

Abstract types in distributed systems
A partial translation of my Ph.D. dissertation

Gilles Peskine

12 June 2008

Abstract

Consider a network of nodes running ML programs that exchange data. How can data which has an abstract type on one node be accepted on another node? A safe approach is to treat abstract types as distinct whenever they are defined on different nodes. However this is too restrictive in practice, for example in the common case where an abstract type enforces a semantic invariant.

The main contributions of this thesis are threefold: I define a notion of hash of an abstract type, whereby abstract types that have the same hash are deemed compatible; I give an operational semantics for a module system that preserves types, including abstract types; I also propose a new, more general module system that is well-suited to distributed applications.

The hash of an abstract type must reflect its intended semantics, which is often not apparent from the program's code. In practice, two modules have the same hash if they have the same code. Compound modules are compatible when they are built from compatible components.

Existing operational semantics for ML modules lose information as they erase abstraction boundaries. I use coloured brackets to track the visibility of abstract types. I study two calculi equipped with brackets, a simply-typed lambda-calculus and a rich ML module calculus.

I use singleton signatures to keep track of not only type but also code sharing, so that module equivalence is defined at arbitrary signatures. A simple effect system limits type constraint to a statically checkable fragment, while permitting both applicative and generative functors. I discuss static and dynamic forms of module sealing.

Contents

Introduction	7
IV TOPHAT: a module calculus suited to distributed environments	11
IV.1 Introduction	11
IV.2 A module calculus $[\mathcal{B}]$	12
IV.2.1 Fundamental constructs	12
IV.2.2 About the base language	12
IV.2.3 Formal description of the core language	12
IV.2.3.1 Syntax	12
IV.2.3.2 Variables	13
IV.2.3.3 Environments	13
IV.2.4 Typing	14
IV.2.4.1 Introduction	14
IV.2.4.2 $\Gamma \vdash \text{ok}$ Environment corrections	14
IV.2.4.3 $\Gamma \vdash T \text{ok}$ Type correctness	14
IV.2.4.4 $\Gamma \vdash E : T$ Expression typing	15
IV.2.4.5 $\langle T \rangle, \text{Type}$ Type fields	15
IV.2.5 Run-time	15
IV.2.5.1 $E \longrightarrow E'$ Expression reduction	16
IV.3 Singletons $[\mathcal{S}]$	17
IV.3.1 Motivation	17
IV.3.1.1 Abstract types, concrete types	17
IV.3.1.2 Type sharing	17
IV.3.1.3 Value singletons	17
IV.3.1.4 Higher-order singletons	18
IV.3.1.5 A practical example	18
IV.3.2 Properties	20
IV.3.3 Typing rules	20
IV.3.3.1 $\Gamma \vdash T <: T' ; \dots$ Subtyping	21
IV.3.3.2 $S(E)$ Singletons	21
IV.3.3.3 $\Gamma \vdash E : T ; \Gamma \vdash T_1 <: T_2$ Expression typing	21
IV.3.3.4 $\Gamma \vdash T \equiv T' ; \Gamma \vdash E \equiv E'$ Convertibility equivalences	22
IV.3.3.5 $\Gamma \vdash T \longrightarrow T'$ Type conversion	22
IV.3.3.6 $\Gamma \vdash E \longrightarrow E'$ Expression conversion	23
IV.3.3.7 Extensionality	24
IV.4 Sealing $[\mathcal{E}]$	24
IV.4.1 Sealing	24
IV.4.2 An effect system	24
IV.4.2.1 Introduction	24
IV.4.2.2 Purity	25
IV.4.2.3 Projectibility, separability and comparability	25
IV.4.3 Formal presentation	26

IV.4.3.1	Syntax		26
IV.4.3.2	$E \longrightarrow E'$	Run-time	27
IV.4.3.3	$\Gamma \vdash \dots$	Typing: correction, equivalences, subtyping	27
IV.4.3.4	$\Gamma \vdash E :^\gamma T$	Expression typing	28
IV.4.4	Applicativity		30
IV.4.4.1	Applicative functors		30
IV.4.4.2	Static sealing: formalisation	[W]	31
IV.4.4.3	Equivalences in the presence of static sealing		32
IV.4.4.4	Other forms of sealing		33
IV.4.4.5	Mutual encodings of static and dynamic sealing		35
IV.4.4.6	On applicativity through functor sealing		36
IV.5	Colours and brackets	[C]	36
IV.5.1	Module identities		37
IV.5.1.1	Nonce generation		37
IV.5.1.2	Lexes		38
IV.5.1.3	From sealing to brackets		38
IV.5.1.4	Abstract types		39
IV.5.1.5	Selfification		40
IV.5.2	Colors		41
IV.5.2.1	Colouring		41
IV.5.2.2	Semantics of a type and dependencies of a nonce		43
IV.5.2.3	Variables in colours		45
IV.5.2.4	Absolute brackets, additive brackets		46
IV.5.3	Polymorphism		47
IV.5.3.1	Coloration of a type		47
IV.5.3.2	Kinds		47
IV.5.3.3	Brackets and function application; polymorphic functions		48
IV.5.3.4	Polymorphic types and values		49
IV.5.3.5	Colour fusion		50
IV.5.3.6	Generative functors		51
IV.5.4	Evaluation		52
IV.5.4.1	Syntax		52
IV.5.4.2	Values and abstract components		53
IV.5.4.3	$B \vdash E \longrightarrow_c B' \vdash E'$	Reduction	54
IV.5.5	Typing		56
IV.5.5.1	$B; \Gamma \vdash_c \text{ok}$	Environment formation	57
IV.5.5.2	$B; \Gamma \vdash_c T : K$	Type kinding	57
IV.5.5.3	$B; \Gamma \vdash_c c_0 \text{ transparent}$	Colour transparency	58
IV.5.5.4	$B; \Gamma \vdash_c A \triangleright E : T ; \dots$	Module components	58
IV.5.5.5	$B; \Gamma \vdash_c E :^\gamma T$	Coloration of expressions	59
IV.5.5.6	$B; \Gamma \vdash_c E \longrightarrow E'$	Conversion and coloured brackets	60
IV.6	Dynamic typing and distributed programs	[D]	60
IV.6.1	Dynamic typing		60
IV.6.2	Formalisation		61
IV.6.2.1	Syntax		61
IV.6.2.2	Reduction		61
IV.6.2.3	Typing		62
IV.6.3	Communication inter-machines		62
IV.6.3.1	Introduction		62
IV.6.3.2	Communication and colours		63
IV.6.3.3	Universals		63
IV.6.3.4	Nonce sharing		63
IV.6.3.5	Static sealing and hashes		64
IV.7	Conclusion		65

V Conclusion	67
V.1 Summary	67
V.2 Related work	67
V.2.1 Theoretic considerations	67
V.2.2 Programming languages	67
V.2.3 Acute and HashCaml	67
V.2.4 Alice ML	67
V.3 Future work	68
V.3.1 Improvements to the theory	68
V.3.1.1 Stratification	68
V.3.1.2 One or two language levels?	68
V.3.1.3 Effect analysis	68
V.3.1.4 Colours and brackets	68
V.3.1.5 Decidability of type-checking	68
V.3.1.6 Parametricity	68
V.3.2 Supplementary features	68
V.3.2.1 Field names and width subsignaturing	68
V.3.2.2 Towards a programming language	68
V.3.2.3 Generic programming	68
V.3.2.4 Security	69
V.3.3 Implementation	69
V.3.3.1 Hash computation	69
V.3.3.2 Typing TOPHAT	69
V.3.3.3 Integration into Objective Caml: the module system	69
V.3.4 Applications of dynamic typing	69
V.3.4.1 The JoCaml name server	69
A Formal definition of TOPHAT	71
Bibliography	81
Index	85

Introduction

Objective

The objective of the present dissertation is to extend an ML-like language to adapt it to distributed systems. Specifically, we are interested in the requirements that the distributed nature of the environment imposes on the type system — we will not concern ourselves with other aspects such as concurrent execution and fault tolerance.

Consider two machines A and B , each executing a program. At some point in time, A and B start exchanging data. The central question of this dissertation is, how can we make sure that A and B agree on the semantics of the exchanged data?

A network link carries sequences of bits (usually arranged in bytes). When A sends data to B , the data is encoded as a sequence of bits. This operation is known as **marshaling** (the words **pickling** and **serialisation** are synonyms). Upon reception of the bit sequence, B must perform the opposite operation (known as unmarshaling, unpickling or deserialisation). Many languages provide a standard representation of data as strings: s-expressions in Lisp, **Marshal** library in Objective Caml [L⁺], **Pickle** library in Standard ML [PSL], **Serializable** interface in Java [Sun]... Several standards (ASN.1, XML) specify language-independent string encodings of data for communication. While the exact set of supported data shapes varies greatly, serialisation libraries and data representation standards usually specify at least how to encode numbers (n-bit integers, little- or big-endian, decimal notation...), strings (character sets and encodings: ASCII, Unicode, ...), sequences of such...

Marshaling data entails transforming it to an unambiguous sequence of bits. Unmarshaling consists of two parts: the bit sequence must be transformed back into a workable representation of the data, and one must verify that the resulting data has the expected type or shape. For example, if the program running on B expects a number, and the program on machine A sends the string "foo", the error must be detected. The usual approach in ML-like languages is to detect such errors as soon as possible, which is as soon as the program has been written (during the type-checking phase of compilation). It seems natural in an ML-like language to express the unmarshal-time check as a type constraint; but how can this constraint be imposed?

According to the ML approach, the error must be detected when compiling the program on A or B . Thus the program on A would declare a communication channel of type **string** (on which one may send "foo"), and the program on B would declare a channel of type **int** (on which only numbers may be received). But this only delays the problem, since the fact that A and B disagree on the type of their shared communication channel cannot be detected until A and B start communicating.

This observation leads us to desire a run-time type-check, specifically a type-check *when establishing a communication channel between programs that have not yet communicated*. (Once the programs have communicated, such a check is no longer necessary, since the programs may now have agreed on the types of future communications. For example JoCaml [MM01] has a static type system [FLMR97]; however two independently started JoCaml programs that wish to communicate

must get a shared channel via a “name server”, which is currently not well-typed.)

Although ML is designed to be statically typed, and most compilers erase types to save memory during execution, there are ways of checking whether a value has a certain type at run-time. However existing systems do not manage abstract types correctly, allowing only types with a predefined structure to be shared between separate programs.

One solution is to forbid values of abstract types to be marshaled. Another is to require the author of the abstract type to provide marshaling and unmarshaling functions. This however does not solve our problem: a serialisation format can usually be deduced automatically from the representation of the type, but this does not fully solve the problem of checking whether the type of the sent data is the type that is expected at the point of reception. Herein lies the gist of the matter: when are two abstract types the same?

There are two main intuitions to the nature of an abstract type. One point of view states that an abstract type is *hidden*. It has an *implementation*, which is a “concrete” type (the implementation may make use of other abstract types, but these can be traced through in turn all the way to built-in types). Hence an abstract type is a concrete type — but we do not know which. Another view is that an abstract type is a *new, fresh* type, distinct from any other type (in particular it is distinct from any concrete type, and it is distinct from its implementation type, in that one may not convert freely between the two).

When are two hidden types the same? One prerequisite that comes to mind immediately is that the implementation types must be the same. But this condition is neither necessary nor sufficient. One may wish to consider two hidden types as the same when their implementations have identical behaviour, even if their code differs. Conversely, just because the implementations match exactly does not mean that the types can be matched freely — for example a `Euro` type and a `Dollar` type may have the same implementation, yet should definitely not be compatible. Type abstraction can play multiple rôles, and usages may differ in terms of ideal degree of compatibility.

When are two fresh types the same? The simplest answer is “when they were created in the same operation”. This approach has often been refined by proposing language constructs that may or may not create fresh types. In ML, control of type freshness is given to the *module* language, which we shall therefore study.

General outline of the dissertation

Chapter I presents the basic concepts upon which the dissertation is based. We first study the notion of abstraction, its uses and how to express it. In ML-like languages, abstraction arises via the module language, and we highlight some points of its rich history. We also study how to add dynamic type-checking to a statically typed language.

Chapter II develops a notion of *imprint*. The imprint of a software component identifies the abstraction that it provides. We examine many sample programs in order to decide how much compatibility is desired in various conditions, and we discuss how to compute imprints so that two components have the same imprint if and only if they are supposed to be compatible.

Chapter III presents a simple first language equipped with imprints, the HAT language. This language extends the simply-typed lambda-calculus with simple modules. We keep track of abstraction domains throughout program execution using *coloured brackets*. The language also includes dynamically typed communication primitives that use imprints to test the equality of abstract types.

Chapter IV describes a new *module system* for ML which is suitable for distributed programs, the TOPHAT language. This language includes central concepts in module calculi, such as functors and sealing. Singleton types with no signature restriction allow for the expression of code equalities as well as type equalities, generalising the usual notion of type sharing. We show how to express

different kinds of sealing, depending on the expected level of generativity. Like HAT, TOPHAT uses imprints to perform run-time type equality tests involving abstract types, and coloured brackets to preserve abstraction barriers during program execution.

We conclude with a survey of related work and future work perspectives.

Appendix A summarises the formal definition of the TOPHAT language introduced in chapter IV. Appendix B contains a proof of the soundness of TOPHAT.

A note about code snippets

We usually present code snippets in Objective Caml syntax. We do not expect the reader to know the fine points of the language, and will in particular explain any subtlety concerning the semantics of modules. When features that Objective Caml does not have are illustrated, we use Objective Caml-like syntax augmented as desired and describe the intended semantics in the text. Readers used to Standard ML may wish to consult a correspondance table between the two dialects [Ros].

Chapter IV

TOPHAT: a module calculus suited to distributed environments

IV.1 Introduction

Objectives

The purpose of this chapter is to present a module system that combines the usual features found in ML-like languages with a flexible management of abstract types that, as in HAT, is suited to distributed programs.

This module system is described as a typed lambda-calculus, for which we provide typing rules and a type-preserving small-step operational semantics. The system described herein purports to be a theoretic model, not a full-blown programming language, although our design choices will be motivated by practical concerns. As such, it lacks some practical features that might be thought of as syntactic sugar. We will also for the most part delay implementation considerations until section V.3.3.

Existing module systems already span a wide range of features and style. With respect to expressivity, our aim is to cover the features that we think are fundamental to our specific objective of coping well with abstract types in a distributed environment. Style-wise, we have tried to provide a compositional approach, where each aspect of the language is embodied in a separate language construct that can be easily understood on its own.

Vocabulary and notations

We will endeavour to distinguish between the words “expression”, meaning a specific language category (expressions typically have types, and can be evaluated), and “term”, which denotes an element of any syntactic category (expression, type, module, environment, etc.). We will usually denote by \mathfrak{N} or \sqsupset a term of unknown syntactic category.

Let \mathfrak{N} be any term, x a variable and E an expression. We will write the substitution of E for x in \mathfrak{N} as $\{x \leftarrow E\}\mathfrak{N}$. Similarly, we will write $\{X \leftarrow M\}\mathfrak{N}$ for the substitution of module expression M for module variable X in \mathfrak{N} .

Outline

We will present the TOPHAT language incrementally. Each of the following five steps refines or extends the previous language.

- ⓑ We will start with a basic module system, including only module-building constructs that would be sufficient for our purpose in the absence of types. This system is simple but lacks expressivity as far as types are concerned.
- Ⓢ We will add singleton types (which generalise singleton kinds) to the basic system, in order to keep track of equalities between types. The resulting system will adequately model modules with no type abstraction.
- ⓔ We will then add a sealing construct to the language in order to permit making types abstract. We will see that sealing makes the language impure, and will equip our type system with a suitable effect system.
- ⓐ The previous system can express type abstraction at the source level, but abstraction is lost when the program is evaluated. We will therefore provide a way of keeping track of abstraction boundaries during program evaluation, in the form of module identities and coloured brackets.
- ⓓ We will finally be able to equip our language with dynamic typing constructs that behave reasonably in a distributed setting.

System \mathcal{D} constitutes the full TOPHAT language¹.

At each stage, we will motivate the features to be introduced with examples, and we will examine how these features can be used in programs. We will discuss the choices we made when designing the theory presented here. We will then state precisely the semantics of the language we define, in the form of typing rules (the static semantics) and small-step reduction rules (the dynamic semantics).

IV.2 A module calculus B

The present section presents the core of a module description language. This core, which we call \mathcal{B} , builds on two essential features: aggregates of values and types, called *structures*; and parametric modules, called *functors*.

IV.2.1 Fundamental constructs

[Sorry, this fragment has not been translated yet.]

IV.2.2 About the base language

[Sorry, this fragment has not been translated yet.]

IV.2.3 Formal description of the core language

IV.2.3.1 Syntax

We can now formally state the syntax and semantics of our core language. We limit ourselves to the features mentioned so far, and delay singletons until section IV.3.

Since we have decided to unify the module and expression languages, objects formerly noted E and objects formerly noted M now belong to the same world, and shall be noted E and called

¹*Total Or Partial Hashed Abstract Types*, in which the words “total” and “partial” refer to total and partial functors, also known as applicative and generative functors.

expressions. Similarly we will write T rather than S and speak of **types**. Nonetheless some expressions will intuitively be seen as modules, and their types as signatures.

$E ::=$	expression or module	$\mathsf{T} ::=$	type or signature
$x \mid y \mid t \mid \dots$	variables		
$()$	unit value	UNIT	unit
$\mathsf{false} \mid \mathsf{true}$	boolean (generically bv)	BOOL	booleans
$0 \mid 1 \mid \dots$	integer (generically \underline{n})	INT	integers
$\langle \mathsf{T} \rangle$	type field	TYPE	abstract type field
(E_1, E_2)	pair	$\mathsf{Type} E$	projection from a type field
$\pi_i E$	projection ($i \in \{1, 2\}$)	$\Sigma x : T_1. T_2$	dependent sum
$\lambda x : T. E$	lambda-abstraction	$\Pi x : T_0. T_1$	dependent product
$E_1 E_2$	application		
$\mathsf{let} x = E_0 \mathsf{in} E : T$	local binding		

We will use the following abbreviations.

$T_1 * T_2 := \Sigma x : T_1. T_2$ product type

$T_1 \rightarrow T_2 := \Pi x : T_1. T_2$ arrow (function) type

In the definitions of $T_1 * T_2$ and $T_1 \rightarrow T_2$, x is a fresh variable, i.e., a variable that is not free in T_2 .

IV.2.3.2 Variables

We use standard definitions of **free variables**, **bound occurrences**, **alpha-conversion** and **substitutions**. A **closed term** is one with no free variable.

We write $\mathsf{fv} \mathfrak{N}$ for the set of free variables of \mathfrak{N} . We write $\{x \leftarrow E\} \mathfrak{N}$ for the substitution of E for x in \mathfrak{N} .

We will systematically work *up to alpha-conversion*, i.e., any term that we write down will formally denote its equivalence class modulo alpha-conversion. Any typing or reduction step may rename variables. For example, if a typing or reduction rule requires more than one instantiation of a metavariable, each instantiation may use different representative for bound variables. This follows the tradition of the *Barendregt variable convention*, which allows for substantially clearer exposition. We will generally not mention the omnipresent possibility of alpha-conversion; syntax descriptions will mention the binding structure in the text.

IV.2.3.3 Environments

An **environment** is a finite list of (variable, type) pairs. We write nil for the empty environment, $x : T$ for an environment of length 1, and use “,” for concatenation, which we treat as associative. For example, an environment binding three variables will usually be written as $x : T, y : T', z : T''$; other ways of writing the same environment are $((x : T, y : T'), z : T'')$ and $(x : T, y : T', z : T'')$ and $((\mathsf{nil}, x : T), y : T', z : T'')$.

Environments are built from the following grammar:

$\Gamma ::=$	environnement
nil	empty
$\Gamma, x : T$	binding of the variable x

Alternatively, environments may be seen as objects of the form $x_1 : T_1, \dots, x_k : T_k$ with $k \in \mathbb{N}$.

The **domain** of an environnement $\Gamma = x_1 : T_1, \dots, x_k : T_k$ is the set of variables $\{x_1, \dots, x_k\}$. It is written **dom** Γ .

An environment binds the variables of its domain, and they are as usual subject to alpha-conversion. Writing (Γ, Γ') supposes that Γ and Γ' have disjoint domains; alpha-conversion must be performed if necessary. In a concatenation (Γ, Γ') , the variables in the domain of Γ bind in Γ' .

Note that our environments must be ordered since we have dependent types. Thus $(x : \text{INT}, y : \text{Typ } x)$ is a well-formed environment, whereas $(y : \text{Typ } x, x : \text{INT})$ is not (it could have been written as $(y : \text{Typ } x, z : \text{INT})$ after renaming the bound variable to z , with the remaining occurrence of x being free).

IV.2.4 Typing

IV.2.4.1 Introduction

We consider a correct program (fragment) to be an expression E associated with a type T such that E has the type T . When E contains free variables, these must be assigned a type through an environment.

We will manipulate several forms of typing judgements, which will always be **local judgements**, of the form $\Gamma \vdash J$ (in system \mathcal{B} — later local judgements will bear more annotations). System \mathcal{B} has three forms of right-hand side for a local typing judgement. .

$\mathcal{J} ::=$ **typing judgement**
 $\Gamma \vdash J$ local judgement

$J ::=$ **local judgement right-hand side**
 ok environment correction
 $T \text{ ok}$ correction of the T
 $E : T$ expression typing

We present typing rule under the usual presentation as a deduction rules.

IV.2.4.2 $\Gamma \vdash \text{ok}$ Environment corrections

Environments are built from left to right, binding by binding. Each type assigned to a variable must be valid in the environment that precedes the binding under consideration. Note that variables bound by an environment are automatically distinct as per our alpha-conversion convention.

$$\frac{}{\text{nil} \vdash \text{ok}} (\mathcal{B}/\text{envok.nil}) \qquad \frac{\Gamma \vdash T \text{ ok}}{\Gamma, x : T \vdash \text{ok}} (\mathcal{B}/\text{envok.x})$$

IV.2.4.3 $\Gamma \vdash T \text{ ok}$ Type correctness

The correctness rules for types are standard: for base types, we require that the environment be well-formed, and for constructed types, we require each part to be well-formed (treating dependencies properly).

$$\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \text{UNIT ok}} (\mathcal{B}/\text{tok.base.unit}) \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \text{BOOL ok}} (\mathcal{B}/\text{tok.base.bool}) \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \text{INT ok}} (\mathcal{B}/\text{tok.base.int})$$

$$\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \text{TYPE ok}} (\mathcal{B}/\text{tok.type})$$

$$\frac{\Gamma \vdash T' \text{ ok} \quad \Gamma, x : T' \vdash T'' \text{ ok}}{\Gamma \vdash \Sigma x : T'. T'' \text{ ok}} (\mathcal{B}/\text{tok.pair}) \qquad \frac{\Gamma \vdash T' \text{ ok} \quad \Gamma, x : T' \vdash T'' \text{ ok}}{\Gamma \vdash \Pi x : T'. T'' \text{ ok}} (\mathcal{B}/\text{tok.fun})$$

IV.2.4.4 $\boxed{\Gamma \vdash E : T}$ Expression typing

Constants Basic constants have their respective type in a correct environment.

$$\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash () : \text{UNIT}} \quad (\mathcal{B}/\text{et.base.unit}) \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \text{bv} : \text{BOOL}} \quad (\mathcal{B}/\text{et.base.bool}) \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \underline{n} : \text{INT}} \quad (\mathcal{B}/\text{et.base.int})$$

Variables Variables have the type stated in the environment.

$$\frac{\Gamma \vdash \text{ok} \quad \text{when } x : T \in \Gamma}{\Gamma \vdash x : T} \quad (\mathcal{B}/\text{et.x})$$

Pairs Although our syntax allows dependent sums, system \mathcal{B} is too restricted to take advantage of them (this defect will be remedied in system \mathcal{S}). We can only give pairs an ordinary pair type.

$$\frac{\Gamma \vdash E_1 : T_1 \quad \Gamma \vdash E_2 : T_2}{\Gamma \vdash (E_1, E_2) : T_1 * T_2} \quad (\mathcal{B}/\text{et.pair}) \qquad \frac{\Gamma \vdash E : T_1 * T_2}{\Gamma \vdash \pi_1 E : T_1} \quad (\mathcal{B}/\text{et.proj.1}) \qquad \frac{\Gamma \vdash E : T_1 * T_2}{\Gamma \vdash \pi_2 E : T_2} \quad (\mathcal{B}/\text{et.proj.2})$$

Functions We state classical rules for typing functions (or functors) and application, keeping in mind that we have dependent types. Note that in order to type the application of a function that has a dependent type, the occurrences of the formal parameter x must be replaced by the actual argument E_0 inside the result type E . Thus an arbitrary expression can appear in a type where a simple variable formally was. This illustrates the difficulty of restricting the presence of expressions in types to certain syntactic categories.

$$\frac{\Gamma, x : T_0 \vdash E : T_1}{\Gamma \vdash \lambda x : T_0. E : \Pi x : T_0. T_1} \quad (\mathcal{B}/\text{et.fun}) \qquad \frac{\Gamma \vdash E_1 : \Pi x : T_0. T \quad \Gamma \vdash E_0 : T_0}{\Gamma \vdash E_1 E_0 : \{x \leftarrow E_0\} T} \quad (\mathcal{B}/\text{et.app})$$

Local binding In order to type the local binding of a variable to a value, we request that the programmer specify the resulting type of the whole expression. Furthermore this type is not allowed to mention the locally bound variable. This last point is easily understood from the fact that while it would make sense for the variable x to be bound in the type of the body E , this variable cannot be free in $(\text{let } x = E_0 \text{ in } E : T)$. In particular, if E_0 were to create abstract types, there is no way to reference them outside the binding. The necessity for the programmer to specify the type is due to the avoidance problem mentioned in section I.2.2.6, which makes inference of T undecidable.

$$\frac{\Gamma \vdash E_0 : T_0 \quad \Gamma, x : T_0 \vdash E : T \quad \Gamma \vdash T \text{ ok}}{\Gamma \vdash (\text{let } x = E_0 \text{ in } E : T) : T} \quad (\mathcal{B}/\text{et.let})$$

IV.2.4.5 $\boxed{\langle T \rangle, \text{Typ } E}$ Type fields

We can see $\langle _ \rangle$ as a constructor for the type `TYPE` and `Typ _` as the corresponding destructor. This approach yields suitable typing rules.

$$\frac{\Gamma \vdash T \text{ ok}}{\Gamma \vdash \langle T \rangle : \text{TYPE}} \quad (\mathcal{B}/\text{et.type}) \qquad \frac{\Gamma \vdash E : \text{TYPE}}{\Gamma \vdash \text{Typ } E \text{ ok}} \quad (\mathcal{B}/\text{tok.field})$$

IV.2.5 Run-time

Values The class of **values** (generically written V) is given as a subgrammar of expressions.

