin{array}{ll} \widehat{\gamma}' \simeq^c \gamma' & \widehat{V}_s \simeq^s V_s \\ FCE(\widehat{\gamma}'_f, V_s) & \widehat{V}_c \simeq_c^c V_c \end{array} \quad \blacksquare$$

Otherwise, M is a module expression. By definition of slicability, M does not syntactically contain any structure. In the general case, we should proceed by induction over module expressions. We will simply present the case of a mixed functor application where the functor returns a mixed structure.

We consider $M = F(X_1) \dots (X_n)$. We have $M^c = F(X_1) \dots (X_n)$ with the following executions:

$$\begin{array}{l} F \xrightarrow{\rho_s} {}_m \text{functor} (\rho'_s)(Y_i : \mathcal{M}_i)_i \text{struct } S \text{ end}_F, \varepsilon, \langle \rangle \\ X_i \xrightarrow{\rho_s} {}_m V_i^s, \varepsilon, \langle \rangle \quad V_i^s(\text{Dyn}) = \mathbf{R}_i \quad \mathbf{R} \text{ fresh} \quad S[\mathbf{f}_i \mapsto \mathbf{R}.\mathbf{f}_i]_i \xrightarrow{\rho'_s + \{Y_i \mapsto V_i^s\}} {}_m V, \mu_0, \theta \\ \hline F(X_1) \dots (X_n) \xrightarrow{\rho_s} {}_m V_s + \{\text{Dyn} \mapsto \mathbf{R}\}, \mu = \left(\begin{array}{l} \text{bind } \mathbf{R} = \text{struct} \\ \quad (\text{module } Y_i = \mathbf{R}_i;)_i \\ \quad \mu_0 \\ \text{end with } F \end{array} \right), \theta \\ \hline \text{module}_m X = F(X_1) \dots (X_n) \xrightarrow{\rho} {}_m V_s, (\mu; \text{module } X = F(X_1) \dots (X_n)), \theta_s \\ \hline \gamma(F) = \rho_F \quad \gamma(\mathbf{R}_i) = V_i^c \quad \mu_0 \xrightarrow{\rho_F + \{Y_i \mapsto V_i^c\} \mid \gamma \rightarrow \gamma'} {}_c V_c, \theta_c \\ \left(\begin{array}{l} \text{bind } \mathbf{R} = \text{struct} \\ \quad (\text{module } Y_i = \mathbf{R}_i;)_i \\ \quad \mu_0 \\ \text{end with } F \end{array} \right) \xrightarrow{\rho_c \mid \gamma \rightarrow \gamma' \cup \{\mathbf{R} \mapsto V_c\}} {}_c \varepsilon, \theta_c \\ \hline F \xrightarrow{\rho_s \mid \dots \rightarrow \dots} {}_c \text{functor} (\rho'_c)(Y_i : \mathcal{M}_i)_i \text{struct } S_c \text{ end}, \langle \rangle \\ X_i \xrightarrow{\rho_c \mid \dots \rightarrow \dots} {}_c V_i, \langle \rangle \quad S_c \xrightarrow{\rho'_c + \{Y_i \mapsto V_i\} \mid \gamma' \cup \{\mathbf{R} \mapsto V_c\} \rightarrow \gamma''} {}_c V_c', \theta'_c \\ \hline \text{module } X = F(X_1) \dots (X_n) \xrightarrow{\rho_c \mid \gamma' \cup \{\mathbf{R} \mapsto V_c\} \rightarrow \gamma''} {}_c \{X \mapsto V_c'\}, \theta'_c \end{array}$$

We note V_F the value of F in γ_s , which is $\text{functor}(\rho'_s)(Y_i : \mathcal{M}_i)_i \text{struct } S \text{ end}$.

We have the following compilations:

$$\begin{array}{l} \langle \text{module}_m X = F(X_1) \dots (X_n) \rangle_s = \text{module } X = F(X_1) \dots (X_n); \\ \quad \text{end } (); \\ \langle \text{module}_m X = F(X_1) \dots (X_n) \rangle_c = \text{module } X = F(X_1) \dots (X_n); \\ \quad \text{exec } (); \end{array}$$

Let us now look at the execution of the server application $F(X_1) \dots (X_n)$. By hypothesis, V_F is a mixed functor. By equivalence, we know that $\widehat{\rho}_s(F) = \widehat{V}_F$ where $\widehat{V}_F \simeq^s V_F$. By definition of the equivalence on server values, \widehat{V}_F is of the following shape:

$$\widehat{V}_F = \left(\begin{array}{l} \text{functor}(\widehat{\rho}'_s)(Y_i : \mathcal{M}_i)_i \text{struct} \\ \text{module Dyn} = \text{fragment}_m \mathbf{F}(Y_i.\text{Dyn})_i; \\ \langle S \rangle_s \\ \text{end} \end{array} \right)$$

where $\widehat{\rho}'_s \simeq^s \rho'_s$. For each i we note $\widehat{V}_i^s = \widehat{\rho}_s(X_i)$. By equivalence, we have $\widehat{V}_i^s \simeq^s V_i^s$.