$V ::=$	value
$() \mid bv \mid \underline{n}$	constant
$\langle T \rangle$	type field
(V_1, V_2)	pair
$\lambda x : T. E$	lambda-abstraction

IV.2.5.1 $E \longrightarrow E'$ Expression reduction

We define the dynamic behaviour of expressions via small-step reduction rules.

Head reduction rules In the language that we have defined so far, head reduction confronts each destructor with a matching constructor, and performs local bindings. We impose a *call-by-value* strategy in the rules ($\mathcal{B}/\text{ered.app}$) and ($\mathcal{B}/\text{ered.let}$). For the time being, we could allow β -reduction in its full generality, and obtain a confluent system; however we will ultimately introduce side effects, which suggests sticking to call-by-value.

$$\begin{array}{ll}
 (\lambda x : T. E) V \longrightarrow \{x \leftarrow V\}E & (\mathcal{B}/\text{ered.app}) \\
 \pi_i (V_1, V_2) \longrightarrow V_i & (\mathcal{B}/\text{ered.proj}) \\
 \text{let } x = V \text{ in } E : T \longrightarrow \{x \leftarrow V\}E & (\mathcal{B}/\text{ered.let})
 \end{array}$$

No reduction in types We do not define any reduction relation on types. Accordingly there is no restriction on T in order for $\langle T \rangle$ to be a value; in particular, if $\langle T \rangle$ contains embedded expressions (as in e.g., $\langle \text{Typ}((\lambda x : \text{TYPE}. x) \langle \text{INT} \rangle) \rangle$) these need not be values. The reason is that computations in types traditionally belong in the compile-time world, hence to typing rules (and where relevant typing algorithms), rather than in the run-time world now under scrutiny. We will later (in system \mathcal{D}) add a construct for run-time type-checking, thus type computations will need to occur during program execution; run-time manipulation of types is also useful for generic programming (see section V.3.2.3).

Evaluation contexts We generically write C for an **evaluation context** of depth 1. These evaluation contexts are defined by the following grammar.

$C ::=$	evaluation context (of depth 1)
$E_1 _$	function argument
$_ V_2$	applied function
$(_, E_2)$	first component of a pair
$(V_1, _)$	second component of a pair
$\pi_i _$	projection ($i \in \{1, 2\}$)
$\text{let } x = _ \text{ in } E : T$	local bound

We have arbitrarily fixed the evaluation order for function application (argument first, then function) and pairing (left to right). This somewhat simplifies the metatheory by not introducing spurious local nondeterminism. We could relax these constraints by authorising the reduction contexts $_ E_2$ et $(E_1, _)$; it is folklore that the resulting reduction relation would be confluent.

The following reduction rule allows expressions to be reduced under contexts. The notation $C \cdot E$ means the expression resulting from placing E inside the context C .

$$\frac{E \longrightarrow E'}{C \cdot E \longrightarrow C \cdot E'} \quad (\mathcal{B}/\text{ered.context})$$

IV.3 Singletons §

IV.3.1 Motivation

IV.3.1.1 Abstract types, concrete types

[Sorry, this fragment has not been translated yet.]

IV.3.1.2 Type sharing

[Sorry, this fragment has not been translated yet.]

IV.3.1.3 Value singletons

So far, we have used singleton types to express type equalities: our singletons were of the form $S(\langle T \rangle)$ for some type T . The purpose of these singletons was to enable making x have the type T' when x has the type T and $\langle T \rangle$ and T' are equivalent: in other words, the judgement $t : S(\langle T \rangle), x : T \vdash x : \text{Typ } t$ should be derivable (one could then substitute $\langle T' \rangle$ for t).

Let us now consider a functor f which creates an abstract type from a type and a value, with a signature of the form $\Pi x : (\Sigma t : \text{TYPE}. T_0). (\Sigma t : \text{TYPE}. T_1)$. As we saw in sections I.2.1.3 and II.5.1.1, $\text{Typ } \pi_1(fx)$ and $\text{Typ } \pi_1(fy)$ are the same types only when x and y have the same behaviour: it is not enough for them to provide the same types.

Let us consider an example potential argument for f , with $T_0 = \text{Typ } t * (\text{Typ } t \rightarrow \text{UNIT})$.

```
module A = struct type t = int let x = (... : t * (t->unit)) end
```

The principal signature of the module A in Objective Caml is

```
module type S = sig type t = int val x : t * (t->unit) end
```

If we want to express that some module B is compatible with A , the best we can do (whether in Objective Caml or in some other ML dialect, or in the language defined so far) is to specify that B has the signature S . Unfortunately this specification is incomplete since it does not distinguish between modules that have x fields with the same type but different values.

One way to illustrate this limitation is to consider an identity functor Id1 capable of taking A as an argument. The principal signature of such a functor (which is based on Leroy's manifest type theory with applicative functors [Ler95]) is the following:

```
functor (A : sig type t          val x : t * (t->unit) end) ->
  sig type t = A.t val x : t * (t->unit) end
```

Notice that the signature of $\text{Id1}(A)$ is the signature S defined above: while it does indicate that $\text{Id1}(A).t$ is equal to $A.t$, nothing connects $\text{Id1}(A).x$ with $A.x$ beyond them having the same type. The higher-order module theory of Dreyer, Cray and Harper [DCH03] does not perform any better on this example. With our notations, the signature of the functor Id1 is $\Pi x : (\Sigma t : \text{TYPE}. \text{Typ } t * (\text{Typ } t \rightarrow \text{UNIT})). (\Sigma t' : S(\pi_1 x). \text{Typ } t' * (\text{Typ } t' \rightarrow \text{UNIT}))$.

In the case of an identity functor E_{Id0} acting solely on a type, i.e., whose argument has the signature TYPE , we get a more precise signature: $\Pi x : \text{TYPE}. S(x)$, clearly indicating that the result of applying the identity functor is equivalent to the argument — $\text{Typ } E_{\text{Id0}}(E_{A0})$ has the signature $S(E_{\text{Id0}}(E_{A0}))$. We shall extend our language so that this property also holds for value fields.

We add singleton types of the form $S(E)$, where E is any value. For example the principal type of the expression `3` is now $S(3)$, the type of values that are equal to 3. We can now give `Id1` a more precise signature, by also giving the second component a singleton type: $\Pi x : (\Sigma t : \text{TYPE}. \text{Type } t * (\text{Type } t \rightarrow \text{UNIT})). (\Sigma t' : S(t). S(\pi_1 x) * S(\pi_2 x))$, i.e., in an Objective Caml-like notation

```
functor (A : sig type t          val x : t * (t->unit) end) ->
      sig type t = A.t    val x = A.x          end
```

Thanks to this signature, `Id1(A)` can have the signature `sig type t = A.t val x = A.x end`, which makes it interchangeable with `A`.

IV.3.1.4 Higher-order singletons

[Sorry, this fragment has not been translated yet.]

IV.3.1.5 A practical example

Let us illustrate higher-order singletons on an example from the author's programming experience. The standard library of Objective Caml provides an implementation of finite sets via a functor `Set.Make`. This functor takes an argument with the following signature:

```
module type Set.OrderedType = sig
  type t
  val compare : t -> t -> int
end
```

A module of signature `Set.OrderedType` provides a type as well as a function which must implement a total order; a set is represented as a search tree. An example of a module with this signature is `String: Set.Make(String).t` is therefore a type for sets of strings. The result returned by `Set.Make` has a signature called `Set.S` from which we quote a relevant excerpt:

```
module type Set.S = sig
  type elt                (*type of elements*)
  type t                  (*type of sets*)
  val add : elt -> t -> t  (*addition function*)
  ...
end
```

An annotation in the definition of `Set.Make` specified that `Set.Make(M).elt = M.t`.

The program under consideration manipulates symbols, which are internally implemented as strings. However only suitable approved strings may be symbols, therefore the type of symbols is an abstract type provided by a module which we call `Syntax`.

```
module Syntax : sig
  type symbol
  val name : symbol -> String.t
  ...
end
```

Since several other modules in the program manipulate sets of symbols, we wish to provide this type alongside `symbol`. How do we mention the type of symbol sets in the signature of `Syntax`? We must specify a module of signature `Set.S`, indicating that the element type is that of symbols. For that purpose we need to define a symbol module.

```
module Syntax : sig
  module Symbol : Set.OrderedType
  module SymSet : Set.S with type elt = Symbol.t
                      and type t = Set.Make(Symbol).t
  ...
end
```

We could also write as follows:

```
module Syntax : sig
  module Symbol : Set.OrderedType
  module SymSet : Set.S
  ...
end with module SymSet = Set.Make(Symbol)
```

The signatures above are equivalent in Objective Caml.

The difficulty arises when writing the implementation of the `Syntax` module. We may write

```
module Syntax = struct
  module Symbol = struct type t = String.t let compare = String.compare end
  module SymSet = Set.Make(Symbol)
  ...
end
```

or even

```
module Syntax = struct
  module Symbol = String
  module SymSet = Set.Make(Symbol)
  ...
end
```

Unfortunately the resulting `SymSet` module is not compatible with `Set.Make(String)`. Since Objective Caml only ever compares type fields of modules, its type analysis remembers the equality between `Symbol.t` and `String.t` but not that between `Symbol` and `String`, therefore the types `Set.Make(String).t` and `Set.Make(Symbol).t` cannot safely be declared compatible.

Since our `Syntax` module calls other, lower-level modules that manipulate sets of strings, the incompatibility of `Set.Make(String).t` with `Set.Make(Symbol).t` is a major problem. The solution we chose was to only expose the symbol type, and not its comparison function:

```
module Syntax = struct
  type symbol = String.t
  module SymSet = Set.Make(String)
  ...
end
```

```

end : sig
  type symbol
  module SymSet : Set.S with type t = symbol
  ...
end

```

The disadvantage of this signature is that it hides the choice of a set implementation: that fact that `SymSet` is the result of an application of the `Set.Make` functor does not appear. This is a problem because some users of the `Syntax` modules manipulate data with more complex structures (e.g., sets of sets of symbols) built from functors that take an argument produced by `Set.Make`. We had to provide these data structures alongside `SymSet` in the `Syntax` module, even though these extra data structures had nothing to do in `Syntax` from a code organisation point of view.

In this case, simply being able to write `module Symbol = String` in the implementation of the `Syntax` module in such a way as to make the types `Set.Make(String).t` and `Set.Make(Symbol).t` interchangeable would have permitted the code to be organised properly, in particular with respect to abstraction. In system \mathcal{S} , this is possible, since the `Symbol` module will have the signature $\mathcal{S}(x$ called "String") which result in `Set.Make(String)` and `Set.Make(Symbol)` being compatible within the implementation of the `Syntax` module.

IV.3.2 Properties

[Sorry, this fragment has not been translated yet.]

IV.3.3 Typing rules

We state typing rules for system \mathcal{S} . The operational semantics (consisting of the reduction rules ($\mathcal{S}/\text{ered.app}$), ($\mathcal{S}/\text{ered.proj}$), ($\mathcal{S}/\text{ered.let}$), ($\mathcal{S}/\text{ered.context}$), as well as the definitions of values and reduction contexts) is unchanged from system \mathcal{B} .

System \mathcal{S} contains new typing judgements for subtyping, conversion and convertibility.

$J ::=$ **local judgement right-hand side**

```

...
T  $\longrightarrow$  T'  typing conversion
T  $\equiv$  T'         convertibility equivalence on types
E  $\longrightarrow$  E'  expression conversion
E  $\equiv$  E'         convertibility equivalence on expressions
T1 <: T2       subtyping

```

Most typing rules of \mathcal{B} are included in \mathcal{S} . The following rules are taken as is from \mathcal{B} (and will not be repeated here):

- environment correction: all rules — ($\mathcal{S}/\text{envok.nil}$), ($\mathcal{S}/\text{envok.x}$);
- type correction: all rules — ($\mathcal{S}/\text{tok.base.unit}$), ($\mathcal{S}/\text{tok.base.bool}$), ($\mathcal{S}/\text{tok.base.int}$), ($\mathcal{S}/\text{tok.type}$), ($\mathcal{S}/\text{tok.pair}$), ($\mathcal{S}/\text{tok.fun}$), ($\mathcal{S}/\text{tok.field}$);
- expression typing: all rules except projections and local binding — ($\mathcal{S}/\text{et.base.unit}$), ($\mathcal{S}/\text{et.base.bool}$), ($\mathcal{S}/\text{et.base.int}$), ($\mathcal{S}/\text{et.x}$), ($\mathcal{S}/\text{et.pair}$), ($\mathcal{S}/\text{et.fun}$), ($\mathcal{S}/\text{et.app}$), ($\mathcal{S}/\text{et.type}$).

We omit local binding in system \mathcal{S} because it is superfluous (see section IV.2.1.4), thus avoiding the need to take them into account in singleton typing rules.

IV.3.3.1 $\Gamma \vdash T <: T' ; \dots$ **Subtyping**

The subtyping relation explains how an expression may have more than one type, some of which are more precise than others. Intuitively, the type T is a subtype of T' whenever any expression that has the type T also has the type T' . We engrave the forward implication with an *implicit subtyping* rule.

$$\frac{\Gamma \vdash E : T \quad \Gamma \vdash T <: T'}{\Gamma \vdash E : T'} \text{ (\$/et.sub)}$$

Our subtyping relation is defined syntactically (by deduction rules) rather than semantically, in that nothing mandates the reverse application: it is possible for $\Gamma \vdash E : T'$ to be derivable whenever $\Gamma \vdash E : T$ is without the judgement $\Gamma \vdash T <: T'$ being derivable. Whether a subtyping relation should be semantic (i.e., fully capture type subsumption for expressions) is debatable. On the one hand, a semantic subtyping relation permits a set-theoretic interpretation of types as sets of expressions. On the other hand, the rules needed to enforce semanticity would be fragile, in that they would not play well with extensions of the system. For example, if V is a value of type T , then with semantic subtyping $T <: S(V)$ must hold whenever T contains the single value V , which may happen coincidentally. Consider for instance the type $\Pi t : \text{TYPE}. \text{Typ } t \rightarrow \text{Typ } t$, which obviously contains the polymorphic identity function $(\lambda t : \text{TYPE}. \lambda x : \text{Typ } t. x)$. In a suitably weak system, such as system \mathcal{S} , a parametricity [Wad89] result ensures that there is no other function of this type. However adding either dynamic type-checking (as we will do in system \mathcal{D}) or an unrestricted fixpoint combinator would let one write other functions of this type.

If two types are interconvertible, they are subtypes of one another. Thus subtyping includes computational equivalences on types.

$$\frac{\Gamma \vdash T \equiv T'}{\Gamma \vdash T <: T'} \text{ (\$/tsub.eq)}$$

Subtyping is a preorder. The rule $(\$/tsub.eq)$ enforces reflexivity; we state transitivity.

$$\frac{\Gamma \vdash T <: T' \quad \Gamma \vdash T' <: T''}{\Gamma \vdash T <: T''} \text{ (\$/tsub.trans)}$$

IV.3.3.2 $S(E)$ **Singletons**

Singleton types appear through three generic rules, which have no constraint on the type of the expression whose singleton is taken. The singleton $S(E)$ is well-formed as soon as E has some type T ; any well-typed expression E thus has the type $S(E)$, and $S(E)$ is a subtype of any of its types. Note that in order to prove that E has the type $S(E)$, one must first find some type T that E has and then apply $(\$/et.sing)$.

$$\frac{\Gamma \vdash E : T}{\Gamma \vdash S(E) \text{ ok}} \text{ (\$/tok.sing)} \quad \frac{\Gamma \vdash E : T}{\Gamma \vdash E : S(E)} \text{ (\$/et.sing)} \quad \frac{\Gamma \vdash E : T}{\Gamma \vdash S(E) <: T} \text{ (\$/tsub.sing)}$$

The rule $(\$/et.sing)$ (combined with the subtyping rules) is a particularly powerful instance of a selfification rule in a module language with manifest types [HL94, Ler94] (see section I.2.2.2).

IV.3.3.3 $\Gamma \vdash E : T ; \Gamma \vdash T_1 <: T_2$ **Expression typing**

As in system \mathcal{B} , we first assign non-dependent types to pairs. A dependent type can be obtained via the subtyping rule $(\$/tsub.cong.pair)$ (recall that $T_1 * T_2$ is an abbreviation for $\Sigma x : T_1. T_2$). In this rule, note that the second premise contains the stronger hypothesis on x , namely $x : T_1'$, which follows from the fact that the hypothesis $x : T_2'$ might not be enough to ensure that T_1'' be valid. A third premise ensures that $\Sigma x : T_2'. T_2''$ is well-formed, which requires that T_2'' be well-formed under the weaker hypothesis $x : T_2'$.

$$\frac{\Gamma \vdash T_1 <: T'_1 \quad \Gamma, x : T_1 \vdash T_2 <: T'_2 \quad \Gamma, x : T'_1 \vdash T'_2 \text{ ok}}{\Gamma \vdash \Sigma x : T_1. T_2 <: \Sigma x : T'_1. T'_2} \text{ (\$/tsub.cong.pair)}$$

In order to type a projection, the argument expression must be given a (dependent) pair type. Typing the first projection is simple, as its type is readily available in the pair type. Typing the second projection is more complicated. If the expression E has the type $\Sigma x : T_1. T_2$ then $\pi_2 E$ only has the type T_2 with a suitably strong hypothesis on the variable x . For instance the expression $(3, 3)$ has the type $\Sigma x : \text{INT}. S(x)$, and while $x : S(3) \vdash \pi_2(3, 3) : S(x)$ holds, $x : \text{INT} \vdash \pi_2(3, 3) : S(x)$ does not.

$$\frac{\Gamma \vdash E : \Sigma x : T_1. T_2}{\Gamma \vdash \pi_1 E : T_1} \text{ (\$/et.proj.1)} \quad \frac{\Gamma \vdash E : \Sigma x : T_1. T_2 \quad \Gamma \vdash E_1 : S(\pi_1 E)}{\Gamma \vdash \pi_2 E : \{x \leftarrow E_1\} T_2} \text{ (\$/et.proj.2)}$$

The premises of the rule $(\$/\text{et.proj.2})$ are usually unduely constraining, and we will often use one of two admissible variants that only require the first premise $\Gamma \vdash E : \Sigma x : T_1. T_2$. The most common rule replaces the variable x by the first projection of E in T_2 . Another variant keeps track of the first component via a variable x which is constrained to the type $S(\pi_1 E)$.

$$\frac{\Gamma \vdash E : \Sigma x : T_1. T_2}{\Gamma \vdash \pi_2 E : \{x \leftarrow \pi_1 E\} T_2} \text{ (et.proj.2s)} \quad \frac{\Gamma \vdash E : \Sigma x : T_1. T_2}{\Gamma, x : S(\pi_1 E) \vdash \pi_2 E : T_2} \text{ (et.proj.2x)}$$

We state the usual congruence rule for subtyping through dependent products. This rule is similar to $(\$/\text{tsub.cong.pair})$, with hypotheses suitably reversed when dealing with the contravariant argument type.

$$\frac{\Gamma \vdash T'_0 <: T_0 \quad \Gamma, x : T'_0 \vdash T_1 <: T'_1 \quad \Gamma, x : T_0 \vdash T_1 \text{ ok}}{\Gamma \vdash \Pi x : T_0. T_1 <: \Pi x : T'_0. T'_1} \text{ (\$/tsub.cong.fun)}$$

IV.3.3.4 $\Gamma \vdash T \equiv T' ; \Gamma \vdash E \equiv E'$ Convertibility equivalences

We define an equivalence relation on types and one on expressions as the smallest equivalence relation containing the appropriate conversion relation. (Strictly speaking these are two families of relations, indexed by environments.) These relations are called **convertibility**.

$$\begin{array}{ll} \frac{\Gamma \vdash T \text{ ok}}{\Gamma \vdash T \equiv T} \text{ (\$/teq.refl)} & \frac{\Gamma \vdash T_2 \equiv T_1}{\Gamma \vdash T_1 \equiv T_2} \text{ (\$/teq.sym)} \\ \frac{\Gamma \vdash T_1 \equiv T_2 \quad \Gamma \vdash T_2 \equiv T_3}{\Gamma \vdash T_1 \equiv T_3} \text{ (\$/teq.trans)} & \frac{\Gamma \vdash T_1 \longrightarrow T_2}{\Gamma \vdash T_1 \equiv T_2} \text{ (\$/teq.conv)} \\ \frac{\Gamma \vdash E : T}{\Gamma \vdash E \equiv E} \text{ (\$/eeq.refl)} & \frac{\Gamma \vdash E_2 \equiv E_1}{\Gamma \vdash E_1 \equiv E_2} \text{ (\$/eeq.sym)} \\ \frac{\Gamma \vdash E_1 \equiv E_2 \quad \Gamma \vdash E_2 \equiv E_3}{\Gamma \vdash E_1 \equiv E_3} \text{ (\$/eeq.trans)} & \frac{\Gamma \vdash E_1 \longrightarrow E_2}{\Gamma \vdash E_1 \equiv E_2} \text{ (\$/eeq.conv)} \end{array}$$

IV.3.3.5 $\Gamma \vdash T \longrightarrow T'$ Type conversion

Type conversion mostly consists in conversion of embedded expressions. Types additionally undergo some slight simplification.

Conversion is only defined on valid types; the rules defining conversion contain correction premises in addition to the conversion premises in context rules.

Contexts At the type level, conversion applies recursively to all subtypes. We state this via context rules which allow conversion of both type and expression subterms of types.

$$\begin{array}{c}
\frac{\Gamma \vdash T_1 \longrightarrow T'_1 \quad \Gamma, x : T_1 \vdash T_2 \text{ ok}}{\Gamma \vdash \Sigma x : T_1. T_2 \longrightarrow \Sigma x : T'_1. T_2} (\$/\text{tconv.cong.pair.1}) \\
\frac{\Gamma, x : T_1 \vdash T_2 \longrightarrow T'_2 \quad \Gamma \vdash T_1 \text{ ok}}{\Gamma \vdash \Sigma x : T_1. T_2 \longrightarrow \Sigma x : T_1. T'_2} (\$/\text{tconv.cong.pair.2}) \\
\frac{\Gamma \vdash T_0 \longrightarrow T'_0 \quad \Gamma, x : T_0 \vdash T_1 \text{ ok}}{\Gamma \vdash \Pi x : T_0. T_1 \longrightarrow \Pi x : T'_0. T_1} (\$/\text{tconv.cong.fun.arg}) \\
\frac{\Gamma \vdash T_0 \text{ ok} \quad \Gamma, x : T_0 \vdash T_1 \longrightarrow T'_1}{\Gamma \vdash \Pi x : T_0. T_1 \longrightarrow \Pi x : T_0. T'_1} (\$/\text{tconv.cong.fun.ret}) \\
\frac{\Gamma \vdash E \longrightarrow E'}{\Gamma \vdash S(E) \longrightarrow S(E')} (\$/\text{tconv.cong.sing}) \qquad \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E : \text{TYPE}}{\Gamma \vdash \text{Typ } E \longrightarrow \text{Typ } E'} (\$/\text{tconv.cong.field})
\end{array}$$

Simplifications The term $\text{Typ } \langle T \rangle$ can be seen as a destructor applied to the matching constructor applied to T ; it is equivalent to T .

$$\frac{\Gamma \vdash T \text{ ok}}{\Gamma \vdash \text{Typ } \langle T \rangle \longrightarrow T} (\$/\text{tconv.field})$$

Semantic rules We declare that the `UNIT` type contains a single value: this type is isomorphic to a singleton, and the rule $(\$/\text{tconv.unit})$ enshrines this equivalence into the syntactic definition of type equivalence. The choice of orientation in this rule does not matter greatly.

$$\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash S(()) \longrightarrow \text{UNIT}} (\$/\text{tconv.unit})$$

IV.3.3.6 $\Gamma \vdash E \longrightarrow E'$ Expression conversion

If E evaluates to E' and E has the type $S(E)$, then type preservation requires that E' have the type $S(E)$. This is ensured by making E convertible to E' . Conversion thus includes run-time reduction².

Conversion is only defined on well-typed expressions; the rules defining conversion contain correction premises in addition to the conversion premises in context rules.

We do not state any conversion rule for local binding expressions (i.e., expressions of the form $\text{let } x = E_0 \text{ in } E : T$). As indicated in section IV.2.1.4, this construct is just syntactic sugar when E_0 is pure; in \mathcal{E} , where local binding becomes useful, it is always judged impure and thus not subject to conversion.

Contexts Unlike for the run-time reduction relation, there is no particular advantage to keeping conversion deterministic, while confluence of conversion is a key property of our metatheoretic study. We therefore permit arbitrary evaluation strategies, and allow reduction in any context. In preparation for the addition of impure constructs to the language, which only come in conversion inside function bodies that hide the impurity, we state the slightly peculiar rule $(\$/\text{econv.cong.fun.body})$ to allow the conversion of any pure subexpression of an impure subexpression of a pure expression.

$$\frac{\Gamma \vdash T_0 \longrightarrow T'_0 \quad \Gamma, x : T_0 \vdash E_1 : T_1}{\Gamma \vdash (\lambda x : T_0. E_1) \longrightarrow (\lambda x : T'_0. E_1)} (\$/\text{econv.cong.fun.arg})$$

²For pure expressions, as we shall see in system \mathcal{E} .

$$\begin{array}{c}
 \frac{\Gamma, x : T_0 \vdash E \longrightarrow E' \quad \Gamma, x : T_0, y : S(E) \vdash E_1 : T_1}{\Gamma \vdash (\lambda x : T_0. \{y \leftarrow E\} E_1) \longrightarrow (\lambda x : T_0. \{y \leftarrow E'\} E_1)} (\mathcal{S}/\text{econv.cong.fun.body}) \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E : \Pi x : T_0. T_1 \quad \Gamma \vdash E_0 : T_0}{\Gamma \vdash E E_0 \longrightarrow E' E_0} (\mathcal{S}/\text{econv.cong.app.fun}) \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E : T_0 \quad \Gamma \vdash E_1 : \Pi x : T_0. T_1}{\Gamma \vdash E_1 E \longrightarrow E_1 E'} (\mathcal{S}/\text{econv.cong.app.arg}) \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E_2 : T_2}{\Gamma \vdash (E, E_2) \longrightarrow (E', E_2)} (\mathcal{S}/\text{econv.cong.pair.1}) \quad \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E_1 : T_1}{\Gamma \vdash (E_1, E) \longrightarrow (E_1, E')} (\mathcal{S}/\text{econv.cong.pair.2}) \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E : \Sigma x : T_1. T_2}{\Gamma \vdash \pi_i E \longrightarrow \pi_i E'} (\mathcal{S}/\text{econv.cong.proj}) \quad \frac{\Gamma \vdash T \longrightarrow T'}{\Gamma \vdash \langle T \rangle \longrightarrow \langle T' \rangle} (\mathcal{S}/\text{econv.cong.field})
 \end{array}$$

Head reduction The following rules describe the usual evaluation of lambda terms with pairs.

$$\frac{\Gamma, x : T_0 \vdash E_1 : T_1 \quad \Gamma \vdash E_0 : T_0}{\Gamma \vdash (\lambda x : T_0. E_1) E_0 \longrightarrow \{x \leftarrow E_0\} E_1} (\mathcal{S}/\text{econv.app}) \quad \frac{\Gamma \vdash E_1 : T_1 \quad \Gamma \vdash E_2 : T_2}{\Gamma \vdash \pi_i (E_1, E_2) \longrightarrow E_i} (\mathcal{S}/\text{econv.proj})$$

IV.3.3.7 Extensionality

We state extensionality rules for system \mathcal{S} . Such rules can have many forms; we choose to use conversion rules, oriented as eta-expansions. For example the rule $(\mathcal{S}/\text{econv.eta.pair})$ may be read as “any expression that is typable as a pair can be rewritten in such a way as to expose the pair structure”.

Given the choice of using conversion rules (η conversions to supplement the β conversions above), there is a further choice between expansions and contractions. A major technical advantage of expansions is that they do not hurt the confluence of the system, unlike eta-contractions [Klo80]. Expansions do however have the obvious defect of breaking normalisation. In practice, it seems preferable to express extensionality using expansions, and when normalisation is required to restrict their use to a finite domain (given by the structure of the type of the converted expression) [Gog05].

$$\frac{\Gamma \vdash E : \text{TYPE}}{\Gamma \vdash E \longrightarrow \langle \text{Typ } E \rangle} (\mathcal{S}/\text{econv.eta.field}) \\
 \frac{\Gamma \vdash E : \Pi x : T_0. T_1}{\Gamma \vdash E \longrightarrow (\lambda x : T_0. E x)} (\mathcal{S}/\text{econv.eta.fun}) \quad \frac{\Gamma \vdash E : \Sigma x : T_1. T_2}{\Gamma \vdash E \longrightarrow (\pi_1 E, \pi_2 E)} (\mathcal{S}/\text{econv.eta.pair})$$

IV.4 Sealing \mathcal{E}

IV.4.1 Sealing

[Sorry, this fragment has not been translated yet.]

IV.4.2 An effect system

IV.4.2.1 Introduction

[Sorry, this fragment has not been translated yet.]

IV.4.2.2 Purity

[Sorry, this fragment has not been translated yet.]

IV.4.2.3 Projectibility, separability and comparability

In the system we are describing, a “well-behaved” module is a pure module. Purity is a very strong notion: a pure module is fully known statically (it has its singleton type, and singleton types fully characterise an object). There are finer notions to determine legitimate uses of a module; in this section we will discuss some of these.

Classical module calculi, notably Harper and Lillibridge’s translucent sums [HL94, Lil97] and Leroy’s manifest types [Ler94, Ler95] focus on the concept of **projectibility** (see section I.2.2.2). The question is, given a module expression M , whether the term $M.t$ may be used to form a type. If so, the module M is said to be **projectible**. In our notation, E is projectible if and only if $\text{Typ } \pi_1 E$ is a correct type. We approximate projectibility by purity: $\text{Typ } \pi_1 E$ is correct if and only if E is pure (and has an appropriate signature). This is indeed an approximation since purity is a stronger notion; for example, in the following code fragment, the modules **A** and **B** are both projectible, but only **A** is pure, while **B** is impure.

```
module A = struct type t = int let x = 3 end
module B = struct type t = int let x = ref 3 end
```

A closely related notion is that of **comparability**: a module is said to be **comparable** when its equivalence with another module can be tested. In the calculus of Dreyer, Crary and Harper [DCH03], the notions of comparability and projectability coincide, since testing the equivalence of two modules amounts to comparing their type components. In our calculus, purity stands for comparability as well as projectability (we treat type and value components identically).

In section I.3.1.1, we mentioned the issue of *phase separation*, i.e., clearly differentiating between the static phase of the program, which includes a type-check that rejects programs that would go wrong, and the dynamic phase, during which computation proceeds without errors. In the core of ML, each phase is closely associated with one part of the language: type-checking is mostly concerned with types, and computation is mostly concerned with expressions. This is no longer true when modules are considered, as they mix types and expressions at the syntactic level. Nevertheless one usually tries to separate types and expressions in the metatheory of modules, in order to distinguish between the static and dynamic aspects.

In his analysis of ML modules [Dre05], Dreyer distinguishes between two levels of purity in modules. A module is said to be **totally pure**³ if it is pure in our sense, i.e., that its evaluation does not trigger an effect of any kind. A module is said to be **partially pure** if its type components can be fully determined without triggering an effect. For example the module **B** above is partially pure but not totally pure. A projectible module must be partially pure.

One difficulty with partial purity is in deciding whether the effects of an expression have an influence on its type parts. Total purity is of course a sufficient condition. A weaker sufficient condition is **separability**. This notion was introduced by Harper, Mitchell and Moggi [HMM90] and is expounded in the context of a module calculus with functors by Dreyer [Dre05]. A module is **separable** if its type components do not depend at all on any computation that may have effects, in particular any core-expression-level computation. A separable module is always partially pure,

³Actually Dreyer uses the wording “dynamically pure” and “statically pure” where we use “totally pure” and “partially pure”. We changed the terminology because we will use “dynamically pure” and “statically pure” in a different sense, following other work by Dreyer [DCH03].

hence projectible, but may be impure, like B above. Conversely, in a language with first-class modules, one can easily write pure inseparable modules, such as the module C in the following program fragment.

```

let n = read_int ()
module A = struct type t = int let x = 3 end
module B = struct type t = int let x = ref 3 end
module A' = struct type t = bool let x = true end
module B' = struct type t = bool let x = ref true end
module C = if n >= 0 then A else A'
module D = if n >= 0 then B else B'

```

The module D is partially pure but neither totally pure nor separable.

An important conclusion of the discussion of type singletons and module equivalence in IV.3.1 was that in the presence of functors, separability is hard to analyse — and this is why we did not try to analyse it, and instead integrated expressions with types when testing module equivalence. Separability looks all the less enticing to us as we eventually want to be able to compare types dynamically, which means that our notion of equivalence must work well even in inseparable cases. In the present work, we do not try to go further than (total) purity. If refinements are desired, rather than introduce separability, we suggest instead to make the effect system more sophisticated, and in particular to make it possible to detect partial purity by “declassifying” effects that do not impact the value of an expression.

IV.4.3 Formal presentation

We give a formal description of system \mathcal{E} , which consists of adding the sealing construct to \mathcal{S} and, more importantly, an effect system.

IV.4.3.1 Syntax

We first give the syntax of effects.

```

 $\gamma ::=$  effect
  P pure
  I impure

```

Recall that effects are ordered: the relation $\gamma_1 \sqsubseteq \gamma_2$ is such that $P \sqsubseteq I$ (but not the converse). We write $\gamma_1 \sqcup \gamma_2$ for the least upper bound of γ_1 and γ_2 , and $\gamma_1 \sqcap \gamma_2$ for their greatest lower bound.

The syntax of system \mathcal{E} extends that of system \mathcal{S} by adding an effect annotation where necessary, viz.,

- on function types, henceforth written $\Pi x : T_0. {}^\gamma T_1$; they are abbreviated as $T_0 \rightarrow^\gamma T_1$ when x is not free in T_1 ;
- on expression typing judgements, henceforth written $E : {}^\gamma T$.

We sometimes omit the effect annotation when it is P, thus we might write the type of a pure function as $\Pi x : T_0. T_1$ or $T_0 \rightarrow T_1$. Furthermore the syntax of expressions now comprises sealing.

```

 $T ::=$  type
  ...
   $\Pi x : T_0. {}^\gamma T_1$  dependent product (also written  $T_1 \rightarrow^\gamma T_2$  when  $x \notin \mathbf{fv} T_1$ )

```

$E ::=$	expression (module)
...	
$E !! T$	sealing
$J ::=$	typing judgement right-hand side
...	
$E :^{\gamma} T$	expression typing

IV.4.3.2 $E \longrightarrow E'$ Run-time

The only run-time novelty of system \mathcal{E} is the need to reduce a sealing construct. The usual intuition in ML-like languages is that types have no bearing on execution, only on static type-checking; in this light, $E !! T$ is equivalent to E at run-time.

$$V !! T \longrightarrow V \quad (\mathcal{E}/\text{ered.seal})$$

The sealed expression is first reduced to a value.

$C ::=$	evaluation context (of depth 1)
...	
$_ !! T$	sealing

The rules $(\mathcal{E}/\text{ered.app})$, $(\mathcal{E}/\text{ered.proj})$, $(\mathcal{E}/\text{ered.let})$ et $(\mathcal{E}/\text{ered.context})$ are inherited from system \mathcal{B} via \mathcal{S} .

IV.4.3.3 $\Gamma \vdash \dots$ Typing: correction, equivalences, subtyping

System \mathcal{E} contains all the typing rules of system \mathcal{S} , and adds one for sealing. However the inherited rules must usually be modified to add an effect annotation. We will restate affected rules and explain the effect of effects.

Inhreited rules The following rules are taken as is from system \mathcal{B} via \mathcal{S} :

- $(\mathcal{E}/\text{envok.nil})$, $(\mathcal{E}/\text{envok.x})$;
- $(\mathcal{E}/\text{tok.base.unit})$, $(\mathcal{E}/\text{tok.base.bool})$, $(\mathcal{E}/\text{tok.base.int})$, $(\mathcal{E}/\text{tok.type})$, $(\mathcal{E}/\text{tok.pair})$;
- $(\mathcal{E}/\text{teq.refl})$, $(\mathcal{E}/\text{teq.sym})$, $(\mathcal{E}/\text{teq.trans})$, $(\mathcal{E}/\text{teq.conv})$;
- $(\mathcal{E}/\text{eeq.sym})$, $(\mathcal{E}/\text{eeq.trans})$, $(\mathcal{E}/\text{eeq.conv})$;
- $(\mathcal{E}/\text{tconv.cong.pair.1})$, $(\mathcal{E}/\text{tconv.cong.pair.2})$, $(\mathcal{E}/\text{econv.cong.field})$, $(\mathcal{E}/\text{tconv.cong.sing})$, $(\mathcal{E}/\text{tconv.field})$, $(\mathcal{E}/\text{tconv.unit})$;
- $(\mathcal{E}/\text{tsub.trans})$, $(\mathcal{E}/\text{tsub.eq})$, $(\mathcal{E}/\text{tsub.cong.pair})$.

Apart from expression typing and from $(\mathcal{E}/\text{tsub.cong.fun})$, the modifications to the rules of system \mathcal{S} consist of requiring expressions embedded in types to be pure, and allow dependent product types to bear effect annotations. The rules $(\mathcal{E}/\text{econv.cong.fun.arg})$ and $(\mathcal{E}/\text{econv.cong.fun.body})$ do however permit the body of the function to be pure, as all that is required is that the function itself be a pure expression.

$$\frac{\Gamma \vdash T' \text{ ok} \quad \Gamma, x : T' \vdash T'' \text{ ok}}{\Gamma \vdash \Pi x : T'. {}^{\gamma}T'' \text{ ok}} (\mathcal{E}/\text{tok.fun}) \quad \frac{\Gamma \vdash E :^{\text{P}} \text{TYPE}}{\Gamma \vdash \text{Typ } E \text{ ok}} (\mathcal{E}/\text{tok.field}) \quad \frac{\Gamma \vdash E :^{\text{P}} T}{\Gamma \vdash S(E) \text{ ok}} (\mathcal{E}/\text{tok.sing})$$

$$\begin{array}{c}
 \frac{\Gamma \vdash E :^P T}{\Gamma \vdash E \equiv E} \text{ (\mathcal{E}/\text{eeq.refl})} \qquad \frac{\Gamma \vdash E :^P T}{\Gamma \vdash S(E) <: T} \text{ (\mathcal{E}/\text{tsub.sing})} \\
 \\
 \frac{\Gamma \vdash T_0 \longrightarrow T'_0 \quad \Gamma, x : T_0 \vdash T_1 \text{ ok}}{\Gamma \vdash \Pi x : T_0. \gamma T_1 \longrightarrow \Pi x : T'_0. \gamma T_1} \text{ (\mathcal{E}/\text{tconv.cong.fun.arg})} \\
 \frac{\Gamma \vdash T_0 \text{ ok} \quad \Gamma, x : T_0 \vdash T_1 \longrightarrow T'_1}{\Gamma \vdash \Pi x : T_0. \gamma T_1 \longrightarrow \Pi x : T_0. \gamma T'_1} \text{ (\mathcal{E}/\text{tconv.cong.fun.ret})} \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E :^P \text{TYPE}}{\Gamma \vdash \text{Typ } E \longrightarrow \text{Typ } E'} \text{ (\mathcal{E}/\text{tconv.cong.field})} \\
 \frac{\Gamma \vdash T_0 \longrightarrow T'_0 \quad \Gamma, x : T_0 \vdash E_1 :^\gamma T_1}{\Gamma \vdash (\lambda x : T_0. E_1) \longrightarrow (\lambda x : T'_0. E_1)} \text{ (\mathcal{E}/\text{econv.cong.fun.arg})} \\
 \frac{\Gamma, x : T_0 \vdash E \longrightarrow E' \quad \Gamma, x : T_0, y : S(E) \vdash E_1 :^\gamma T_1}{\Gamma \vdash (\lambda x : T_0. \{y \leftarrow E\} E_1) \longrightarrow (\lambda x : T_0. \{y \leftarrow E'\} E_1)} \text{ (\mathcal{E}/\text{econv.cong.fun.body})} \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E_2 :^P T_2}{\Gamma \vdash (E, E_2) \longrightarrow (E', E_2)} \text{ (\mathcal{E}/\text{econv.cong.pair.1})} \qquad \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E_1 :^P T_1}{\Gamma \vdash (E_1, E) \longrightarrow (E_1, E')} \text{ (\mathcal{E}/\text{econv.cong.pair.2})} \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E :^P \Sigma x : T_1. T_2}{\Gamma \vdash \pi_i E \longrightarrow \pi_i E'} \text{ (\mathcal{E}/\text{econv.cong.proj})} \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E :^P \Pi x : T_0. ^P T_1 \quad \Gamma \vdash E_0 :^P T_0}{\Gamma \vdash E E_0 \longrightarrow E' E_0} \text{ (\mathcal{E}/\text{econv.cong.app.fun})} \\
 \frac{\Gamma \vdash E \longrightarrow E' \quad \Gamma \vdash E :^P T_0 \quad \Gamma \vdash E_1 :^P \Pi x : T_0. ^P T_1}{\Gamma \vdash E_1 E \longrightarrow E_1 E'} \text{ (\mathcal{E}/\text{econv.cong.app.arg})} \\
 \frac{\Gamma \vdash E_1 :^P T_1 \quad \Gamma \vdash E_2 :^P T_2}{\Gamma \vdash \pi_i (E_1, E_2) \longrightarrow E_i} \text{ (\mathcal{E}/\text{econv.proj})} \qquad \frac{\Gamma, x : T_0 \vdash E_1 :^P T_1 \quad \Gamma \vdash E_0 :^P T_0}{\Gamma \vdash (\lambda x : T_0. E_1) E_0 \longrightarrow \{x \leftarrow E_0\} E_1} \text{ (\mathcal{E}/\text{econv.app})} \\
 \frac{\Gamma \vdash E :^P \text{TYPE}}{\Gamma \vdash E \longrightarrow \langle \text{Typ } E \rangle} \text{ (\mathcal{E}/\text{econv.eta.field})} \\
 \frac{\Gamma \vdash E :^P \Pi x : T_0. \gamma T_1}{\Gamma \vdash E \longrightarrow (\lambda x : T_0. E x)} \text{ (\mathcal{E}/\text{econv.eta.fun})} \qquad \frac{\Gamma \vdash E :^P \Sigma x : T_1. T_2}{\Gamma \vdash E \longrightarrow (\pi_1 E, \pi_2 E)} \text{ (\mathcal{E}/\text{econv.eta.pair})}
 \end{array}$$

Subtyping for functions The congruence rule for subtyping across dependent products needs to take effects into account. A function type is smaller than another function type when the domain of the first is smaller, the image of the first is larger, and the first allows fewer effects to occur during execution.

$$\frac{\Gamma \vdash T'_0 <: T_0 \quad \Gamma, x : T'_0 \vdash T_1 <: T'_1 \quad \Gamma, x : T_0 \vdash T_1 \text{ ok} \quad \text{when } \gamma \sqsubseteq \gamma'}{\Gamma \vdash \Pi x : T_0. \gamma T_1 <: \Pi x : T'_0. \gamma' T'_1} \text{ (\mathcal{E}/\text{tsub.cong.fun})}$$

IV.4.3.4 $\Gamma \vdash E :^\gamma T$ Expression typing

Expression typing judgements now carry an effect annotation.