Furthermore, since $\widehat{\gamma} \simeq^c \gamma$ and by hypothesis, \mathbf{F} is also bound in $\widehat{\gamma}$. We note V_F the corresponding value. Since $FCE(\widehat{\gamma}, V_c)$ (via ρ_s), then \widehat{V}_F is of the following shape:

$$\widehat{V}_F = \left(\begin{array}{l} \text{functor}(\widehat{\rho}_F)(Y_i : \mathcal{M}_i)_i \text{struct} \\ \langle S \rangle_c \\ \text{end} \end{array} \right)$$

where $\widehat{\rho}_F \simeq^c \rho_F = \gamma(\mathbf{F})$. Additionally, for each i we note $V_i^c = \widehat{\gamma}(\mathbf{R}_i)$. By equivalence, we have $\widehat{V}_i^c \simeq^s V_i^c$.

We can now proceed by induction on S and μ_0 in the environment $\widehat{\gamma}$, $\widehat{\rho}_c \cup \{Y_i \mapsto \mathbf{R}_i\}$ and $\widehat{\rho}_s \cup \{\text{Dyn} \mapsto \mathbf{R}\} \cup \{Y_i \mapsto \widehat{V}_i^s\}_i$. We obtain the following reductions:

$$\langle S \rangle_s \xrightarrow{\widehat{\rho}_s \cup \{Y_i \mapsto \widehat{V}_i^s\}_i} \text{ML}_s \xi_\bullet, \zeta_\bullet, \widehat{\theta}_s, \text{ and } \langle S \rangle_c, \xi_\bullet + \xi \xrightarrow{\widehat{\rho}_c \cup \{Y_i \mapsto \mathbf{R}_i\}_i \mid \zeta_\bullet \cup \widehat{\gamma} \rightarrow \widehat{\gamma}'} \text{ML}_c \widehat{V}_c, \xi, \widehat{\theta}_c$$

with the usual invariants. We can now build the following executions:

$$\begin{array}{l} \text{fragment}_m \mathbf{F}(Y_i.\text{Dyn})_i \xrightarrow{\widehat{\rho}_s \cup \{Y_i \mapsto \widehat{V}_i^s\}_i} \text{ML}_s \mathbf{R}, \xi_{\mathbf{R}}, \varepsilon, \langle \rangle \\ \widehat{\rho}_s(F) = \widehat{V}_f \quad \widehat{\rho}_s(X_i) = \widehat{V}_i^s \quad \langle S \rangle_s \xrightarrow{\widehat{\rho}_s \cup \{\text{Dyn} \mapsto \mathbf{R}\} \cup \{Y_i \mapsto \widehat{V}_i^s\}_i} \text{ML}_s \xi_\bullet, \zeta_\bullet, \widehat{\theta}_s \\ \text{module } X = F(X_1) \dots (X_n); \text{end } () \xrightarrow{\widehat{\rho}_s} \text{ML}_s \{X \mapsto \widehat{V}_s\}, \xi_{\mathbf{R}} + \xi_\bullet + \text{end}, \zeta_\bullet, \widehat{\theta}_s \\ \widehat{\gamma}(\mathbf{F}) = \widehat{V}_F \quad \langle S \rangle_c, \xi_\bullet + \text{end} + \xi \xrightarrow{\widehat{\rho}_F \cup \{Y_i \mapsto \mathbf{R}_i\}_i \mid \zeta_\bullet \cup \widehat{\gamma} \rightarrow \widehat{\gamma}'} \text{ML}_c \widehat{V}_c, \text{end} + \xi, \widehat{\theta}_c \\ \text{exec } (), \xi_{\mathbf{R}} + \xi_\bullet + \text{end} + \xi \xrightarrow{\widehat{\rho}_c \mid \zeta_\bullet \cup \widehat{\gamma} \rightarrow \widehat{\gamma}' \cup \{\mathbf{R} \mapsto \widehat{V}_c\}} \text{ML}_c \varepsilon, \xi, \widehat{\theta}_c \end{array}$$

where $\xi_{\mathbf{R}} = \{\mathbf{R} \mapsto \mathbf{F}(\mathbf{R}_1) \dots (\mathbf{R}_n)\}$. We respect the following invariants:

$$\begin{array}{lll} \widehat{\gamma}' \simeq^c \gamma' & \widehat{V}_s \simeq^s V_s & \widehat{\theta}_s \simeq^s \theta_s \\ FCE(\widehat{\gamma}'_f, V_s) & \widehat{V}_c \simeq^c_{\widehat{\gamma}'} V_c & \widehat{\theta}_c \simeq^c_{\widehat{\gamma}'} \theta_c \end{array}$$