Constants, variables, type fields Constantes, variables and type fields are always pure.

$$\begin{array}{c}
\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash () :^P \text{UNIT}} \quad (\mathcal{E}/\text{et.base.unit}) \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \text{bv} :^P \text{BOOL}} \quad (\mathcal{E}/\text{et.base.bool}) \qquad \frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \underline{n} :^P \text{INT}} \quad (\mathcal{E}/\text{et.base.int}) \\
\frac{\Gamma \vdash \text{ok} \quad \text{when } x : T \in \Gamma}{\Gamma \vdash x :^P T} \quad (\mathcal{E}/\text{et.x}) \qquad \frac{\Gamma \vdash T \text{ ok}}{\Gamma \vdash \langle T \rangle :^P \text{TYPE}} \quad (\mathcal{E}/\text{et.type})
\end{array}$$

Pairs Pairs are simple data structures: the type of a pair simply indicates the types of its components. The information as to which component of the pair carries which effect is lost. Therefore typing a pair requires a common effect annotation to be found for its components (the rule $(\mathcal{E}/\text{et.sub})$ used on each side allows one to use the least upper bound). Similarly, the effects of a first projection are the effects of the original expression. The second projection can only be used on a pure expression since the expression appears in a type⁴.

$$\begin{array}{c}
\frac{\Gamma \vdash E_1 :^\gamma T_1 \quad \Gamma \vdash E_2 :^\gamma T_2}{\Gamma \vdash (E_1, E_2) :^\gamma T_1 * T_2} \quad (\mathcal{E}/\text{et.pair}) \\
\frac{\Gamma \vdash E :^\gamma \Sigma x : T_1. T_2}{\Gamma \vdash \pi_1 E :^\gamma T_1} \quad (\mathcal{E}/\text{et.proj.1}) \qquad \frac{\Gamma \vdash E :^P \Sigma x : T_1. T_2 \quad \Gamma \vdash E_1 :^P S(\pi_1 E)}{\Gamma \vdash \pi_2 E :^P \{x \leftarrow E_1\} T_2} \quad (\mathcal{E}/\text{et.proj.2})
\end{array}$$

Functions An immediate function is always pure. The abstraction construct suspends the effects of the function, which are reflected in the effect annotation on the function type. When a function is applied, the effects of the body are released and add to the effects of the function expression. We require the argument of a function to be pure as it is substituted into the result type.

$$\frac{\Gamma, x : T_0 \vdash E :^\gamma T_1}{\Gamma \vdash \lambda x : T_0. E :^P \Pi x : T_0. ^\gamma T_1} \quad (\mathcal{E}/\text{et.fun}) \qquad \frac{\Gamma \vdash E_1 :^{\gamma_1} \Pi x : T_0. ^{\gamma_2} T \quad \Gamma \vdash E_0 :^P T_0}{\Gamma \vdash E_1 E_0 :^{\gamma_1 \sqcup \gamma_2} \{x \leftarrow E_0\} T} \quad (\mathcal{E}/\text{et.app})$$

Local binding In the rule $(\mathcal{E}/\text{et.app})$, the argument E_0 must be pure. In order to lift this restriction, one may use a local binding construct instead, as discussed in section IV.2.1.4. One may then no longer substitute E_0 in the result type, thus a premise of $(\mathcal{E}/\text{et.let})$ imposes that the result type does not mention the locally bound variable x — T is the type of the whole expression as well as the type of the body E . The effects of the expression are the union of that of E_0 and E . Since we only distinguish between two effects, and the case where both E_0 and E is not useful as function application can be used instead, we directly state the rule assuming the impure effect I throughout. With a richer effect system, we would use the least upper bound of the effects of E_0 and E as the result effect, although we would impose that this bound not be P in order to keep the pure fragment of the language as small as possible.

$$\frac{\Gamma \vdash E_0 :^I T_0 \quad \Gamma, x : T_0 \vdash E :^I T \quad \Gamma \vdash T \text{ ok}}{\Gamma \vdash (\text{let } x = E_0 \text{ in } E : T) :^I T} \quad (\mathcal{E}/\text{et.let})$$

If E is an impure expression of type $T_1 * T_2$, the second projection of E can be encoded as $\text{let } x = E \text{ in } \pi_2 x : T_2$. If E_0 is an impure expression and E has the type T_1 , then the application of E to E_0 can be encoded as $\text{let } x = E_0 \text{ in } E x : T_2$. Note that in both cases T_2 is not allowed to contain x : the type of E may not be dependent.

Subtyping The rule $(\mathcal{E}/\text{et.sub})$ combines implicit subtyping with implicit subeffectuation: any expression whose effects are constrained by γ has its effects constrained by any γ' such that $\gamma \sqsubseteq \gamma'$.

⁴Actually it would suffice to require E_1 to be pure, with effects allowed in E , however we will not have a use for this generalisation.

$$\frac{\Gamma \vdash E :^\gamma T \quad \Gamma \vdash T <: T' \quad \text{when } \gamma \sqsubseteq \gamma'}{\Gamma \vdash E :^{\gamma'} T'} \text{ (\mathcal{E}/\text{et.sub})}$$

Singletons We qualify the rule ($\mathcal{E}/\text{et.sing}$) to restrict singletons to pure expressions.

$$\frac{\Gamma \vdash E :^P T}{\Gamma \vdash E :^P S(E)} \text{ (\mathcal{E}/\text{et.sing})}$$

Sealing A new rule describes how to type a sealing. In general, in the expression $E !! T$, the “natural” type for E is a subtype of T , and the subtyping rule must be applied. The effects of $E !! T$ are those of E , plus the effect of the sealing; since we do not distinguish between effects, the effect annotation on $E !! T$ is always I .

$$\frac{\Gamma \vdash E :^\gamma T}{\Gamma \vdash (E !! T) :^I T} \text{ (\mathcal{E}/\text{et.seal})}$$

IV.4.4 Applicativity

IV.4.4.1 Applicative functors

Consider a functor whose body is sealed, i.e., $\lambda x : T_0. (E !! T)$. According to our description of sealing, each application of this functor causes the expression $E !! T$ to be evaluated, producing a fresh batch of abstract types. Therefore the types $\text{Typ } \pi_1 x_1$ and $\text{Typ } \pi_1 x_2$ in the following program are incompatible:

```
let f = λx : UNIT. ((⟨T'⟩, E) !! Σx : TYPE. T1) in
let x1 = f () in let x2 = f () in ...
```

This means that f is a **generative functor** (see section I.2.2.3). A functor whose body is sealed is always generative; this is reflected by its type, which has the form $\Pi x : T_0. {}^I T$ indicating that the application of the functor has a side effect (namely type creation). On the other hand, a functor whose body is pure (thus in particular not sealed) has a type of the form $\Pi x : T_0. {}^P T$, and does not participate in creating abstraction: it is a **transparent functor**.

Sometimes we would like for a functor to create abstraction, but for repeated applications of the functor to the same arguments to produce compatible abstract types. A typical example is a functor creating a data structure, where the arguments describe the elements of the data structure. Such functors are called **applicative functors** [Ler95]. We shall see two ways of supporting applicative functors.

One method is to add a new notion of sealing to the language, such that sealing the same module twice yields compatible results (in the sense that their abstract types are equivalent). This notion of sealing is called **weak sealing**, and we will write a weak sealing as $E :: T$. The form of sealing which always creates fresh abstract types, which we already know as $E !! T$, is called **strong sealing**. We will see how to formalise weak sealing in section IV.4.4.2. Using it, the types $\text{Typ } \pi_1 x_1$ and $\text{Typ } \pi_1 x_2$ in the following program are compatible:

```
let g = λx : UNIT. ((⟨T'⟩, E) :: Σx : TYPE. T1) in
let x1 = f () in let x2 = f () in ...
```

There is actually another way to build applicative functors without extending the language. One may seal the functor itself, rather than its body. Then all abstraction happens when the functor

is defined, and none when it is applied. For example, in the following program, g is an applicative functor, and the types $\text{Typ } \pi_1 y_1$ and $\text{Typ } \pi_1 y_2$ are compatible (we assume that E is pure).

```
let g = ( $\lambda x : \text{UNIT}. (\langle T' \rangle, E)$ ) !! ( $\Pi x : \text{UNIT}. \text{P}\Sigma x : \text{TYPE}. T_1$ ) in
let  $y_1 = g ()$  in let  $y_2 = g ()$  in ...
```

In either case, an applicative functor has a pure functor type $\Pi x : T_0. \text{P}T_1$, as opposed to the impure functor type $\Pi x : T_0. \text{I}T_1$ of a generative functor. This does not prevent an applicative functor from creating abstract types, if the result signature T_1 contains `TYPE` fields. The method for defining an applicative functor using strong sealing conveys an interesting idea about how abstract types are created: the side effect of creating the abstract types happens when the functor is defined, or more precisely when the functor is sealed, making an applicative functor from a transparent functor. If the functor is defined and sealed at the top level of the program, the side effect happens during program initialisation.

An applicative functor can be transformed into a generative functor at any time by sealing it to the appropriate generative functor signature (just like a transparent functor can be made applicative): $\Pi x : T_0. \text{P}T_1$ is a subtype of $\Pi x : T_0. \text{I}T_1$ (this is contained in the subtyping rule ($\mathcal{E}/\text{tsub.cong.fun}$)). The converse transformation, of an abstraction-creating functor to a less-abstracting functor (generative to applicative, or applicative to transparent), is undesirable, since there would be a loss of abstraction. Our effect system prevents such loss: applying a generative functor triggers an effect, and the only way to hide this effect is to wrap it in a lambda-abstraction whose type records the effect.

IV.4.4.2 Static sealing: formalisation $\boxed{\mathcal{W}}$

Introduction We shall define a new, “weak” notion of sealing such that sealing the same module twice produces compatible results. The purpose of this notion is to define applicative functors, and these provide a good way to understand weak sealing. We have seen how to seal a transparent functor to make it applicative: there is then an effect when the applicative functor is defined. The family of abstract types created by an applicative functor (which is indexed by the arguments of the functor) is fixed once and for all. We mentioned earlier that when encoding applicative functors using strong sealing, the effect happens during program initialisation. In fact the effect of weak sealing can be considered to happen at compile-time (more precisely during type-checking). Such a weak sealing is called **static sealing**, as opposed to the **dynamic sealing** $E !! T$. (We will discuss other weak forms of sealing in section IV.4.4.4). Let us now formally define static sealing.

Syntax We define system \mathcal{W} which extends system \mathcal{E} . The expression language has the new static sealing construct $E :: T$. Local binding now carries an effect annotation, whose meaning we will explain when commenting on the typing rule ($\mathcal{W}/\text{et.let}$); we will often omit this extra annotation in examples where it does not matter. The type language remains unchanged.

```
 $E ::=$ 
    expression (module)
    ...
    let  $x = E_0$  in  $E :^{\gamma} T$   local binding
     $E :: T$                   static sealing
```

The main novation is that the effect system is now larger.

$\gamma ::=$ **effect**
 P pure
 I dynamic effect
 S static effect
 IS both a static and a dynamic effect

The order relation on effects is given by $P \sqsubseteq I \sqsubseteq IS$ and $P \sqsubseteq S \sqsubseteq IS$. Only P and I may appear in function types, i.e., as the γ in $\Pi x : T_0. {}^\gamma T_1$.

The type system of \mathcal{W} is mostly identical with \mathcal{E} : the differences are confined to a few rules that feature impure expressions. Rules that are parametric over an effect instantiation may have the effect instantiated by S or IS , with (for $(\mathcal{W}/\text{et.sub})$) the extended order relation.

Typing a static sealing is similar to typing a dynamic sealing. In either case, the effect of the sealing is added to the effects of the function body. The effect of the sealing is either I for dynamic sealing or S for static sealing.

$$\frac{\Gamma \vdash E : {}^\gamma T}{\Gamma \vdash (E !! T) : {}^{\gamma \sqcup I} T} (\mathcal{W}/\text{et.seal.dyn}) \qquad \frac{\Gamma \vdash E : {}^\gamma T}{\Gamma \vdash (E :: T) : {}^{\gamma \sqcup S} T} (\mathcal{W}/\text{et.seal.stat})$$

Lambda-abstraction hides dynamic effects (although the type of the abstraction remembers the effect), but static effects always remain apparent. The functor result effect $\gamma \sqcap S$ can be seen as “the static part of γ ”, while the effect of the whole functor $\gamma \sqcup I$ is the dynamic part of γ . We recall the rule for application, which is unchanged, but do note that $\gamma_2 \sqsubseteq I$ always holds.

$$\frac{\Gamma, x : T_0 \vdash E : {}^\gamma T_1}{\Gamma \vdash \lambda x : T_0. E : {}^{\gamma \sqcap S} \Pi x : T_0. {}^{\gamma \sqcap I} T_1} (\mathcal{W}/\text{et.fun}) \qquad \frac{\Gamma \vdash E_1 : {}^{\gamma_1} \Pi x : T_0. {}^{\gamma_2} T \quad \Gamma \vdash E_0 : {}^P T_0}{\Gamma \vdash E_1 E_0 : {}^{\gamma_1 \sqcup \gamma_2} \{x \leftarrow E_0\} T} (\mathcal{W}/\text{et.app})$$

In system \mathcal{E} , we forced local binding expressions $\text{let } x = E_0 \text{ in } E : T$ to be impure, in order to keep the set of pure, hence comparable expressions small (as soon as E_0 is pure the expression can be written $(\lambda x : T_0. E) E_0$). We will keep doing this here; however we cannot simply force the expression to have the effect I since there are now other effects: if the expression should have the effect S , we do not want to force it to IS (nor “forget” the static effect and only keep I). We now require that the programmer specify the overall effect along with the type (because of the avoidance problem); since we do not want the expression to be pure we simply forbid this effect annotation from being P .

$$\frac{\Gamma \vdash E_0 : {}^\gamma T_0 \quad \Gamma, x : T_0 \vdash E : {}^I T \quad \Gamma \vdash T \text{ ok} \quad \text{when } \gamma \neq P}{\Gamma \vdash (\text{let } x = E_0 \text{ in } E : {}^\gamma T) : {}^\gamma T} (\mathcal{W}/\text{et.let})$$

Execution The dynamic semantics of system \mathcal{W} is the same as that of \mathcal{E} , with the rule $(\mathcal{W}/\text{ered.seal})$ indifferently accepting a dynamic or static sealing. Since this reduction simply erases the abstraction, the degree of generativity does not matter.

IV.4.4.3 Equivalences in the presence of static sealing

We saw in section IV.4.4.1 that dynamic sealing does not commute with functor abstraction: $\lambda x : T_0. (E !! T)$ and $(\lambda x : T_0. E) !! (\Pi x : T_0. {}^\gamma T)$ are not equivalent — if $\gamma = I$, they have the same type (they are both generative functors), but the first expression is pure while the second is impure. In contrast, static sealing commutes with functor abstraction: the expressions $E_1 = \lambda x : T_0. (E :: T)$ and $E_2 = (\lambda x : T_0. E) :: (\Pi x : T_0. {}^\gamma T)$ can be used interchangeably. ML (or at least Objective Caml) programmers often take this equivalence for granted. We will say that E_1 and E_2 are **equitypable**, meaning that for any Γ , γ' and T' , the typing judgement $\Gamma \vdash E_1 : {}^{\gamma'} T'$ is derivable if and only if $\Gamma \vdash E_2 : {}^{\gamma'} T'$ is. Equitypability will be the notion of interest in the present section, as we will look

at program fragments that have the same run-time behaviour apart from the usage of sealing, but are distinguishable at the typing level as they create abstraction in different amounts.

To prove that E_1 and E_2 are equitytypable, first note that both E_1 and E_2 require $\Gamma, x : T_0 \vdash E :^\gamma T$ to hold in order for each of them to be well-typed (by case analysis on their potential typing derivations). Then E_1 is typable by the rule ($\mathcal{W}/\text{et.seal.stat}$) followed by ($\mathcal{W}/\text{et.fun}$), while E_2 is typable by the rule ($\mathcal{W}/\text{et.fun}$) followed by ($\mathcal{W}/\text{et.seal.stat}$) (extra applications of ($\mathcal{W}/\text{et.sub}$) may be inserted, but they do not have significant impact as all operations involved are covariant in the result type). Both E_1 and E_2 have the principal type $\Pi x : T_0. \gamma \Pi T$ and the principal effect S .

Static sealing also commutes with other constructs. For example, $(E_1, E_2) :: (T_1 * T_2)$ is equitytypable with $(E_1 :: T_1, E_2 :: T_2)$, as well as with $(E_1 :: T_1, E_2)$ and $(E_1, E_2 :: T_2)$, assuming that each E_i has the type T_i . The key reason is that the presence of static sealing in any position makes the whole pair statically impure. Similarly the expressions $\pi_i(E :: T_1 * T_2)$ and $(\pi_i E) :: T_i$ are equitytypable when E has the type $T_1 * T_2$.

Let us now consider a local binding $E' = \text{let } x = E_0 \text{ in } E :^\gamma T$. Sealing only E is manifestly not equivalent to sealing E' , since the set of signatures that E or E' may be sealed to is different: only E may be sealed by a type mentioning x . However the difference is but of little importance, the reason being that the influence of sealing on effects is the same in both cases, viz., introducing S if sealing statically, or introducing I if sealing dynamically. As for the type of the expression, it remains T if E is sealed (assuming the whole expression remains well-typed), and it becomes some subtype of T if E' is sealed. In particular, $\text{let } x = E_0 \text{ in } (E :: T_1) : T$ is equitytypable with $(\text{let } x = E_0 \text{ in } E :^\gamma T) :: T$ as long as E has the type T_1 in the appropriate environment.

A sealing (whether dynamic or static) cannot appear in a function argument. These considerations show us that we do not lose expressivity if we limit the presence of static sealing inside a program to just two kinds of places: on a locally bound module $\text{let } x = (E_0 :: T_0) \text{ in } E :^\gamma T$ and on an applied functor $(E_1 :: T_2) E_0$. In the first case, removing the sealing could allow E to make use of a more precise type for x (any type of E_0 , not limited by E_0) — in other words the abstraction provided by the sealing would vanish with it. The usefulness of sealing an applied functor is of the same order: in order for $(E_1 :: T_2) E_0$ to be well-typed, T_2 must be a function⁵ $\Pi x : T_0. \gamma T_1$. If T_1 is smaller than necessary, the abstraction could migrate above the application (one could write $(E_1 E_0) :: \{x \leftarrow E_0\} T_1$ instead). If T_0 is larger than necessary, the choice of possible types of E_0 is limited, so E_0 becomes more abstract than as seen by the function. In fact, $(E_1 :: \Pi x : T_0. \gamma T_1)$ is equitytypable with $\text{let } x = (E_0 :: T_0) \text{ in } E_1 x :^\gamma T_1$. Hence, in summary, *static sealing is only useful on a locally bound module*.

IV.4.4.4 Other forms of sealing

We have so far formalised two forms of sealing: static sealing, which creates a new family of abstract types for each syntactic occurrence of the sealing operator, and dynamic sealing, which creates a new family of abstract types each time the sealing operator is evaluated. These two forms correspond to the *weak sealing* $E :: T$ and *strong sealing* $E :> T$ proposed by Dreyer, Crary and Harper [DCH03], and our effect system follows theirs⁶ (with one minor difference: they declare strong sealing as having a static effect as well as a dynamic effect, which uselessly strict but of little incidence).

Dreyer [Dre05] distinguishes between three forms of sealing:

⁵ T_2 could actually also be a singleton, but then the sealing would not be creating any abstraction.

⁶We use different notations however: we see effect annotations as indicating effects, whereas they see these annotations as purity annotations; thus we write S for a static effect where they write D for dynamic purity, and we write I for a dynamic effect where they write S for static purity.

- impure sealing $\text{impure}(E \text{ :> } T)$ is the strong sealing of Dreyer, Crary and Harper [DCH03] mentioned in the previous paragraph;
- separable sealing $E \text{ :> } T$ and inseparable sealing $\text{pure}(E \text{ :> } T)$ both correspond to our static sealing, being only distinguished by their separability, which we do not take into account (see section IV.4.2.3).

Many other variants can be conceived, with varying strengths. **Ascription** consists in constraining an expression to a type without restricting the view to it, i.e., without introducing abstraction: if E has the type T , the ascription $E \text{ :}_a T$ has any type that E has; in particular, if E is pure, then $E \text{ :}_a T$ is also pure and has the type $S(E)$. Ascription can be seen as a degenerate form of sealing.

Minimal sealing creates a comparable abstract type: $E \text{ :}_s T$ has the type T and the purity of E . Thus minimal sealing is a new way to construct pure expressions; two minimally sealed expressions are comparable. Minimal sealing cheerfully generates coincidental type equivalences, whenever the same expression happens to be sealed to the same type. A variant of minimal sealing consists in declaring $E \text{ :}_s T$ to be equivalent to $E \text{ :}_s T'$ whenever both are valid. Yet another variant consists in having a whole family of minimal sealings indexed by a name, considering two minimal sealings of the same expression to be equivalent if and only if they carry the same name.

Note that minimal sealing can be emulated using static sealing. All we need is a standard library function providing an applicative functor

$$f_{\text{minseal}} = \lambda t : \text{TYPE}. \lambda x : \text{STRING}. \\ ((t, ((\lambda x : \text{Type } t. x), (\lambda x : \text{Type } t. x)))) :: \\ \Sigma t' : \text{TYPE}. (\text{Type } t \rightarrow \text{Type } t') * (\text{Type } t' \rightarrow \text{Type } t)$$

or in Objective Caml syntax

```
let MinSeal = functor (A : sig type t end) ->
  struct type t = A.t let a x = x let c x = x end :
  sig type t val a : t->A.t val c : A.t->t end
end
```

For any type T and any name x , $f_{\text{minseal}} \langle T \rangle x$ provides an abstract type and conversion functions between that type and T . (In Objective Caml, we should define once `module $M_T^{\text{name}} = \text{struct type } t = T \text{ end}$` for every type T and name name since structures are generative.) This defines the named variant of minimal sealing; getting rid of x yields the basic variant. As we remarked in section IV.4.4.1, f_{minseal} could equally well be defined using dynamic sealing.

We have limited our exposition to two forms of sealing because, together with the easily definable minimal sealing, they seem to be sufficient for all practical purposes. We can roughly partition uses of sealing into three categories:

- abstract datatypes, in which abstraction enforces algebraic properties that go beyond the expressive power of the type system: static sealing is the perfect match;
- named variants of isomorphic types (e.g., `dollar` and `euro`): minimal sealing is suitable;
- abstract types used to limit access to some resource, which require dynamic sealing.

IV.4.4.5 Mutual encodings of static and dynamic sealing

Can we go even further and be content with a single form of sealing? The answer is “sort of”: while static and dynamic sealing can be encoded in terms of one another, a global program transformation is required either way.

Let us first express static sealing in terms of dynamic sealing. We start with two observations. Firstly, the two forms of sealing are equivalent if the sealing is executed exactly once. Secondly, we saw in section IV.4.4.3 that it suffices to study static sealing on locally bound expressions $\text{let } x = E_0 :: T_0 \text{ in } E :^{\gamma} T$.

We can see any program as a sequence of local bindings $\text{let } x_1 = E_1 :: T_1 \text{ in } \dots \text{let } x_k = E_k :: T_k \text{ in } E$ (we omit the return type on lets as they are not important here). We call E the body of the program, and the exposed local bindings are said to be toplevel. A program is said to be in *prenex sealings form* if none of the expressions E_1, \dots, E_k contain any static sealing. Finally a *clean* program is one in prenex sealings form where E contains no static sealing either, i.e., all static sealings are toplevel. In a clean program, static sealings can be replaced by dynamic sealings without affecting the typing of the program. We will show how to transform any program into an equitypable clean program.

Let $E_0 :: T_0$ be a subexpression of the program, so that the body of the program is $E_0 :: T_0$ in some context C , which we write $E = C \cdot (E_0 :: T_0)$. We can replace $E_0 :: T_0$ by $E'_0 = (\lambda y_1 : T_1. \dots \lambda y_j : T_j. E_0 :: T_0 \dots) y_1 \dots y_j$ where y_1, \dots, y_j are the variables bound by C from outermost to innermost, omitting toplevel bindings. If $j = 0$ we instead take $j = 1$, $T_1 = \text{UNIT}$ and $E'_0 = (\lambda y : \text{UNIT}. E_0 :: T_0) ()$. In every case, E'_0 is the application of a lambda-abstraction to one or more parameters. All free variables in this lambda-abstraction are bound at the toplevel. We can therefore extract it out of the context C in order to bind it above, going by beta-expansion and let lifting from $E = C \cdot (E_0 :: T_0)$ to $E' = (\text{let } f = E'_0 \text{ in } C \cdot (f y_1 \dots y_j))$ where f is a fresh variable. Now, as we saw earlier, static sealing can be lifted out of a lambda-abstraction, so E'_0 is equitypable with some expression $E''_0 :: T''_0$. Let $E'' = (\text{let } f = E''_0 :: T''_0 \text{ in } C \cdot (f y_1 \dots y_j))$. Provided E_0 itself contains no static sealing, E'' is in prenex sealings form, and the number of toplevel bindings has increased by 1.

Iterating the transformation we have just described over all of the static sealings in the initial program (from inside out), we can put any program in prenex sealings form. By replacing each toplevel static sealing with a dynamic one, we obtain an equitypable program that does not use static sealing. One intuitive view of this transformation is that each static sealing creates a single family of abstract types, and we lift the creation of this family to be performed exactly once during program initialisation.

We now turn to the dual problem of encoding dynamic sealing into static sealing. The difference between static sealing and dynamic sealing is the effect of the construction. One way to force a dynamic effect is to apply a generative functor, and a generative functor can be created by any sealing, including static sealing, of a transparent functor. This leads Dreyer, Cray and Harper [DCH03] to propose the following encoding of strong sealing (which is almost our dynamic sealing) into weak sealing (identical to our static sealing):

$$E :> T = ((\lambda x : \text{UNIT}. E) :: (\prod x : \text{UNIT}. {}^1 T)) ()$$

Unlike dynamic sealing $E !! T$, strong sealing $E :> T$ has a static effect in addition to its dynamic effect: $E :> T$ is equitypable with $E !! T :: T$. This effect cannot be discharged by a lambda-abstraction, so that generative functors cannot be pure in Dreyer, Cray and Harper’s system. Nonetheless we can move the extra static sealings to the toplevel by applying the transformation above.