Let us now consider the client application. Since $\rho_c \simeq^c \widehat{\rho}_c$, we have that the body of the functor F is equivalent. We can thus build the following reduction:

$$\begin{array}{l} \widehat{\rho}_c(F) = \text{functor}(\widehat{\rho}'_c)(Y_i : \mathcal{M}_i)_i \text{struct } \widehat{S}_c \text{end} \\ \widehat{\rho}_c(X_i) = \widehat{V}_i^c \quad \widehat{S}_c \xrightarrow{\widehat{\rho}'_c + \{Y_i \mapsto \widehat{V}_i^c\}_i \mid \widehat{\gamma}' \cup \{\mathbf{R} \mapsto \widehat{V}_c\} \rightarrow \widehat{\gamma}''} \text{ML}_c \widehat{V}'_c, \widehat{\theta}'_c \\ \text{module } X = F(X_1) \dots (X_n), \xi \xrightarrow{\widehat{\rho}_c \mid \widehat{\gamma}' \cup \{\mathbf{R} \mapsto \widehat{V}_c\} \rightarrow \widehat{\gamma}''} \text{ML}_c \{X \mapsto \widehat{V}'_c\}, \xi, \widehat{\theta}'_c \end{array}$$

By equivalence of F and X_i in ρ_c and $\widehat{\rho}_c$, we have that $\widehat{\gamma}' \simeq^c \gamma'$, $\widehat{V}'_c \simeq^c V'_c$ and $\widehat{\theta}_c \simeq^c \theta_c$.

We already built the reduction for the compiled server program. We can now build the compiled client program:

$$\frac{\begin{array}{l} \text{exec } (), \{\mathbf{R} \mapsto \mathbf{F}(\mathbf{R}_1) \dots (\mathbf{R}_n)\} + \xi_{\bullet} \text{++end++}\xi \xrightarrow{\widehat{\rho}_c \mid \xi_{\bullet} \cup \widehat{\gamma}' \rightarrow \widehat{\gamma}' \cup \{\mathbf{R} \mapsto \widehat{V}'_c\}}_{\text{ML}_c} \varepsilon, \xi, \widehat{\theta}_c \\ \text{module } X = F(X_1) \dots (X_n), \xi \xrightarrow{\widehat{\rho}_c \mid \widehat{\gamma}' \cup \{\mathbf{R} \mapsto \widehat{V}'_c\} \rightarrow \widehat{\gamma}''}_{\text{ML}_c} \{X \mapsto \widehat{V}'_c\}, \xi, \widehat{\theta}'_c \end{array}}{\langle \text{module}_m X = F(X_1) \dots (X_n) \rangle_c, \xi_{\mathbf{R}} + \xi_{\bullet} \text{++end++}\xi \xrightarrow{\widehat{\rho}_c \mid \xi_{\bullet} \cup \widehat{\gamma}' \rightarrow \widehat{\gamma}''}_{\text{ML}_c} \{X \mapsto \widehat{V}'_c\}, \xi, \widehat{\theta}_c @ \widehat{\theta}'_c}$$

where the invariants still hold. This concludes. ■

□

B.5 Proof of the main theorem

Finally, we prove [Theorem 4](#). This is a direct consequence of [Lemma 3](#).

PROOF OF THEOREM 4. We have that $P \xrightarrow{\{\}} v, \theta$. By definition of an ELIOM_ε program execution, we can decompose this rule as following:

$$\frac{P \xrightarrow{\{\}}_m(), \mu, \theta_s \quad \mu \xrightarrow{\{\} \mid \varepsilon \rightarrow \gamma}_c v, \theta_c}{P \xrightarrow{\{\}} v, \theta_s @ \theta_c}$$

We trivially have the following invariants:

$$\{\} \simeq^c_{\{\}} \{\} \quad \{\} \simeq^s \{\} \quad \widehat{\gamma}' \simeq^c \gamma \quad \text{FCE}(\{\}, P) \quad \text{FCE}(\{\}, \{\})$$

which allow us to apply [Lemma 3](#) and conclude. □