In summary, static and dynamic sealing can be encoded in terms of one another, albeit incurring the cost of a global transformation. In the remainder of this document, we will study a language with only dynamic sealing (which we prefer due to its more compositional semantics). We propose to treat static sealing as an additional construct which should be provided in a programming language alongside dynamic sealing, and then elaborated away inside the compiler.

IV.4.4.6 On applicativity through functor sealing

The observation that the position of the sealing determines whether a functor is applicative or generative does not seem to be universally known in the ML community. It requires being able to seal a functor, whereas early module systems for ML only had sealing of structures. In Objective Caml, where functors are systematically applicative, it is customary to seal the body of a functor, and sealing the functor is considered equivalent (and needlessly complex as the type of the argument is then repeated) [Ler]. Note that if sealing is interpreted as an effect, the fact that it does not commute with lambda-abstraction is unsurprising.

Early module systems for ML only defined sealing on structures, and the possibilities of functor sealing seeped in slowly and with a low profile. Russo [Rus98] distinguishes applicative functors from generative functors by their definition rather than by their signature, with the defect that a generative functor can be directly seen as an applicative functor [Dre02] as seen in section I.2.2.4. Shao [Sha99] remarks in passing that sealing a transparent functor is a way of building an applicative functor. This possibility is also mentioned by Dreyer, Crary and Harper [DCH03, Dre05] but they recommend using weak sealing to build applicative functors.

One argument in favor of weak sealing ([DCH03], §2, p. 7) is that it can be applied to a single member of a structure in the body of the functor, which makes some types abstract already in following members, whereas sealing the functor only makes types abstract once the functor is applied. In our notation, the functors under discussion are of the form $\lambda x:T_0.\text{let } y = E_1 :: T_1 \text{ in } E_2 : T_2$. We saw in section IV.4.4.5 that this is precisely the case when the transformation of static sealing into dynamic sealing requires a global code reorganisation.

Dreyer [Dre05] (§1.2.7) mentions in particular the case of a functor whose body defines and uses a “generative” declared type (`datatype` in Standard ML, ordinary variant or record type in Objective Caml). If declared types are modelled by an abstract type obtained through dynamic sealing, such a functor is automatically generative. However we do not see any reason to insist on dynamic sealing: minimal sealing would do just fine, as the generative nature of declared types only serves to differentiate between the constructor and destructor names of different declared types. Since minimal sealing can easily be modelled by dynamic sealing, the lack of sealing other than dynamic is not a problem on this count.

IV.5 Colours and brackets C

In system \mathcal{E} , the sealing construct affects typing but not evaluation, as witnessed by the reduction rule for sealing which just forgets the sealing:

$$E \text{ !! } T \longrightarrow E \quad (\mathcal{E}/\text{ered.seal})$$

This rule is type-preserving in the sense that if the left-hand side is well-typed then the right-hand is also well-typed and has the same type. However information is clearly lost: this rule is not abstraction-preserving. This lack is no concern when evaluating a single program, as type preservation ensures that nothing can “go wrong” and the whole program text is available for any

If T is of the form ... $\Sigma t : \text{TYPE}. (\text{INT} \rightarrow \text{Typ } t)$ $\Sigma t : \text{TYPE}. \Sigma t' : \text{TYPE}. (\text{Typ } t * \text{Typ } t' * \text{INT})$ $\Pi x : T_0. \Sigma t : \text{TYPE}. (\text{Typ } t * \text{INT})$	then $E !! T$ reduces to an expression of type... $S(\pi_1 a) * (\text{INT} \rightarrow \text{Typ } \pi_1 a)$ $S(\pi_1 a) * S(\pi_1 \pi_2 a) * \text{Typ } t * \text{Typ } t' * \text{INT}$ $\Pi x : T_0. S(\pi_1(a x)) * \text{Typ } t * \text{INT}$
--	--

In each case, a is the nonce (fresh module identity) created by the sealing operation.

Figure IV.1: Examples of nonce generation

analysis that might rely on the typing of the program. However we aim to rid ourselves of the strict phase separation between type-checking and evaluation, by introducing a facility for type-checking. This facility requires additional information to remember the distinction between an abstract type and its representation type as long as it matters, which is as long as dynamic type-checking might be performed, i.e., throughout program evaluation. We will now study system \mathcal{C} , which is based on \mathcal{E} but where the reduction of a sealing preserves the abstraction.

IV.5.1 Module identities

IV.5.1.1 Nonce generation

The rule ($\mathcal{E}/\text{ered.seal}$) does not properly reflect our intention regarding the semantics of sealing. We described sealing as creating a new type. Consider the example sealing expression $(\langle \text{INT} \rangle, 3) !! \Sigma t : \text{TYPE}. \text{Typ } t$. It reduces by ($\mathcal{E}/\text{ered.seal}$) to $(\langle \text{INT} \rangle, 3)$, which has the type $\Sigma t : S(\langle \text{INT} \rangle). \text{Typ } t = S(\langle \text{INT} \rangle) * \text{INT}$, whereas we would like a type of the form $\Sigma t : S(\langle T' \rangle). \text{Typ } t = S(\langle T' \rangle) * T'$ where T' is distinct from any previously existing type (especially INT).

More precisely, we do not need to create a type but a *module identity*, as can be seen by looking at the sealing of modules with a more complex structure. For example, if the same sealed module defines several abstract types, these types share a common unique identity. If the sealed module is a functor, a new identity must be created but once when the sealing construct is evaluated, and this identity will be used each time the functor is applied. Each module identity characterises one instantiation of the abstraction, which may produce any number of fresh types: as many as there are abstract type fields for a structure, an unbounded number for a functor (since the argument must be taken into account)...

Figure IV.1 shows a few examples of uses of module identities. A fresh module identity is called a **nonce** (or (h)apax), and written a . These nonces have the same universal uniqueness property as those used in models of security. They generalise the *stamps* of MacQueen [AM91]. Unlike the stamps of many module systems, our nonces can designate modules of arbitrary signatures (in particular functors). Also the creation of a nonce is performed when the sealing construct is evaluated, and not whenever a module is built. Nonces correspond to the singularised identities described in section II.6.1.2.

At the syntactic level, we will for the time being consider nonces an extra construct in the syntax of modules, which should not appear in source programs⁷. However the purpose of nonces is to designate abstract types, and we will eventually restrict their presence to “type components” (see section IV.5.1.4). We assume an infinite supply of distinct nonces (similar to the infinite supply of variable names).

⁷Note that (as we shall see) nonce-free programs have nonce-free types, so types of source programs are expressible in the source language.

IV.5.1.2 Lexes

Evaluating a sealing construct $E !! T$ requires a fresh nonce, i.e., one that is not present in the original term. We manage this freshness requirement by using a store of nonces, called a **lexis** (or stamp book), and written B . A lexis keeps track of nonces in use as well as the module implementation and signature from the sealing construct from which the nonce originated. The syntax of a lexis is thus

$$B = (\alpha_1 = E_1 : T_1, \dots, \alpha_k = E_k : T_k)$$

As with environments, we treat lexis concatenation as associative, and the empty lexis (written nil) is a neutral element for this operation.

Lexes adorn evaluation judgements as well as typing judgement. A nonce α is well-typed and has type T when the ambient lexis contains the binding $\alpha = E : T$ for some E (just as a variable x has the type T when the environment contains the binding $x : T$). The reduction relation for system \mathcal{C} is formally defined on pairs consisting of a lexis and a type; a reduction will be written $B \vdash E \longrightarrow B' \vdash E'$. However in most instances the lexis does not affect the reduction and does not change, and we will then continue to write $E \longrightarrow E'$. Reducing a sealing augments the lexis:

$$B \vdash E !! T \longrightarrow B, \alpha = E : T \vdash E'$$

The nonce α is chosen fresh, i.e., outside the domain of B . Since the language does not include binders for nonces, any nonce appearing in E or T must be recorded in B .

An alternative to this global store would be to introduce a “new” binder for nonces, classically writing $(\nu \alpha = E_0 : T_0)E$. We prefer the use of a global store not only because we have no need for a nonce binder, but also because managing the migration of ν binders accross other syntax nodes and above environments⁸ would be problematic.

IV.5.1.3 From sealing to brackets

We have seen that reducing a sealing $E !! T$ creates a fresh nonce α , and the result is an expression E' of a type T' , where T' uses the nonce α as the identity of the newly created module in order to precisely specify the abstract parts of T .

This type T' is called a **strengthening** of T , or more precisely it is what we will call the **selfification** of the type T for the module identity α (see section I.2.2.2). We will write this selfification as $\mathbf{self}^T(\alpha)$. The general idea behind selfification is to mirror the original structure of the type, but replace the parts originally left abstract by a reference to the newly created nonce. Figure IV.1 shows a few examples of selfification; we will defer the task of formulating a precise definition until section IV.5.1.5.

Sealing must transform the expression E of type T into an expression E' with essentially the same behaviour as E but a different type $\mathbf{self}^T(\alpha)$. This new type is more precise than T : it is a subtype of T . Although T is usually not the most precise type of E , E cannot have the type $\mathbf{self}^T(\alpha)$ in general: if the same expression is sealed twice, the resulting expressions E'_1 and E'_2 should have the respective incompatible types $\mathbf{self}^T(\alpha_1)$ et $\mathbf{self}^T(\alpha_2)$. Thus E' must contain a reference to the specific choice of nonce α .

The most obvious way to construct E' is to start with E and apply a type coercion to it: $E' = \text{coerce } E \text{ to } \mathbf{self}^T(\alpha)$. However how this coercion should interact with the rest of the language is not obvious. What does an expression of the form $\text{coerce } E \text{ to } T'$ reduce to? With such little information in a readily extractible form, when is $\text{coerce } E \text{ to } T'$ well-typed?

⁸A nonce may appear in the type of a bound variable.

We were already confronted with this problem in the simpler setting of the HAT language, as discussed in III.1.1. We will use the same solution as then, to wit, **coloured brackets** [ZGM99]. The expression $[E]_{\alpha}^{\mathbf{self}^T(\alpha)}$ denotes the coercion of E to the type $\mathbf{self}^T(\alpha)$, but records the justification α for equating the implementation E with the abstraction. (Recall that E is recorded in the entry for α in the ambient lexis.) More generally, if E_1 is any expression, and if T_2 is equivalent to the type of E_1 modulo the equivalence between α and its implementation, the expression $[E_1]_{\alpha}^{T_2}$ has type T_2 . This expression is called the coloured bracket (or brackets) surrounding E_1 and annotated with the colour α and the type T_2 .

The expression E_1 should be seen as *inside* the bracket, while the annotations α and T_2 are carried *by* the bracket. Nonces (or fresh module identities) α have the same rôle as the hashes (or structural module identities) h in HAT. The colour α indicates the possibility of using the extra typing equality between α and its implementation E inside the bracket. We say that α is transparent inside the bracket. We will discuss the syntax and semantics of colours in system \mathcal{C} more fully in section IV.5.2.

The reduction rule for sealing is thus as follows:

$$B \vdash E !! T \longrightarrow B, \alpha = E : T \vdash [E]_{\alpha}^{\mathbf{self}^T(\alpha)}$$

As in HAT, the next reduction steps are devoted to pushing the coloured brackets towards the inside of E .

IV.5.1.4 Abstract types

In the expression $[E]_{\alpha}^{\mathbf{self}^T(\alpha)}$, the nonce α can only appear in two positions: as the colour annotation on the bracket, and in building the type annotation on said bracket. It seems therefore possible to restrict the places where nonces may appear in the syntax. But do we have to?

Treating a nonce α as a full-fledged expression definitely simplifies the overall language structure. Under a lexis containing the binding $\alpha = E : T$, the expression α has the type T , and the transparency of α can be simply expressed by the conversion $\alpha \longrightarrow E$.

Such uninhibited use of nonces nonetheless causes two problems, one theoretic, one practical. Both problems arise from reducing certain expressions containing nonces.

Consider for instance the sealing $E !! T$ where $E = (\langle \text{INT} \rangle, (2, 3))$ and $T = \Sigma t : \text{TYPE}. \text{Typ } t * \text{Typ } t$, resulting in the lexis binding $\alpha = E : \Sigma t : S(\pi_1 \alpha). \text{Typ } t * \text{Typ } t$. While it may be reasonable to treat such expressions as α , maybe even $\pi_1 \alpha$, as values — although $\pi_1 \alpha$ has a destructor at the head, which is odd in a value — the same does not go for $\pi_2 \alpha$. Typing excepted, the expected behaviour of this expression is the same as $(2, 3)$. Evaluating $\pi_2 (E !! T)$ yields the expression $[(2, 3)]_{\alpha}^{\text{Typ } \pi_1 \alpha * \text{Typ } \pi_1 \alpha}$. We might reduce α to $[E]_{\alpha}^{\Sigma t : S(\pi_1 \alpha). \text{Typ } t * \text{Typ } t}$ and then let bracket pushing do the rest; however, for type preservation, this expression must still have the original type $S(\alpha)$.⁹ Experience with a preliminary version of this system shows that obtaining such a typing requires a substantially more complex metatheory as analysing the types that are equivalent to $\text{Typ } \pi_1 \alpha$ becomes intractable.

On the practical side, such a reduction requires the implementation of α to be available at any time, whereas a nonce is intuitively an opaque piece of data on which only an equality predicate is defined. In particular, the implementation of a nonce might be cryptographically hidden, as we will see in section IV.5.2.1.

We saw that the principle of selfification is to mirror the original structure of the type while substituting the appropriate projection of the nonce for abstract type fields (TYPE) in the signature.

⁹Note that α must be pure, since it is meant for use in types.

Therefore the nonce α only appears in $\mathbf{self}^T(\alpha)$ as a projection of a field for which the signature indicates `TYPE`. Such a projection has the form `Typ A` where A may be a nonce, a pair projection $\pi_i A$, or the application of a functor to an argument $A E_0$. A term A of the grammar we just describe shall be called a **module component**.

From now on, we will not allow α to be an expression, preferring to add a new entry to the syntax of types. A **component type** $\langle A \rangle$ is a type, denoting what we would have written `Typ A`. Where the expression A was desired, we can now write $\langle \langle A \rangle \rangle$ (provided that the projection goes all the way to a single field of type `TYPE`).

Expressing the transparency of a nonce takes on a more complicated form since the equivalence $\alpha \equiv E$ (where E is the implementation of α) is no longer grammatical. Transparency must be expressed separately for each component type. Let α be the underlying nonce of the module component A_1 (which we shall write as $\alpha = \mathbf{under1}(A_1)$). When α is transparent, $\langle A_1 \rangle$ is equivalent to `Typ E1` where E_1 is the projection of E with the same shape as A_1 . We will say that A_1 is **revealed** as E_1 , which we write $E_1 = \mathbf{reveal}^B(A_1)$ (where B is the ambient lexis). For instance, if $T = \Sigma t_1 : \text{TYPE}. \Sigma t_2 : \text{TYPE}. \Sigma t_3 : S(\langle \text{INT} \rangle). T_4$, then transparency of α entails the equivalences $\langle \pi_1 \alpha \rangle \equiv \text{Typ } \pi_1 E$ and $\langle \pi_1 \pi_2 \alpha \rangle \equiv \text{Typ } \pi_1 \pi_2 E$.

IV.5.1.5 Selffication

The core of the sealing operation consists in replacing the abstract components of a type by the corresponding projections of a certain nonce. Generally speaking, let us study the selffication of a type T for a module component A , written $\mathbf{self}^T(A)$.

Base cases There are three kinds of elementary signatures: manifest type fields $S(\langle T_1 \rangle)$ (`sig type t = T1 end`), abstract type fields `TYPE` (`sig type t end`), and term fields (`sig val x : T2 end`). The purpose of selffication is to transform abstract type fields into manifest type fields: the selffication of `TYPE` by A is $S(\langle \langle A \rangle \rangle)$. Type fields that were already manifest (i.e., $S(\langle T_1 \rangle)$) are left unchanged, as are term fields since no extra information is required. Type fields thus always gain singleton types, while term fields are unchanged.

Structures Selffying a structure conserves its decomposition into fields, with each field selffied

separately. For example, in ML notation,

```

sig
  type t1
  type t2
  type u = int * t1
  val f : int -> t1 -> u
end

```

selffied by the name

```

M is
sig
  type t1 = M.t1
  type t2 = M.t2
  type u = int * t1
  val f : int -> t1 -> u
end

```

Note that the name of the module appears more than once: the same name M is used as part of the global designation of all the abstract types in the signature.

Selffying a pair $T_1 * T_2$ by a component A naturally yields the pair $\mathbf{self}^{T_1}(\pi_1 A) * \mathbf{self}^{T_2}(\pi_2 A)$. A natural generalisation to dependent pairs is achieved by independently selffying each component, thus $\mathbf{self}^{\Sigma x : T_1. T_2}(A) = \Sigma x : \mathbf{self}^{T_1}(\pi_1 A). \mathbf{self}^{T_2}(\pi_2 A)$.

It is tempting to try to go further: since x is now fully known, why not substitute it in T_2 ? Thus in the example above references to $t1$ and u could become $M.t1$ and $M.u$ respectively. However this is not possible in our language, which includes signatures that cannot be expressed in ML — to wit, dependencies (and in particular equalities) on term fields. If the first component of a pair contains term fields, its signature does not become a singleton after selfification, and the effort to specialise the second component cannot proceed further. For example the selfification of $\Sigma x : \text{INT}. S(x)$ by a is $\Sigma x : \text{INT}. S(x)$, no more. The selfification of a dependent pair is therefore still dependent in general.

Functors The notation `TYPE` admits two radically different interpretations, depending on whether it is used in the argument or in the result of a functor. In the argument, `TYPE` denotes a type that will remain unknown until the functor is applied, and will vary from application to application: a functor with a `TYPE`-qualified argument is polymorphic. Thus selfifying a functor type $\Pi x : T_0. \gamma T_1$ does not restrict the domain of the arguments of the functor¹⁰: the selfification has the form $\Pi x : T_0. \gamma T'_1$. As for the interpretation of `TYPE` in the result type of the functor, it depends on whether the functor is applicative or generative. With an applicative type, the effective identity of the type field is fully determined by the argument passed to the functor, and selfification consists of eta-expanding the body as $\Pi x : T_0. \gamma \mathbf{self}^{T_1}(A\ x)$. On the other hand, if the functor is generative, each application generates a new identity; the selfification operation must then be delayed until the application is performed, and T_1 must remain abstract for the duration.

Definition Selfification is defined by structural induction on the type as follows:

$$\begin{aligned} \mathbf{self}^{\text{BT}}(A) &= \text{BT} && \text{if BT is a base type (UNIT, BOOL, INT)} \\ \mathbf{self}^{\Sigma x : T_1. T_2}(A) &= \Sigma x : \mathbf{self}^{T_1}(\pi_1 A). \mathbf{self}^{T_2}(\pi_2 A) \\ \mathbf{self}^{\Pi x : T_0. P T_1}(A) &= \Pi x : T_0. P(\mathbf{self}^{T_1}(A\ x)) \\ \mathbf{self}^{\Pi x : T_0. I T_1}(A) &= \Pi x : T_0. I T_1 \\ \mathbf{self}^{S(E')}(A) &= S(E) \\ \mathbf{self}^{\text{TYPE}}(A) &= S(\langle\langle A \rangle\rangle) \end{aligned}$$

Evaluation of type fields The table above does not state how to compute $\mathbf{self}^{\text{Typ } E}(A)$. The reason for this omission is that the computation depends entirely in the value of E : the definition of selfification cannot be purely syntactic. Intuition whispers that the selfification of two equivalent types at the same identity should be equivalent; this implies for instance $\mathbf{self}^{\text{Typ } \langle T \rangle}(A) = \mathbf{self}^T(A)$. Therefore E must be reduced to a value, which will have the form $\langle T \rangle$ for typing reasons, in order to compute the selfification of `Typ E`.

IV.5.2 Colors

IV.5.2.1 Colouring

Coloured bracket According to their introduction in section IV.5.1.3, a coloured bracket $[E]_a^T$ lets the expression E be given the type T using the typing equations resulting from the knowledge of the implementation of a . As in `HAT`, we will have a wider view of coloured brackets as the syntactic manifestation of a boundary between the inside and the outside of the module. This boundary takes the form of the sealing construct in the source code, and coloured brackets embody it at run time.

¹⁰The selfification is a functor $\Pi x : T'_0. \gamma T'_1$ per conservation of the overall structure.

Coloured syntax One way to describe the syntax of system \mathcal{C} from the syntax of \mathcal{E} would be to associate a colour to each node of the syntax. This colour would represent the origin of that node, i.e., from which sealed module the node comes from. The brackets provide this information in a different form: the colour of a syntax node is that carried by the innermost surrounding bracket.

In the absence of a surrounding bracket, the colour is the **ambient colour** carried by the judgement in which the term under consideration is placed. The final form of an expression typing judgement in system \mathcal{C} includes a colour annotation:

$$B; \Gamma \vdash_c E :^Y T$$

So does the final form of an expression reduction rule:

$$B \vdash E \longrightarrow_c B' \vdash E'$$

Colours appear in other places in the syntax. The rule of thumb is that anywhere a type is attributed to an expression, a colour must also be provided. For instance a lexis binding has the form $(a = E :_c T)$, and an environment binding has the form $(x :_c T)$.

Colour semantics In a typing judgement, the colour determines which abstract types may be revealed. Following the intuition outlined in section IV.5.1.4, they are the components of transparent nonces in the indicated colour. A reduction rule ($\mathcal{C}/\text{ered.colAbs}$) allows the revelation of transparent nonces.

Reduction of a sealing We can finally state the rule for sealing reduction in full. Let us first consider our usual example consisting of the module `struct type t = int let x = 3 end` sealed to the signature `sig type t val x : t end`.

$$\begin{aligned} \text{nil} \vdash (\langle \text{INT} \rangle, 3) !! (\Sigma t : \text{TYPE. Typ } t) &\longrightarrow \bullet \\ \mathbf{a} = (\langle \text{INT} \rangle, 3) : \bullet (\Sigma t : \text{TYPE. Typ } t) \vdash [(\langle \text{INT} \rangle, 3)]_{\mathbf{a}}^{\Sigma t : S(\langle \langle \pi_1 \mathbf{a} \rangle \rangle). \text{Typ } t} \end{aligned}$$

Subsequent reduction steps push the brackets towards the inside of the value, as in HAT (see section III.1.2.2).

In general a sealing $V !! T$ is evaluated in the ambient colour c to $[V]_{c'}^{T'}$, where $c' = c \cup \{\mathbf{a}\}$ et $T' = \mathbf{self}^T(\mathbf{a})$, with \mathbf{a} being a fresh nonce. Thus the reduction rule ($\mathcal{C}/\text{ered.seal}$)¹¹ is as follows:

$$B \vdash V !! T \longrightarrow_c B, \mathbf{a} = V :_c T \vdash [V]_{c \cup \{\mathbf{a}\}}^{\mathbf{self}^T(\mathbf{a})}$$

The ambient colour c may be necessary just to ensure that V has the type T ; for V to have the type $\mathbf{self}^T(\mathbf{a})$ requires the transparency of \mathbf{a} in addition¹². A colour can thus be a (finite) set of nonces, which we shall write as $c = \{\mathbf{a}_1, \dots, \mathbf{a}_k\}$. The semantics of a colour is to render its elements transparent. Until now, we had only seen singleton colours $\{\mathbf{a}\}$, abbreviated as \mathbf{a} ; we call such colours **primary colours**. By synecdoche we will also call an element of a colour a primary colour. The **empty colour**, for which we will prefer the notation \bullet , makes no nonce transparent.

¹¹Specialised to an uncoloured sealing — see section IV.5.3.6.

¹²This situation is possible because a sealing can be embedded inside another, possibly with a functor interposed which prevents from pushing the brackets induced by the outermost sealing before reaching the evaluation of the functor body containing the innermost sealing. This scenario did not arise in HAT where module definitions were always sequential.

Transparency A nonce is said to be **transparent** or **opaque** (in a colour c , often obvious from context) depending on whether it is, or is not, an element of the colour. (This definition will be generalised under a more semantic form for a larger class of colours in section IV.5.2.3.)

Colour weakening Reducing a sealing moves the sealed expression from a colour to a larger colour. Intuitively, this should not cause any typing trouble: the larger colour provides more typing equations, so more typings are possible. We shall indeed state a **colour weakening** lemma: if E has the type T in the colour c , and c' is a well-formed colour containing c , then E has the type T in c' .

Border colour In a coloured bracket $[E]_{c'}^T$, the type annotation T lives on the border between the inner colour c' and the outer (ambient) colour c . We can formalise this by requiring that T be valid in both c and c' . By the colour weakening lemma, it suffices that T be valid in $c \cap c'$, and we will use this requirement in typing rules.¹³

Colours and security Let us briefly mention the security interpretation of coloured brackets. Colours can be seen as capabilities handed to expressions — in our application these capabilities unlock type equations. The brackets mark and maintain the boundaries of privileged chunks of code. Nonces are the usual nonces of abstract cryptography. This interpretation was first formulated early in the history of coloured brackets [PS00] and has been studied, in particular, under the name λ_{seal} [SP04].

IV.5.2.2 Semantics of a type and dependencies of a nonce

Semantics of a type in a colour A colour denotes equalities between types, so that there may be more than one way to express a type in a given colour. For instance, if the ambient lexis contains $\alpha_1 = ((\text{INT}), E_1) : \bullet \Sigma t : \text{TYPE}. T_1$, then $(\pi_1 \alpha_1)$ is equivalent to INT in the colour $\{\alpha_1\}$ but not in the empty colour \bullet . This possibility for a term to be a valid type in different colours with different semantics (in terms of what expressions have that type) is the key to the expressivity of coloured brackets, as the type annotation is considered in both the inner and the outer colour. Note that the set of terms having a given type is a monotone function of the ambient colour.

Semantics vs. validity Since a type may contain embedded expressions, the very validity of a type (and not just its semantics) may depend on the colour. In the lexis above, the type $S((\lambda x : (\pi_1 \alpha_1). x) 3)$ is only valid in a colour that makes α_1 transparent. Our type system can apparently accomodate this phenomenon, simply by virtue of annotating every typing judgement with a colour, including type correctness $B; \Gamma \vdash_c T \text{ ok}$.

Lexis binding Consider in particular the reduction of a sealing $B \vdash V !! T \longrightarrow_c B' \vdash [V]_{c \cup \{\alpha\}}^{\text{self}^T(\alpha)}$. Since V and T are only known to be valid in the colour c , c must be recorded together with V and T in B' , hence $B' = B, \alpha = V :_c T$. The elements of c are called the **dependencies** of the nonce α .

¹³Using $c \cap c'$ rather than replicating a premise in the colour c and c' has technical advantages, mainly in that just because the judgements $B; \Gamma \vdash_c J$ and $B; \Gamma \vdash_{c'} J$ are derivable does not automatically mean their derivations have the same shape. We suspect that a common shape can be found, and would work equally well to derive $B; \Gamma \vdash_{c \cap c'} J$, but proving such a result looks very difficult in our syntactic approach. Furthermore, given the presence of variables in colours (see section IV.5.2.3), we would have to use a semantic intersection that operates on the transparent closures of the colours (see section B.1.2).

Colour well-formation If the nonce a is used (as part of a type $\langle A \rangle$) in some colour, the equivalences between projections of a and of its implementation V must hold. Hence the colour in question must contain c (the only colour where V is known to be valid). Therefore a nonce can only be transparent if its dependencies are also transparent. We express this constraint in a side condition in the rule for colour well-formation ($(\mathcal{C}/\text{envok.c.a})$). Another option would be to automatically make all dependencies transparent, so that $\{a\}$ and $\{a\} \cup c$ would have the same semantic. We choose the restriction on well-formation as simpler and because it is similar to the condition for using a nonce in an expression.

Use of a nonce in an expression If the nonce a is used in an expression, its type T must be valid. Therefore the ambient colour must contain c (the only colour where T is known to be valid). Therefore a nonce can only be used if its dependencies are transparent, which is expressed in a side condition in the rule $(\mathcal{C}/\text{ac.a})$.

Concretisation We might attempt another approach to uses of nonces in an expression. Instead of giving the nonce a the type T taken directly from the sealing expression that created a , we might create a type T' that is equivalent to T in colours containing c but is valid in any colour. In the colour c , we can replace any dependency of a by its implementation. The type resulting from performing all such possible replacements meets the stated requirement. This operation is called **concretisation** of the type T for the colour c . Concretisation is also mutually recursively defined on types, expressions and module components. Concretisation is a copy function, except for the following cases:

$$\begin{aligned} \mathbf{conc}_c^B(\langle A_1 \rangle) &= \text{Typ reveal}^B(A_1) \quad \text{if } \mathbf{underl}(A_1) \in c \\ \mathbf{conc}_c^B(\langle A_1 \rangle) &= \langle A_1 \rangle \quad \text{if } \mathbf{underl}(A_1) \notin c \\ \mathbf{conc}_c^B([E]_{c'}^T) &= [E]_{c'}^{\mathbf{conc}_{c'}^B(T)} \\ &\quad (\text{other cases by simple induction}) \end{aligned}$$

We get $B; \Gamma \vdash_{c'} T \equiv \mathbf{conc}_c^B(T)$ as soon as c' contains c .

The advantage of concretisation is to allow the nonce a (with dependencies c) to be used in any colour c' , whether or not $c \subseteq c'$ holds. This is achieved by giving a the type $\mathbf{conc}_c^B(T)$. Unfortunately, when c is not included in the ambient colour, $\mathbf{conc}_c^B(T)$ is generally not equivalent to T even if the latter is well-formed, which is somewhat confusing. Forcing concretisation instead of restricting nonce use to a suitable colour does not on balance simplify the system design, which is why we eschew it here.

Seal-time concretisation Could we concretise the type of a nonce when the nonce is created, rather than when the nonce is used? The reduction rule for sealing would look like the following:

$$B \vdash V !! T \longrightarrow_c B, a = [V]_c^{\mathbf{conc}_c^B(T)} : \mathbf{conc}_c^B(T) \vdash [V]_{c \cup \{a\}}^{\text{self } T(a)}$$

The type of the sealed expression is then universal, i.e., valid in any colour. A lexis binding is also universal, and so need not be annotated with a colour. The slight loss of expressivity triggered by the forced concretisation is thus compensated by a simplification of the type system.

Unfortunately we are here using coloured brackets outside their operating parameters. Usually the type T will contain some abstract components (otherwise the sealing is useless). However the point of selfification was to replace these abstract components in the type annotation on the bracket by a manifest type (using the abstract module identity a). In fact, the type annotation on coloured brackets must be monomorphic, i.e., completely specified, free of TYPE fields (section IV.5.3 will discuss the concept further).

IV.5.2.3 Variables in colours

We saw that in HAT substitution of a value of type T for a variable assumed to have the type T does not always result in a well-typed term. An additional hypothesis is required stating that the value have the type T in any colour at which the variable is used. One way to ensure that this additional requirement is met is to associate a colour to each variable, which will be the colour of the syntax node at which the variable is bound, and then only allow using the variable in this colour. Then type preservation by substitution only require that the substituted-in value have the right type in the colour of the variable. Reduction rules affecting colours push brackets inside data constructors; they do not affect the colours of variables inside reduced terms (the case of pushing a bracket inside a lambda-abstraction will be discussed in section IV.5.3.3). Compared with HAT, we gain the advantage that beta-reduction (`@/ered.app`) does not mention any bracket. Annotating variable binding sites with colours is consistent with the principle stated above that anything that is attributed a type is also attributed a colour. However we shall see that confining variables to a colour is not, in itself, sound, and requires additional machinery to make the system sound.

As usual, the colour of a variable binding is given by the innermost surrounding bracket, or in the absence of one by the surrounding colour. A variable binding in an environment records that colour, so that an environment binding in system \mathcal{C} has the form $(x :_c T)$.

The simplest way to state the variable typing rule would be $x :_c T \vdash_c x :^P T$, i.e., the variable x is (only) usable in its colour of definition. In order for colour weakening to hold, this condition needs to be relaxed to allow using x in any colour containing the colour of definition, i.e.,

$$x :_c T \vdash_{c'} x :^P T \quad \text{when } c \subseteq c'$$

Unfortunately this formulation does not suffice to ensure that colour weakening holds, as shown by the following example:

$$B; \text{nil} \vdash_c (\lambda x : \text{INT}. [[x]_c^{\text{INT}}]_{c_1}^{\text{INT}}) :^P \text{INT} \rightarrow \text{INT}$$

Provided that B , c and c_1 are well-formed, this typing judgement is correct. There is no requirement for the intermediate colour c_1 to have any connection with c . If c is widened to some colour c' (such that $c \subseteq c'$), the judgement above becomes

$$B; \text{nil} \vdash_{c'} (\lambda x : \text{INT}. [[x]_c^{\text{INT}}]_{c_1}^{\text{INT}}) :^P \text{INT} \rightarrow \text{INT}$$

The colour of the innermost bracket remains c , as there is no indication that it should change during weakening. As a result, the term is no longer well-typed — the variable x is now used in a colour that is smaller than its colour of definition.

One way to perceive this problem is to consider the colour of a variable occurrence as explicitly bound to the colour of the binding of the variable, as opposed to these colours merely having to always being related. One might say that occurrence colours must be computed by name rather than by value. Each variable is then assigned a “symbolic primary colour” (contrast with nonces as constant primary colours). This new primary colour is written identically with the variable. A colour therefore has the form $c = \{a_1, \dots, a_k, x_1, \dots, x_n\}$, i.e., a finite set of nonces and colours¹⁴. A variable can only be used if it is present in the ambient colour, which gives us the following rule:

$$x :_c T \vdash_{c'} x :^P T \quad \text{when } x \in c'$$

A further point to note is that the colour c in which x is defined may itself contain other variables — and it may of course contain nonces. These variables and nonces will automatically be considered

¹⁴Another, symbolic, notation might be $c = \{a_1, \dots, a_k\} \cup \text{col}(x_1) \cup \dots \cup \text{col}(x_n)$ where the symbol `col` notes the colour of definition of its argument and the symbol \cup is interpreted as set union.

transparent whenever x is. In other words, if $x \in c'$, then any element of $c \cup \{x\}$ is transparent in c' . We note this by the typing judgement $x :_c \top \vdash_{c'} c \cup \{x\}$ transparent. The rule for using a variable in an expression is finally

$$\Gamma_0, x :_c \top, \Gamma_1 \vdash_{c'} x :^P \top \quad \text{when } \Gamma_0, x :_c \top, \Gamma_1 \vdash_{c'} x \text{ transparent}$$

Since variables now occur in colours, colours are subject to alpha-conversion and substitution. Additionally, a substitution must now include a target colour along with a target expression for the variable: the substitution of x by E in the colour c will be written $\{x \leftarrow_c E\}$. The following example illustrates the interaction of substitutions and coloured brackets, assuming that $x \in c$:

$$\{x \leftarrow_{c_0} E_0\}[x]_c^{S(x)} = [E_0]_{(c \setminus \{x\}) \cup c_0}^{S(E_0)}$$

IV.5.2.4 Absolute brackets, additive brackets

A coloured bracket $[E]_c^T$ lets the expression E (of colour c') be used in any ambient colour c . This is a special case of a relation between the inner colour c' and the outer colour c . Let \mathcal{R} be any binary relation on colours. We can write $[E]_{\mathcal{R}}^T$ for a coloured bracket with relation \mathcal{R} , which is well-typed in the ambient colour c if and only if there exists a colour c' such that E has the type T in c' and $(c, c') \in \mathcal{R}$ (we assume throughout this discussion that T is valid in the outer colour).

When \mathcal{R} can be an arbitrary relation, an important piece of information is lost, as the inner colour becomes ambiguous. We will therefore limit our analysis to the case where the inner colour is a (partial) function of the outer colour: $[E]_f^T$ is well-typed in the ambient colour c if and only if E has the type T in $f(c)$. Furthermore an important property that ensures that colour weakening will hold is that any widening of the inner colour matches a widening of the outer colour, in other words f must be monotone increasing ($c_1 \subseteq c_2$ implies $f(c_1) \subseteq f(c_2)$).

The coloured brackets that we have seen so far correspond to the case where the function f is total and constant. Such brackets are known as **absolute brackets**. At this point the only way to introduce a bracket in an expression is the reduction of a sealing, in which the bracket has some arbitrary outer colour c_0 and an internal colour of the form $c_0 \cup \{a\}$. Rather than an absolute bracket, we could use an **additive bracket** whose relation is $c \mapsto c \cup \{a\}$ (a total, one-to-one function from the inner colour to the outer colour, thus a (partial) one-to-one function from the outer colour to the inner colour¹⁵). We write such an additive bracket as $[E]_{+a}^T$.

Additive brackets blend in nicely with the rest of the language. In particular occurrences of variables under additive brackets trigger none of the problems discussed in section IV.5.2.3: in an expression such as $\lambda x:T. [[x]_{+a_2}^T]_{+a_1}^T$, the colour of the bound occurrence is automatically a superset of the binding colour. Hence we could dispense with the additional complexity resulting from having variables in colour, provided additive brackets were sufficient for our purposes.

Unfortunately, additive brackets are not expressive enough. They can never restrict the set of typing equations accessible in a subexpression, and in particular do not provide a way to enforce that an expression is independent from any surrounding colour. Although this does not impact the intrinsic validity of \mathcal{C} , it does limit possible applications. In the security interpretation of brackets, an additive bracket provides additional privileges to the surrounding code, which prohibits any modelling of a `,` and means that any ordinary code called by privileged code would inherit the privileges. In the context of our objective to cope with distributed system, we will need universal brackets, which ensure that their contents is usable in any context: these can be expressed naturally as absolute brackets annotated with the empty colour \bullet (see section IV.6.3.1). This second application motivates our choice of only including absolute brackets in the language (other forms of brackets being then superfluous).

¹⁵Since a is fresh, c_0 cannot contain a .

IV.5.3 Polymorphism

IV.5.3.1 Coloration of a type

We saw that how a coloured bracket can be used to build a value of an abstract type, thanks to a suitable type annotation on the bracket. For instance, if \mathbf{a} is a nonce whose implementation is $(\langle \text{INT} \rangle, 3)$ et and signature $\Sigma \mathbf{t} : \text{TYPE. Typ } \mathbf{t}$, the expression $[3]_{\{\mathbf{a}\}}^{\langle \pi_1 \mathbf{a} \rangle}$ is a value of the abstract type $\langle \pi_1 \mathbf{a} \rangle$, whereas $[3]_{\{\mathbf{a}\}}^{\text{INT}}$ evaluates to 3. Now consider a bracket around a type field: $[\langle \text{INT} \rangle]_{\{\mathbf{a}\}}^{S(\langle \pi_1 \mathbf{a} \rangle)}$ is the abstract type field written more simply as $\langle \langle \pi_1 \mathbf{a} \rangle \rangle$, while $[\langle \text{INT} \rangle]_{\{\mathbf{a}\}}^{S(\langle \text{INT} \rangle)}$ is equivalent to the simple $\langle \text{INT} \rangle$. Another expression that might be written is $[\langle \text{INT} \rangle]_{\{\mathbf{a}\}}^{\text{TYPE}}$; but what does it mean?

In an expression of the form $[\langle \mathbf{T} \rangle]_{\mathbf{c}'}^{\text{TYPE}}$, the annotation carried by the bracket is not sufficient to decide between the abstract and the concrete version of the type \mathbf{T} . The annotation TYPE does not provide any abstraction — any abstraction would result from the use of an abstract type (i.e., of a nonce) in \mathbf{T} . Rather TYPE is here an incompletely specified type, a mere indication of the signature of the module expression rather than a full specification of its type field.

The expression $[\langle \mathbf{T} \rangle]_{\mathbf{c}'}^{\text{TYPE}}$ denotes a type field, the type in question being described as the type \mathbf{T} as seen in the colour \mathbf{c}' . As we saw in section IV.5.2.2, the colour has a double influence on the type. On the one hand, it contributes in determining whether the type is well-formed (i.e., whether the judgement $\mathbf{B}; \Gamma \vdash_{\mathbf{c}} \mathbf{T} \text{ok}$ holds). On the other hand, it contributes in determining the semantics of the type, that is, which expressions have this type. For instance, if \mathbf{a} is the above nonce, the type $\langle \pi_1 \mathbf{a} \rangle$ is well-formed in any colour; in the colour \mathbf{a} , the values 3 and $[3]_{\{\mathbf{a}\}}^{\text{Typ } \pi_1 \mathbf{a}}$ both have this type, while only the latter does in the empty colour. In general, a larger colour makes more types valid¹⁶ makes more expressions have a given type.

We have seen a way to transform a type into an equivalent universal type, i.e., a type that is valid in any colour and, in the original colour, characterises the same expressions: this is the concretisation operation $\mathbf{conc}_{\mathbf{c}'}^{\mathbf{B}}(\mathbf{T})$. However $\mathbf{conc}_{\mathbf{c}'}^{\mathbf{B}}(\mathbf{T})$ does not have the same semantics in other colours, for all it is valid; hence $[\langle \mathbf{T} \rangle]_{\mathbf{c}'}^{\text{TYPE}}$ cannot be replaced by $\langle \mathbf{conc}_{\mathbf{c}'}^{\mathbf{B}}(\mathbf{T}) \rangle$.

The semantics of the expression $[\langle \mathbf{T} \rangle]_{\mathbf{c}'}^{\text{TYPE}}$ is novel: it cannot be expressed by previously seen means. It is not clear whether such an expression should be accepted at all. We shall evaluate the pros and cons of allowing such expressions. But first, we need to characterise them precisely, which we do by assigning kinds to types.

IV.5.3.2 Kinds

We would like to recognise types that fully specify values. Obvious such types are singletons: all pure expressions having a given singleton types are essentially equivalent. Following a very strict interpretation, one might say that the only fully specified types are singletons. However, under extensional equivalence, a type can be a singleton semantically without being one syntactically; for example the type $S(3) * S(4)$ does not contain any more values than $S((3, 4))$. Some types may even contain a single value up to observational equivalence for reasons having to do with the language as a whole, such as parametricity [Wad89] results in ML that ensure that the only function of type $\forall \alpha, \alpha \rightarrow \alpha$ (which we would write $\Pi \mathbf{t} : \text{TYPE. Typ } \mathbf{t} \rightarrow \text{Typ } \mathbf{t}$) is the identity function.

In fact, we are trying to characterise the signatures that fully specify the type fields that they contain. For example, although the type BOOL contains two observationally distinguishable values, we are content with BOOL as a specification of a boolean value: we do not treat BOOL as an abstract

¹⁶The colour has a bearing on embedded expressions. For instance $S((\lambda x : \text{Typ } \pi_1 \mathbf{a}. x) 3)$ is only well-formed in a colour that reveals \mathbf{a} .

type. We aim to describe a programming language, not a proof language; as a consequence we limit abstraction to types, and allow the revelation of $[\text{true}]_{\mathcal{C}}^{\text{bool}}$ as true.

The paradigm of an incomplete signature is `TYPE`, which designates an unspecified type field. More generally, any type containing `TYPE` in a covariant position, such as `INT * TYPE` or `INT → TYPE`, is incompletely specified. The presence of `TYPE` in a contravariant position does not indicate incompleteness, as shown for example in the constant function type `Πt : TYPE. S(V)`.

In order to formalise this notion, we equip system \mathcal{C} with two **kinds**. The kind λ contains fully specified types such as `S(E)`, `INT` or `Πt : TYPE. T` with `T` of kind λ . A type of kind λ is said to be **fully specified** or **completely specified**. The kind $*$ contains all types irregardless of their level of specification; a type that does not have the kind λ is said to be **partially specified** or **incompletely specified**. Kinds, written K , are equipped with an order relation $\lambda \leq *; \lambda \vee * = *$ and the greatest lower bound is written $K_1 \wedge K_2$ ($\lambda \wedge * = \lambda$). Type kinding comes with a very simple subkinding relationship: if T has the kind K_1 and $K_1 \leq K_2$ then T has the kind K_2 .

It is tempting to call a fully specified type **monomorphic** (as opposed to **polymorphic** for a partially specified type). Another tempting designation is *concrete* (vs. *abstract*). These terminologies is slightly misleading out of context (as attested by their multiplicity). In a way the type `TYPE` is a type variable (under which interpretation λ is the kind of closed types), and the meaning of its presence depends on how type variables are quantified. If they are quantified universally, types of kind $*$ are polymorphic; if they are quantified existentially, types of kind $*$ are abstract. For aesthetic reasons, we will usually use the words “monomorphic” and “polymorphic” (the latter usually meaning non-monomorphic rather than just having the kind $*$). We do however warn the reader to take this terminology with a grain of salt.

Type kinding rules are fairly simple: `TYPE` is polymorphic, any other constructor is monomorphic if and only if its components are. In particular, all base types (`BOOL`, `INT`, etc.) other than `TYPE` are monomorphic, as are singletons. A product type is monomorphic if its components are. A function type is monomorphic if its result type is.

The remaining case is that of the projection of a type field from an expression: when is `Typ E` monomorphic? Since the types `Typ T` and `T` are equivalent, `Typ T` must be monomorphic if and only if `T` is. Given the presence in our language of the dependent type `Typ E`, we need to reflect type kinding at the expression typing level. We annotate the type `TYPE` with a kind annotation, writing `TYPEλ` or `TYPE*`: the type `TYPEK` characterises type fields whose contents is a type of kind K . For example `⟨INT⟩` and `⟨TYPE* → INT⟩` have the type `TYPEλ`, and so does `⟨INT * Typ x⟩` when `x` has the type `TYPEλ`; whereas `⟨TYPEK⟩` and `⟨INT → TYPEλ⟩` only have the type `TYPE*`. One must be careful not to confuse the type assigned to the expression with the type contained in the field: for example `⟨TYPE*⟩` is an expression containing a type field, which happens to be the type of arbitrary type fields (in ML, this would be a signature field in a module, e.g., `struct module type S : sig end end` in Objective Caml). The expression `⟨TYPE*⟩` has the monomorphic type `S(⟨TYPE*⟩)`, but not the type of monomorphic fields `TYPEλ` — nor does `⟨TYPEλ⟩` since `TYPEλ` does not have the kind λ .

IV.5.3.3 Brackets and function application; polymorphic functions

Reducing a sealing introduces a coloured bracket during evaluation. The type on this bracket is produced by the selfification operation, which creates a monomorphic type. This is the central point of selfification: the type on a sealing is usually incompletely specified, and selfification completes the selfification, by replacing the unspecified parts by projections of the new name. When a coloured bracket is pushed inside a data structure, the monomorphic nature of the type annotations is preserved — all the terms being manipulated have monomorphic brackets (as opposed to polymorphic

brackets whose type annotation is partially specified). One aspect of bracket pushing in system \mathcal{C} remains to be described however, namely pushing brackets inside a function.

Consider the expression $[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. P T_1}$ in some ambient colour c . In order for it to be well-typed, T_2 must be a subtype of T_0 and E must have the type T_1 in the colour c' . When this expression is applied to an argument V of type T_0 , the result must be that of beta-reduction $\{x \leftarrow V\}E$, with any necessary coloured brackets thrown in. Let us study how to manage brackets during evaluation.

We announced in section IV.5.2.3 that beta-reduction would remain unadorned, which forces us to rule on the the fate of brackets as soon as they are pushed under the lambda. This does not constrain our latitude regarding the choice of semantics: we are effectively giving a symbolic name x to the effective argument (as well as the outer colour c , which is the colour of the argument). The question is therefore how to reduce $[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. P T_1}$.

The most obvious target uses coloured brackets both around the body of the function (to mark the border on exit from the function) and around each occurrence of the parameter (to mark the border when entering the function). The argument must also be protected in the return type.

$$[\lambda z : T_2. E]_{c'}^{\Pi y : T_0. P T_1} \longrightarrow_c \lambda x : T_0. \{z \leftarrow \{x\} [x]_{c \cup \{x\}}^{T_0}\} E_{c'}^{\{y \leftarrow \{x\} [x]_{c \cup \{x\}}^{T_0}\} T_1} \text{ (ered.col.fun.P-POLY)}$$

The bracket around the parameter x in the function body and in the return type must allow x to be used inside, hence the colour annotating the bracket must contain x (adding c as well is technically useless since x automatically brings in the colour of the binding c). The choice of what type annotation to put on this bracket is not so obvious. We know that E and T_1 are well-typed as soon as their variable has the type T_0 (note that T_2 is a subtype of T_0); T_0 is valid in c given the annotation on the bracket in the redex, therefore T_0 is a possible choice.

Nothing however requires T_0 to be monomorphic, even if $\Pi y : T_0. P T_1$ is. Therefore this rule introduces brackets carrying polymorphic type annotations. In fact, for T_0 to be partially specified means that the function $\lambda z : T_0. E$ is a **polymorphic function**¹⁷. This terminology follows that of ML: in ML, a polymorphic function has a type scheme $\forall \alpha, T_0 \rightarrow T_1$, which corresponds in our dependently typed system to $\Pi t : \text{TYPE}^l. T_0 \rightarrow T_1$ (with $\text{Typ } t$ corresponding to α).

If we want to adopt bracket pushing into functions as stated above, we must accept polymorphic brackets. Let us now study these further, in the light of how they can appear. We will then look for a way of avoiding them.

IV.5.3.4 Polymorphic types and values

We still assume that bracket pushing into functions happens according to the rule (ered.col.fun.P-POLY) from section IV.5.3.3. Polymorphic brackets result from applying a polymorphic function inside a colour other than its colour of definition: a bracket $[E]_{c'}^{\text{TYPE}^l}$ appears when applying a function $[f]_{c'}^{\Pi t : \text{TYPE}^l. T_1}$. Let us first look at a very simple example: the identity function on the type TYPE^l , published from c' under the type $\Pi t : \text{TYPE}^l. S(t)$.

$$\begin{aligned} \left([\lambda t : \text{TYPE}^l. t]_{c'}^{\Pi t : \text{TYPE}^l. S(t)} \right) \langle \text{INT} \rangle &\longrightarrow_c \left(\lambda t : \text{TYPE}^l. [[t]_{c \cup \{t\}}^{\text{TYPE}^l}]_{c'}^{S([t]_{c \cup \{t\}}^{\text{TYPE}^l})} \right) \langle \text{INT} \rangle \\ &\longrightarrow_c [[\langle \text{INT} \rangle]_{c'}^{\text{TYPE}^l}]_{c'}^{S([\langle \text{INT} \rangle]_{c'}^{\text{TYPE}^l})} \end{aligned}$$

¹⁷This is a case of the strong connection between partial specification and polymorphism when the variable is universally quantified that we mentioned in section IV.5.3.2.

In this case, the final value is $[\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}$ — the identity function returns its argument protected by a spurious bracket annotated by the ambient colour.

Let us now examine the polymorphic identity function $\lambda t : \text{TYPE}^l. \lambda x : \text{Typ } \mathbf{t}. x$ published from \mathbf{c}' under the type $\Pi t : \text{TYPE}^l. \text{Typ } \mathbf{t} \rightarrow \text{Typ } \mathbf{t}$.

$$\begin{aligned}
 & \left([\lambda t : \text{TYPE}^l. \lambda x : \text{Typ } \mathbf{t}. x]_{\mathbf{c}'}^{\Pi t : \text{TYPE}^l. \text{Typ } \mathbf{t} \rightarrow \text{Typ } \mathbf{t}} \right) \langle \text{INT} \rangle 3 \\
 & \longrightarrow_{\mathbf{c}} \left(\lambda t : \text{TYPE}^l. [\lambda x : \text{Typ } [\mathbf{t}]_{\mathbf{c} \cup \{\mathbf{t}\}}^{\text{TYPE}^l}. x]_{\mathbf{c}'}^{\text{Typ } [\mathbf{t}]_{\mathbf{c} \cup \{\mathbf{t}\}}^{\text{TYPE}^l} \rightarrow \text{Typ } [\mathbf{t}]_{\mathbf{c} \cup \{\mathbf{t}\}}^{\text{TYPE}^l}} \right) \langle \text{INT} \rangle 3 \\
 & \longrightarrow_{\mathbf{c}} [\lambda x : \text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}. x]_{\mathbf{c}'}^{\text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l} \rightarrow \text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}} 3 \\
 & \longrightarrow_{\mathbf{c}} \left(\lambda x : \text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}. [[x]_{\mathbf{c} \cup \{\mathbf{x}\}}^{\text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}}]_{\mathbf{c}'}^{\text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}} \right) 3 \\
 & \longrightarrow_{\mathbf{c}} [[3]_{\mathbf{c}}^{\text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}}]_{\mathbf{c}'}^{\text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}}
 \end{aligned}$$

The argument 3 is now surrounded by two brackets. The inner bracket places the value into the colour \mathbf{c}' with a polymorphic value; the outer bracket, in spite of its similar appearance, has a different rôle as the seemingly polymorphic type $\text{Typ } [\langle \text{INT} \rangle]_{\mathbf{c}}^{\text{TYPE}^l}$ is in fact monomorphic in the outside colour \mathbf{c} .

In order to generalise upon these examples, a few concepts are worth noting. A bracket $[\langle T \rangle]_{\mathbf{c}}^{\text{TYPE}^l}$ is a **polymorphic type parameter** for a function. A bracket $[V]_{\mathbf{c}}^{\text{Typ } [\langle T \rangle]_{\mathbf{c}}^{\text{TYPE}^l}}$ is a **polymorphic value**. A polymorphic value has no apparent structure, since it is protected by a bracket (which cannot be reduced away since its type annotation itself has no apparent structure). The body of the function is unable to manipulate polymorphic values in a way other than polymorphic. This approach should work to model parametrically polymorphic languages such as ML. However it is problematic in a non-parametric language with generics or dynamic type-checking (which we will introduce in system \mathcal{D}).

We shall not go any further along the lines of studying polymorphic values. It remains to be seen how to reduce polymorphic brackets. In particular, how can the polymorphic identity function return its argument with no superfluous bracket? (Note that the argument passes through the colour \mathbf{c}' ; how can we make sure that this passage is harmless?)

IV.5.3.5 Colour fusion

We present a solution to the problem of managing brackets around polymorphic function calls. This solution lacks expressivity and finesse, but remains attractive in a certain light — not least because of its simplicity. The idea is to merge the colour of the argument with the colour of the function body.

$$[\lambda x : T_2. E]_{\mathbf{c}'}^{\Pi x : T_0. P T_1} \longrightarrow_{\mathbf{c}} \lambda x : T_0. [E]_{\mathbf{c}' \cup \{\mathbf{x}\}}^{T_1} \quad (\mathcal{C}/\text{ered.col.fun.P})$$

We abandon any thought of protecting the argument: all type equations required to type the argument are allowed when executing the function. This rule is very simple, technically speaking: one bracket turns into one bracket, with a smaller type annotation and smaller contents.

This rule enjoys a certain symmetry: applying $[\lambda x : T_2. E]_{\mathbf{c}'}^{\Pi x : T_0. P T_1}$ to an argument V yields

$$[[x \leftarrow_{\mathbf{c}} V] E]_{\mathbf{c} \cup \mathbf{c}'}^{\{x \leftarrow_{\mathbf{c}} V\} T_1}$$

so that the computations are simply performed in the union of the colours of the expressions that come into contact (V and E)

Let us check the result of an application of the polymorphic identity function using this rule.

$$\begin{aligned}
 & \left([\lambda t : \text{TYPE}^l. \lambda x : \text{Typ } t. x]_{c'}^{\Pi t : \text{TYPE}^l. \text{Typ } t \rightarrow \text{Typ } t} \right) \langle \text{INT} \rangle 3 \\
 & \longrightarrow_c \left(\lambda t : \text{TYPE}^l. [\lambda x : \text{Typ } t. x]_{c' \cup \{t\}}^{\text{Typ } t \rightarrow \text{Typ } t} \right) \langle \text{INT} \rangle 3 \\
 & \longrightarrow_c [\lambda x : \text{Typ } \langle \text{INT} \rangle. x]_{c \cup c'}^{\text{Typ } \langle \text{INT} \rangle \rightarrow \text{Typ } \langle \text{INT} \rangle} 3 \\
 & \longrightarrow_c \left(\lambda x : \text{Typ } \langle \text{INT} \rangle. [\lambda x : \text{Typ } \langle \text{INT} \rangle. x]_{c \cup c' \cup \{x\}}^{\text{Typ } \langle \text{INT} \rangle} \right) 3 \\
 & \longrightarrow_c [3]_{c \cup c'}^{\text{Typ } \langle \text{INT} \rangle} \longrightarrow_c [3]_{c \cup c'}^{\text{INT}} \longrightarrow_c 3
 \end{aligned}$$

In our study of system C, we will retain this fusion formulation of bracket pushing around a function.

IV.5.3.6 Generative functors

We saw in section IV.5.1.5 that selfification does not affect generative functors. Since selfification produces a type that is meant to annotate a coloured bracket (as it is used in (C/ered.seal)), the resulting type must be monomorphic. Therefore a generative functor type $\Pi x : T_0. {}^1T_1$ must be monomorphic even if T_1 is polymorphic¹⁸.

We have stated a rule (C/ered.col.fun.P) to push coloured brackets bearing an applicative functor type. The transposition to a generative functor type is not straightforward. A naive proposal would be

$$\overline{[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. {}^1T_1} \longrightarrow_c \lambda x : T_0. [E]_{c' \cup \{x\}}^{T_1}}$$

However T_1 may be polymorphic, in which case the right-hand side is ill-typed. Intuitively, the rule above cannot be suitable because the left-hand side is a generative functor, whose every application triggers the creation of a new nonce, whereas this aspect is simply not present in the right-hand side.

Nonces are generated by the reduction rule (C/ered.seal) for sealing expressions. Let us therefore introduce a sealing in the right-hand side. We may attempt to place the sealing inside or outside the coloured brackets:

$$\begin{aligned}
 & \overline{[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. {}^1T_1} \longrightarrow_c \lambda x : T_0. ([E]_{c' \cup \{x\}}^{T_1} !! T_1)} \\
 & \overline{[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. {}^1T_1} \longrightarrow_c \lambda x : T_0. [E !! T_1]_{c' \cup \{x\}}^{T_1}}
 \end{aligned}$$

In both cases, although the intuitive behaviour is acceptable, formal correction is lacking, as the type annotation on the coloured bracket may still be polymorphic. However this is a benign form of ill typing, as reduction will change the type into a monomorphic one before the coloured bracket is reduced.

A sealing construct expresses a static border between abstraction domains, whereas a coloured bracket is a dynamic border. We have here a border that is both static and dynamic. We will note it by a **coloured sealing**, written $E !!_{c'} T$. This sealing acts like normal sealing, except that it gives the expression E the additional knowledge of abstract types designated by c' . A coloured sealing

¹⁸In a way a generative functor type is amorphous: it is not yet fully specified, but will give rise to a monomorphic type when the functor is applied.

therefore includes the effect of an additive bracket (see section IV.5.2.4). A normal sealing is the special case where the coloured sealing adds no extra knowledge: $E !! T = E !!_{\bullet} T$.

Pushing a coloured bracket bearing a generative functor type shall produce a coloured sealing:

$$[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. P T_1} \longrightarrow_c \lambda x : T_0. (E !!_{c'} T_1) \quad (\mathcal{C}/\text{ered.col.fun.l})$$

Note that thanks to the additivity of c' we do not need to include x , in contrast with $(\mathcal{C}/\text{ered.col.fun.P})$. A coloured sealing is reduced by the rule $(\mathcal{C}/\text{ered.seal})$ which we can finally state in its full glory:

$$\frac{}{B \vdash V !!_{c'} T \longrightarrow_c B, a = V :_{c \cup c'} T \vdash [V]_{c \cup c' \cup \{a\}}^{\text{self } T(a)} \quad (\mathcal{C}/\text{ered.seal})}$$

IV.5.4 Evaluation

IV.5.4.1 Syntax

The syntax of system \mathcal{C} extends system \mathcal{E} with two new constructs that should not appear in source programs: abstract types and coloured brackets. Furthermore the signature `TYPE` now carries a kind annotation, and sealing now carries a colour annotation (the notation $E !! T$ is kept as an abbreviation for $E !!_{\bullet} T$).

$K ::=$	kind
\wr	monomorphic (fully specified)
$*$	polymorphic (partially specified)
$T ::=$	type
\dots	
TYPE^K	abstract type field
(A)	abstract type
$E ::=$	expression
\dots	
$E !!_c T$	sealed and coloured module
$[E]_c^T$	coloured bracket
$A ::=$	module component
a	nonce
ΛE	application
$\pi_i A$	projection ($i \in \{1, 2\}$)
$\xi ::=$	primary colour
a	nonce
x	variable
$c ::=$	colour
\bullet	empty colour (also written $\{\}$)
$\{a_1, \dots, a_k, x_1, \dots, x_k\}$	finite set of primary colours

Recall that kinds are equipped with an order relation, written $K_1 \leq K_2$, such that $\wr \leq *$. We write $K_1 \vee K_2$ for the least upper bound of K_1 and K_2 , and $K_1 \wedge K_2$ for their greatest lower bound.

If A is a module component, its underlying nonce **under1**(A) is formally defined as follows:

$\mathbf{underl}(a) = a$

$\mathbf{underl}(A E) = \mathbf{underl}(A)$

$\mathbf{underl}(\pi_i A) = \mathbf{underl}(A)$

Revelation of a module component is defined as follows:

$\mathbf{reveal}^B(a) = E$ where $a = E :_c T \in B$

$\mathbf{reveal}^B(A E) = (\mathbf{reveal}^B(A)) E$

$\mathbf{reveal}^B(\pi_i A) = \pi_i(\mathbf{reveal}^B(A))$

Typing judgements now carry a lexis and a colour. Additional right-hand sides to those in system \mathcal{E} are colour transparency, revelation of a module component and conversion and convertibility for components.

$\mathcal{J} ::=$ **typing judgement**
 $B; \Gamma \vdash_c J$ local judgement

$J ::=$ **local judgement right-hand side**

...
 $T : K$ type kinding (generalising Tok)
 c_0 transparent colour transparency
 $A \triangleright E : T$ component revelation
 $A \longrightarrow A'$ component conversion
 $A \equiv A'$ convertibility equivalence on components

We write ξ transparent for $\{\xi\}$ transparent.

Environments now contain colour annotations. We also state the syntax of lexes.

$B ::=$ **lexis**
 nil empty
 $B, a = E :_{c_0} T$ nonce a with implemented by E with the signature T

$\Gamma ::=$ **environnement**
 nil empty
 $\Gamma, x :_c T$ binding of the variable x

Following the definition for environments, the domain of a lexis B , i.e., the set of nonces that it records, is written $\mathbf{dom} B$.

Since colours may contain variables, they are affected by substitutions. A substitution specifies both an expression and a colour to replace the variable with. The substitution of E for x under c in \mathfrak{N} is written $\{x \leftarrow_c E\} \mathfrak{N}$.

IV.5.4.2 Values and abstract components

Brackets and values As in HAT, the set of values depends on the ambient colour. We first define a grammatical notion of *quasi-value*, which is a value with some possibly-eliminatable brackets. In addition to the values of system \mathcal{E} (which are the same as in system \mathcal{B}), coloured brackets may appear in quasi-values and (with restrictions) in values. In order for an expression of the form $[V]_{c'}^T$ to be a value, V must be a value (in the colour c'), and T must have an appropriate form. If T has apparent structure, the bracket pushing rules allow $[V]_{c'}^T$ to be reduced. The only case where $[V]_{c'}^T$ may be a value is when T is an abstract type $\langle A \rangle$. Even then, $[V]_{c'}^{\langle A \rangle}$ may be reducible in some colours.

Quasi-values The language of **quasi-values** in system \mathcal{C} is a supergrammar of the one for \mathcal{B} , with brackets carrying an abstract type annotation thrown in. The abstract type annotation must itself be in a reduced form, called **component value**, where functor arguments are all values.

$V ::=$ **quasi-value**
 \dots
 $[V]_{c'}^{(A^V)}$ potentially abstraction-making coloured bracket

$A^V ::=$ **component value**
 a nonce
 $A^V V$ application to a quasi-value
 $\pi_i A^V$ projection ($i \in \{1, 2\}$)

Irreducible coloured brackets A quasi-value of the form $[V]_{c'}^{(A^V)}$ is only a value if the bracket cannot be eliminated. Intuitively a coloured bracket is indispensable only if it actually creates abstraction, which translates as the requirement that the underlying nonce of A^V must be opaque in the ambient colour yet transparent in the inside colour c' . We will analyse the behaviour of a coloured bracket expression according to the transparency of the underlying nonce when presenting bracket elimination rules in section IV.5.4.2.

Values and abstract components The set of values depends on the ambient colour: we write V^c for a value in the colour c . The set of values is described as a family of grammars parametrised by a colour; it is a subset of quasi-values. In order for a quasi-value $[V]_{c'}^{(A^V)}$ to be a value, the bracket must be indispensable in the sense described above, and the quasi-values in A^V must themselves be values in the appropriate colour.

$V^c ::=$ **value in c**
 $() \mid bv \mid \underline{n}$ constant
 $\langle T \rangle$ type field
 (V_1^c, V_2^c) pair
 $\lambda x : T. E$ lambda-abstraction
 $[V^{c'}]_{c'}^{(A^{V^{c \cap c'}})}$ coloured bracket, if $A^{V^{c \cap c'}}$ is abstract in c but concrete in c'

$A^{V^c} ::=$ **abstract component in c**
 a nonce, if opaque in c
 $A^{V^c} V^c$ application of a functor to a value
 $\pi_i A^{V^c}$ projection ($i \in \{1, 2\}$)

Strictly speaking, since transparency of a nonce is a semantic value depending on a lexis and an environment, the notions of values and abstract components should be indexed by a lexis and an environment. In practice the lexis and environments will always be clear from context, so we omit them.

IV.5.4.3 $B \vdash E \longrightarrow_c B' \vdash E'$ **Reduction**

Evaluatin contexts We reduce expressions under brackets. Although brackets are initially introduced around values, this property is not preserved by reduction; specifically, pushing a bracket inside a function body results in a bracket surrounding an arbitrary expression. When the type annotation on a bracket is the type field of a module, the module expression must also be reduced.

$\mathcal{C} ::=$ **evaluation context (of depth 1)**

\dots
 $\underline{\quad} !!_{c_1} \top$ sealing
 $\underline{\quad}^T_{c_1}$ coloured bracket
 $[\mathbf{V}^{c_1}]_{c_1}^{\text{Typ}} \underline{\quad}$ type field on a bracket

Formally speaking, the set of evaluation contexts, like values, depends on a lexis and an environment. In the context $\underline{\quad}^T_{c_1}$, the expression inside is reduced in the colour c_1 . In the context $[\mathbf{V}^{c_1}]_{c_1}^{\text{Typ}} \underline{\quad}$, the expression inside is reduced in the colour $c \cap c_1$ (the intersection of the colours outside and inside the border upon which the expression lies).

Computational rules System \mathcal{C} inherits the rules that were already present in \mathcal{B} , viz., ($\mathcal{C}/\text{ered.app}$), ($\mathcal{C}/\text{ered.proj}$), ($\mathcal{C}/\text{ered.let}$), ($\mathcal{C}/\text{ered.context}$). These rules can be used in any colour and any lexis; the colour is added to substitution when required. Values are also considered in their ambient colour. Reduction inside contexts happens in the inside colour of the context. Appending A lists all the rules of system \mathcal{C} , included the inherited rules.

The rule for reducing a sealing is modified to surround the value with a coloured bracket, and to take the colour annotation on the sealing into account.

$$B \vdash \mathbf{V}^{c \cup c'} !!_{c'} \top \longrightarrow_c B, a = \mathbf{V}^{c \cup c'} :_{c \cup c'} \top \vdash [\mathbf{V}^{c \cup c'}]_{c \cup c' \cup \{a\}}^{\text{self}^T(a)} (\mathcal{C}/\text{ered.seal})$$

where a is fresh (i.e., $a \notin \text{dom } B$)

Reductions in types Until system \mathcal{E} , types contained in expressions did not influence reduction. This is no longer the case in system \mathcal{C} , since the reduction of a coloured bracket depends on the type annotation carried by the bracket, specifically on the head constructor on this type. Nonetheless our computational needs on types are small enough — we only need to reach a weak head normal form, and only in a single context within expressions, so we do not need to introduce a reduction on types. The only destructor in the syntax of types is $\text{Typ } \underline{\quad}$; its argument can be reduced via the context $[\mathbf{V}^{c_1}]_{c_1}^{\text{Typ}} \underline{\quad}$, after which the destructor can be eliminated with ($\mathcal{C}/\text{ered.colTyp}$).

$$[\mathbf{V}^{c'}]_{c'}^{\text{Typ}(T)} \longrightarrow_c [\mathbf{V}^{c'}]_{c'}^T \quad (\mathcal{C}/\text{ered.colTyp})$$

Abstract types are peculiar, as $\langle A \rangle$ is in weak head normal form if and only if the underlying nonce is opaque, but must be revealed if it is transparent.

$$B \vdash [\mathbf{V}^{c'}]_{c'}^{\langle A \rangle} \longrightarrow_c B \vdash [\mathbf{V}^{c'}]_{c'}^{\text{Typ reveal}^B(A)} \quad (\mathcal{C}/\text{ered.colAbs})$$

if $\text{underl}(A) \in c \cap c'$

Bracket pushing When a bracket surrounds a value, and the type annotation on the bracket is not an abstract type, the bracket is pushed inside a value. The bracket pushing rules mostly follow the same principle as in HAT (see sections III.1.2.2 and III.2.5). The selection of the bracket pushing rule relies on the type that is apparent on the bracket (specifically its head constructor); the effect

on the expression is to push the bracket inside the constructor for this type.

$$\begin{aligned}
 [()]_{c'}^{\text{UNIT}} &\longrightarrow_c () && (\mathcal{C}/\text{ered.col.base.unit}) \\
 [bv]_{c'}^{\text{BOOL}} &\longrightarrow_c bv && (\mathcal{C}/\text{ered.col.base.bool}) \\
 [\underline{n}]_{c'}^{\text{INT}} &\longrightarrow_c \underline{n} && (\mathcal{C}/\text{ered.col.base.int}) \\
 [V^{c'}]_{c'}^{\text{S(E)}} &\longrightarrow_c E && (\mathcal{C}/\text{ered.col.sing}) \\
 [(V_1^{c'}, V_2^{c'})]_{c'}^{\Sigma x:T_1.T_2} &\longrightarrow_c ([V_1^{c'}]_{c'}^{T_1}, [V_2^{c'}]_{c'}^{\{x \leftarrow_c [V_1^{c'}]_{c'}^{T_1}\} T_2}) && (\mathcal{C}/\text{ered.col.pair})
 \end{aligned}$$

In the case of functions, we adopt the colour fusion rule explained in section IV.5.3.5. In the case of a generative functor, new types must be created whenever the functor is applied, so we add a sealing to the body of the functor; the colour annotation on the sealing plays the role of a coloured bracket.

$$\begin{aligned}
 [\lambda x : T_2. E]_{c'}^{\Pi x:T_0.P T_1} &\longrightarrow_c \lambda x : T_0. [E]_{c' \cup \{x\}}^{T_1} && (\mathcal{C}/\text{ered.col.fun.P}) \\
 [\lambda x : T_2. E]_{c'}^{\Pi x:T_0.I T_1} &\longrightarrow_c \lambda x : T_0. (E !!_{c' \cup \{x\}} T_1) && (\mathcal{C}/\text{ered.col.fun.I})
 \end{aligned}$$

When a bracket immediately surrounds another bracket and neither bracket can be reduced by one of the already mentioned pushing rules, i.e., given an expression of the form $[[V^{c_2}]_{c_2}^{(A_2)}]_{c_1}^{(A_1)}$ where $[V^{c_2}]_{c_2}^{(A_2)}$ is a value, there are three possible behaviours.

- If the annotation on the outer bracket makes it simplifiable, the outer bracket disappears. In HAT, this is performed by $(\mathcal{H}/\text{ered.col.le})$. Here the rule $(\mathcal{C}/\text{ered.col.Abs})$ is used, followed by computations on the revealed expression in the type annotation and possibly later bracket pushing.
- If the annotation on the outer bracket is abstract outside but concrete inside, the expression is a value.
- The remaining case is when the annotation on the outer bracket is abstract outside as well as inside. In HAT, typing ensures that A_1 and A_2 are equal, and the outer bracket is erased by the rule $(\mathcal{H}/\text{ered.col.col})$.

In our present systems, which includes functors and colours with non-trivial intersections, the situation is more complex. A new possibility arises that $A_1 = a_1 V_1$ and $A_2 = a_2 V_2$; then typing ensures that (as in HAT) $a_1 = a_2$, but the arguments are only known to be equivalent in the intermediate colour c_1 . The arguments may not be equivalent in c , so c_1 is (sometimes) an obligatory intermediate. We state a weaker rule, which (as in function application) merges the colours in play.

$$\begin{aligned}
 [[V^{c_2}]_{c_2}^{(A_2)}]_{c_1}^{(A_1)} &\longrightarrow_c [V^{c_2}]_{c_1 \cup c_2}^{(A_1)} && (\mathcal{C}/\text{ered.col.merge}) \\
 &\text{if } A_1 \text{ et } A_2 \text{ are both opaque in } c_1 \text{ but } A_2 \text{ is concrete in } c_2
 \end{aligned}$$

IV.5.5 Typing

The type system of system \mathcal{C} inherits from that of \mathcal{E} , but all rules must be modified to add lexes and colours. For most rules, this modification is done mechanically by permitting an arbitrary lexis and colour. Each judgement $\Gamma \vdash J$ becomes $B; \Gamma \vdash_c J$. When a variable is bound by the environment, it must be added to the colour: $\Gamma, x : T \vdash J$ becomes $B; \Gamma, x :_c T \vdash_{c \cup \{x\}} J$. Substitutions must also be decorated with the appropriate colours. Typical examples are given by the rules $(\mathcal{C}/\text{et.fun})$ et $(\mathcal{C}/\text{et.app})$ given below:

$$\frac{B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} E :^{\gamma} T_1}{B; \Gamma \vdash_c \lambda x : T_0. E :^P \Pi x : T_0. {}^{\gamma} T_1} \text{ (C/et.fun)} \quad \frac{B; \Gamma \vdash_c E_1 :^{\gamma_1} \Pi x : T_0. {}^{\gamma_2} T \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c E_1 E_0 :^{\gamma_1 \sqcup \gamma_2} \{x \leftarrow_c E_0\} T} \text{ (C/et.app)}$$

In addition to the addition of colours, type correction judgements Tok now become type kinding judgements $T : *$. Contexts that required an expression of type TYPE now require TYPE^* . The complete list of adapted rules is the following:

- all conversion, convertibility and subtyping rules: (C/econv.cong.fun.arg), (C/econv.cong.fun.body), (C/econv.cong.app.fun), (C/econv.cong.app.arg), (C/econv.cong.pair.1), (C/econv.cong.pair.2), (C/econv.cong.field), (C/econv.cong.proj), (C/econv.app), (C/econv.proj), (C/econv.eta.field), (C/econv.eta.fun), (C/econv.eta.pair), (C/tconv.cong.pair.1), (C/tconv.cong.pair.2), (C/tconv.cong.fun.arg), (C/tconv.cong.fun.ret), (C/tconv.cong.sing), (C/tconv.cong.field), (C/tconv.field), (C/tconv.unit), (C/eeq.refl), (C/eeq.sym), (C/eeq.trans), (C/eeq.conv), (C/teq.refl), (C/teq.sym), (C/teq.trans), (C/teq.conv), (C/tsub.trans), (C/tsub.eq), (C/tsub.cong.fun), (C/tsub.cong.pair), (C/tsub.sing);
- most expression typing rules: (C/et.base.unit), (C/et.base.bool), (C/et.base.int), (C/et.fun), (C/et.app), (C/et.pair), (C/et.proj.1), (C/et.proj.2), (C/et.let), (C/et.sub), (C/et.sing).

Appendix A contains a complete list of the rules of system \mathcal{C} in their final form, including inherited rules.

IV.5.5.1 B; $\Gamma \vdash_c \text{ok}$ Environment formation

We describe how to build lexes, environments and colours. These components are built from left to right, both inside lexes and environments (which are built binding by binding from left to right) and in that the lexis is built first, then the environment, then the colour.

Lexis and environment validity follow a similar principle: all entered information must be checked for validity, and a fresh name must be used to label each binding. Note that the colour of a lexis binding may include previous nonces, while that of an environment binding may use any nonce in the lexis as well as previous variables.

$$\frac{}{\text{nil}; \text{nil} \vdash_{\bullet} \text{ok}} \text{ (C/envok.nil)} \quad \frac{B; \text{nil} \vdash_{c_0} E :^P T \quad \text{when } a \notin \mathbf{dom} B}{B, a = E :_{c_0} T; \text{nil} \vdash_{\bullet} \text{ok}} \text{ (C/envok.a)}$$

$$\frac{B; \Gamma \vdash_c T : * \quad \text{when } x \notin \mathbf{dom} \Gamma}{B; \Gamma, x :_c T \vdash_{\bullet} \text{ok}} \text{ (C/envok.x)}$$

A colour may contain nonces and variables taken respectively from the lexis and the environment. As a colour is an unordered set, there is no constraint on the order in which a colour is built (other than the validity of intermediate colours when adding a nonce). Nonces may only be added to the colour if their dependencies are already present (or otherwise transparent), while variables may be added at any time — see section IV.5.2.3 and the transparency rules below).

$$\frac{B; \Gamma \vdash_{c'} \text{ok} \quad \text{when } a = E :_{c_0} T \in B \wedge c_0 \subseteq c'}{B; \Gamma \vdash_{c' \cup \{a\}} \text{ok}} \text{ (C/envok.c.a)}$$

$$\frac{B; \Gamma \vdash_{c'} \text{ok} \quad \text{when } x :_{c_0} T \in \Gamma}{B; \Gamma \vdash_{c' \cup \{x\}} \text{ok}} \text{ (C/envok.c.x)}$$

IV.5.5.2 B; $\Gamma \vdash_c T : K$ Type kinding

Type kinding rules refine the type correction rules of system \mathcal{E} . Adding colours is straightforward.

As discussed in section IV.5.3.6, a generative functor is always monomorphic (whereas an applicative functor has the same kind as its result type); the rule ($\mathcal{E}/\text{tok.fun}$) is split to treat each case correctly.

$$\begin{array}{c}
 \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{BOOL} : \lambda} \text{ (}\mathcal{E}/\text{tok.base.bool)} \quad \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{INT} : \lambda} \text{ (}\mathcal{E}/\text{tok.base.int)} \quad \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{UNIT} : \lambda} \text{ (}\mathcal{E}/\text{tok.base.unit)} \\
 \\
 \frac{B; \Gamma \vdash_c E : {}^P \text{TYPE}^K}{B; \Gamma \vdash_c \text{Typ } E : K} \text{ (}\mathcal{E}/\text{tok.field)} \\
 \\
 \frac{B; \Gamma \vdash_c T' : K' \quad B; \Gamma, x :_c T' \vdash_{c \cup \{x\}} T'' : K''}{B; \Gamma \vdash_c \Pi x : T'. {}^P T'' : K''} \text{ (}\mathcal{E}/\text{tok.fun.P)} \\
 \\
 \frac{B; \Gamma \vdash_c T' : K' \quad B; \Gamma, x :_c T' \vdash_{c \cup \{x\}} T'' : K''}{B; \Gamma \vdash_c \Pi x : T'. {}^I T'' : \lambda} \text{ (}\mathcal{E}/\text{tok.fun.I)} \\
 \\
 \frac{B; \Gamma \vdash_c T' : K' \quad B; \Gamma, x :_c T' \vdash_{c \cup \{x\}} T'' : K''}{B; \Gamma \vdash_c \Sigma x : T'. T'' : K' \vee K''} \text{ (}\mathcal{E}/\text{tok.pair)} \\
 \\
 \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{TYPE}^K : *} \text{ (}\mathcal{E}/\text{tok.type)} \quad \frac{B; \Gamma \vdash_c E : {}^P T}{B; \Gamma \vdash_c S(E) : \lambda} \text{ (}\mathcal{E}/\text{tok.sing)}
 \end{array}$$

An extra rule indicates that any monomorphic type is also polymorphic. Similarly a type field containing a monomorphic type can be seen as a type field containing a polymorphic type, so the type of the former is a subtype of the type of the latter.

$$\frac{B; \Gamma \vdash_c T : K' \quad \text{when } K' \leq K}{B; \Gamma \vdash_c T : K} \text{ (}\mathcal{E}/\text{tok.sub)} \quad \frac{B; \Gamma \vdash_c \text{ok} \quad \text{when } K_1 \leq K_2}{B; \Gamma \vdash_c \text{TYPE}^{K_1} <: \text{TYPE}^{K_2}} \text{ (}\mathcal{E}/\text{tsub.cong.type)}$$

Let us also state the rule for forming a type field, which is also modified to account for kinding.

$$\frac{B; \Gamma \vdash_c T : K}{B; \Gamma \vdash_c \langle T \rangle : {}^P \text{TYPE}^K} \text{ (}\mathcal{E}/\text{et.type)}$$

IV.5.5.3 B; $\Gamma \vdash_c c_0$ transparent **Colour transparency**

A primary colour (nonce or variable) can be transparent if it is directly present in the ambient colour. It can also be transparent if it is indirectly made so, via a variable that is present in the ambient colour and whose colour of definition makes the primary colour under consideration transparent.

$$\frac{B; \Gamma \vdash_c \text{ok} \quad \text{when } \xi \in c}{B; \Gamma \vdash_c \xi \text{ transparent}} \text{ (}\mathcal{E}/\text{vis.in)} \\
 \frac{B; \Gamma \vdash_c \text{ok} \quad B; \Gamma_0 \vdash_{c_0} \xi \text{ transparent} \quad \text{when } \Gamma = (\Gamma_0, x :_{c_0} T, \Gamma_1) \wedge x \in c}{B; \Gamma \vdash_c \xi \text{ transparent}} \text{ (}\mathcal{E}/\text{vis.env)}$$

A colour is transparent if and only if all of its elements are transparent.

$$\frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \bullet \text{ transparent}} \text{ (}\mathcal{E}/\text{vis.o)} \quad \frac{B; \Gamma \vdash_c c_1 \text{ transparent} \quad B; \Gamma \vdash_c c_2 \text{ transparent}}{B; \Gamma \vdash_c c_1 \cup c_2 \text{ transparent}} \text{ (}\mathcal{E}/\text{vis.union)}$$

Nonce transparency is used in ($\mathcal{E}/\text{tconv.abs}$) to justify revealing it. Colour transparency is used in several rules ($(\mathcal{E}/\text{ac.a})$, ($\mathcal{E}/\text{et.x}$)) to express the transparency of the dependencies of a primary colour.

IV.5.5.4 B; $\Gamma \vdash_c A \triangleright E : T ; \dots$ **Module components**

Revelation judgements $B; \Gamma \vdash_c A \triangleright E : T$ assign two pieces of information to a component A : the

expression E to which it is revealed, and the apparent signature Γ mechanically derived from the signature of the underlying nonce in the lexis. The structure of the revelation derivation follows that of this signature.

$$\begin{array}{c}
 \frac{B; \Gamma \vdash_c c_0 \text{ transparent} \quad \text{when } a = E :_{c_0} \Gamma \in B}{B; \Gamma \vdash_c a \triangleright E : \Gamma} \text{ (c/ac.a)} \quad \frac{B; \Gamma \vdash_c A \triangleright E : \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_1 A \triangleright \pi_1 E : T_1} \text{ (c/ac.proj.1)} \\
 \\
 \frac{B; \Gamma \vdash_c E_1 :^P S(\pi_1 E) \quad B; \Gamma \vdash_c A \triangleright E : \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_2 A \triangleright \pi_2 E : \{x \leftarrow_c E_1\} T_2} \text{ (c/ac.proj.2)} \\
 \\
 \frac{B; \Gamma \vdash_c A \triangleright E : \Pi x : T_0. {}^P T_1 \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c A E_0 \triangleright E E_0 : \{x \leftarrow_c E_0\} T_1} \text{ (c/ac.app)}
 \end{array}$$

When a component has the apparent signature TYPE^K , it can be used to form an abstract type. If the underlying nonce is transparent, this abstract type can be converted to the revealed representation.

$$\begin{array}{c}
 \frac{B; \Gamma \vdash_c A \triangleright E : \text{TYPE}^K}{B; \Gamma \vdash_c \langle A \rangle : K} \text{ (c/tok.abs)} \\
 \\
 \frac{B; \Gamma \vdash_c A \triangleright E : \text{TYPE}^K \quad B; \Gamma \vdash_c \text{underl}(A) \text{ transparent}}{B; \Gamma \vdash_c \langle A \rangle \longrightarrow \text{Typ } E} \text{ (c/tconv.abs)}
 \end{array}$$

A component is almost inert: the only conversion that might significantly affect it is its revelation. Context rules are however needed to enable conversion of embedded expressions.

$$\begin{array}{c}
 \frac{B; \Gamma \vdash_c E_0 \longrightarrow E'_0 \quad B; \Gamma \vdash_c E_0 :^P T_0 \quad B; \Gamma \vdash_c A \triangleright E : \Pi x : T_0. {}^P T_1}{B; \Gamma \vdash_c A E_0 \longrightarrow A E'_0} \text{ (c/aconv.cong.app.arg)} \\
 \\
 \frac{B; \Gamma \vdash_c A \longrightarrow A' \quad B; \Gamma \vdash_c A \triangleright E : \Pi x : T_0. {}^P T_1 \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c A E_0 \longrightarrow A' E_0} \text{ (c/aconv.cong.app.fun)} \\
 \\
 \frac{B; \Gamma \vdash_c A \longrightarrow A' \quad B; \Gamma \vdash_c A \triangleright E : \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_i A \longrightarrow \pi_i A} \text{ (c/aconv.cong.proj)} \\
 \\
 \frac{B; \Gamma \vdash_c A \longrightarrow A' \quad B; \Gamma \vdash_c A \triangleright E : \text{TYPE}^K}{B; \Gamma \vdash_c \langle A \rangle \longrightarrow \langle A' \rangle} \text{ (c/tconv.cong.abs)}
 \end{array}$$

Convertibility equivalence for components follows the same model as for types and expressions, with four rules (c/aeq.refl), (c/aeq.sym), (c/aeq.trans) and (c/aeq.conv) following the model of (teq.*) and (eeq.*).

IV.5.5.5 $B; \Gamma \vdash_c E :^Y \Gamma$ Coloration of expressions

The rule for typing variables contains a novel side condition which requires the transparency of the variable in the ambient colour. This ensures that the dependencies of the variable keep being present if the ambient colour is weakened.

$$\frac{B; \Gamma \vdash_c x \text{ transparent} \quad \text{when } x :_{c'} \Gamma \in \Gamma}{B; \Gamma \vdash_c x :^P \Gamma} \text{ (c/et.x)}$$

Coloured brackets surround an expression with a different colour from its surroundings. The type annotation must be valid both in both the outer and inner colours, for which we use the intersection of the two colours.

$$\frac{B; \Gamma \vdash_{c'} E :^Y \Gamma \quad B; \Gamma \vdash_{c \cap c'} T : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E]_{c'}^T :^Y \Gamma} \text{ (c/et.col)}$$

The sealing typing rule takes the colour annotation into account. The colour is added when typing the body of the module.

$$\frac{B; \Gamma \vdash_{c \cup c'} E :^{\gamma} T \quad B; \Gamma \vdash_c T : *}{B; \Gamma \vdash_c (E !!_{c'} T) :^{\bar{1}} T} \text{ (C/et.seal)}$$

IV.5.5.6 $B; \Gamma \vdash_c E \longrightarrow E'$ Conversion and coloured brackets

New conversion rules reflect the new reduction rules concerning brackets: new reduction contexts, and bracket pushing rules.

$$\frac{B; \Gamma \vdash_{c'} E \longrightarrow E' \quad B; \Gamma \vdash_{c'} E :^P T \quad B; \Gamma \vdash_{c \cap c'} T : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E]_{c'}^{\bar{1}} \longrightarrow [E']_{c'}^{\bar{1}}} \text{ (C/econv.cong.col.e)}$$

$$\frac{B; \Gamma \vdash_{c'} E :^P T_1 \quad B; \Gamma \vdash_{c \cap c'} T_1 \longrightarrow T_2 \quad B; \Gamma \vdash_{c \cap c'} T_1 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E]_{c'}^{\bar{1}} \longrightarrow [E']_{c'}^{\bar{2}}} \text{ (C/econv.cong.col.t)}$$

$$\frac{B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [()]_{c'}^{\text{UNIT}} \longrightarrow ()} \text{ (C/econv.col.base.unit)} \quad \frac{B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\text{bv}]_{c'}^{\text{BOOL}} \longrightarrow \text{bv}} \text{ (C/econv.col.base.bool)}$$

$$\frac{B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\underline{n}]_{c'}^{\text{INT}} \longrightarrow \underline{n}} \text{ (C/econv.col.base.int)}$$

$$\frac{B; \Gamma \vdash_{c'} T_0 <: T_2 \quad B; \Gamma, x :_{c'} T_2 \vdash_{c' \cup \{x\}} E :^P T_1 \quad B; \Gamma \vdash_c \text{ok} \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 : \lambda}{B; \Gamma \vdash_c [\lambda x : T_2. E]_{c'}^{\Pi x : T_0. P T_1} \longrightarrow \lambda x : T_0. [E]_{c' \cup \{x\}}^{\bar{1}}} \text{ (C/econv.col.fun.P)}$$

$$\frac{B; \Gamma \vdash_{c'} T_0 <: T_2 \quad B; \Gamma, x :_{c'} T_2 \vdash_{c' \cup \{x\}} E :^{\bar{1}} T_1 \quad B; \Gamma \vdash_c \text{ok} \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 : \lambda}{B; \Gamma \vdash_c [\lambda x : T_2. E]_{c'}^{\Pi x : T_0. \bar{1} T_1} \longrightarrow \lambda x : T_0. E !!_{c' \cup \{x\}} T_1} \text{ (C/econv.col.fun.I)}$$

$$\frac{B; \Gamma \vdash_{c'} E_1 :^P T_1 \quad B; \Gamma \vdash_{c \cap c'} T_1 : \lambda \quad B; \Gamma, x :_{c'} T_1 \vdash_{c' \cup \{x\}} E_2 :^P T_2 \quad B; \Gamma \vdash_c \text{ok} \quad B; \Gamma, x :_{c \cap c'} T_1 \vdash_{(c \cap c') \cup \{x\}} T_2 : \lambda \quad B; \Gamma \vdash_{c'} E_2 :^P \{x \leftarrow_{c'} [E_1]_{c'}^{\bar{1}}\} T_2}{B; \Gamma \vdash_c [(E_1, E_2)]_{c'}^{\Sigma x : T_1. T_2} \longrightarrow ([E_1]_{c'}^{\bar{1}}, [E_2]_{c'}^{\{x \leftarrow_{c'} [E_1]_{c'}^{\bar{1}}\} T_2})} \text{ (C/econv.col.pair)}$$

$$\frac{B; \Gamma \vdash_{c_2} E :^P T_2 \quad B; \Gamma \vdash_{c_1 \cap c_2} T_2 : \lambda \quad B; \Gamma \vdash_{c_1} T_2 <: T_1 \quad B; \Gamma \vdash_{c \cap c_1} T_1 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [[E]_{c_2}^{\bar{1}}]_{c_1}^{\bar{1}} \longrightarrow [E]_{c_1 \cup c_2}^{\bar{1}}} \text{ (C/econv.col.merge)}$$

$$\frac{B; \Gamma \vdash_{c'} E' :^P S(E) \quad B; \Gamma \vdash_{c \cap c'} E :^P T \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E']_{c'}^{\text{S}(E)} \longrightarrow E} \text{ (C/econv.col.sing)}$$

$$\frac{B; \Gamma, x :_c T_0 \vdash_{c'} E_1 :^{\gamma} T_1 \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 \longrightarrow T'_1 \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 : \lambda}{B; \Gamma \vdash_c (\lambda x : T_0. E_1 !!_{c'} T_1) \longrightarrow (\lambda x : T_0. E_1 !!_{c'} T'_1)} \text{ (C/econv.cong.fun.seal)}$$

IV.6 Dynamic typing and distributed programs \mathcal{D}

IV.6.1 Dynamic typing

[Sorry, this fragment has not been translated yet.]

IV.6.2 Formalisation

IV.6.2.1 Syntax

We define a new language, system \mathcal{D} , which is a conservative extension of system \mathcal{C} . The new features are the type **DYN** and two constructors and a destructor for this type.

$T ::=$	type	
...		
DYN		dynamically typed values
$E ::=$	expression	
...		
dyn E at T		dynamic
dynned E at T		universal dynamic
undyn E at T else E'		dynamic type verification

IV.6.2.2 Reduction

The universal dynamic of a value is a value. The new constructors and destructor are evaluation contexts.

$V ::=$	quasi-value	
...		
dynned V^\bullet at T		universal dynamic
$V^c ::=$	value in c	
...		
dynned V^\bullet at T		universal dynamic
$C_{c'}^c ::=$	evaluation context with inner colour c' and outer colour c	
...		
dyn $__$ at T		dynamic
dynned $__$ at T		universal dynamic, when $c' = \bullet$
undyn $__$ at T else E'		dynamic type verification

We saw in section IV.6.1.4 how to produce a universal dynamic from a dynamic. Evaluating a dynamic type verification can either result in accepting the underlying value if the types are compatible, or evaluating the alternate expression otherwise. A new rule lets a bracket be pushed into a dynamic. A coloured bracket around a universal dynamic can simply be erased, since its contents are already protected.

$$\begin{aligned}
 \text{dyn } V^c \text{ at } T &\longrightarrow_c \text{dynned } [V^c]_c^{\text{conc}_c^B(T)} \text{ at } \text{conc}_c^B(T) && (\mathcal{D}/\text{ered.dyn}) \\
 B \vdash \text{undyn } (\text{dyn } V^c \text{ at } T) \text{ at } T' \text{ else } E' &\longrightarrow_c B \vdash \begin{cases} V^c & \text{if } B; \text{nil} \vdash_c T <: T' \\ E' & \text{otherwise} \end{cases} && (\mathcal{D}/\text{ered.undyn}) \\
 [\text{dynned } V^\bullet \text{ at } T]_{c'}^{\text{DYN}} &\longrightarrow_c \text{dynned } V^\bullet \text{ at } T && (\mathcal{D}/\text{ered.col.dynned})
 \end{aligned}$$

IV.6.2.3 Typing

As we saw in section IV.6.1.4, ordinary dynamics $\text{dyn } E \text{ at } T$, but universal dynamics $\text{dynned } E \text{ at } T$ (which are only useful with E pure) are pure. We do not wish to have to manage dynamic typing inside a compiler¹⁹, thus we declare that any expression of the form $\text{undyn } E \text{ at } T \text{ else } E'$ is impure.

$$\frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{DYN} : K} \text{ (D/tok.base.dyn)} \quad \frac{B; \Gamma \vdash_c E :^\gamma T \quad B; \Gamma \vdash_c T : \lambda}{B; \Gamma \vdash_c \text{dyn } E \text{ at } T :^I \text{DYN}} \text{ (D/et.dyn)}$$

$$\frac{B; \Gamma \vdash_\bullet E :^P T \quad B; \Gamma \vdash_\bullet T : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{dynned } E \text{ at } T :^P \text{DYN}} \text{ (D/et.dynned)}$$

$$\frac{B; \Gamma \vdash_c E :^\gamma \text{DYN} \quad B; \Gamma \vdash_c E' :^\gamma T}{B; \Gamma \vdash_c \text{undyn } E \text{ at } T \text{ else } E' :^I T} \text{ (D/et.undyn)}$$

Since the language has a new constructor, we need corresponding conversion rules: congruence rules to rewrite the arguments of the constructor, and a bracket pushing rule (reflecting the reduction rule (D/ered.col.dynned)).

$$\frac{B; \Gamma \vdash_\bullet E \longrightarrow E' \quad B; \Gamma \vdash_\bullet T : \lambda \quad B; \Gamma \vdash_\bullet E :^P T \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{dynned } E \text{ at } T \longrightarrow \text{dynned } E' \text{ at } T} \text{ (D/econv.cong.dynned.e)}$$

$$\frac{B; \Gamma \vdash_\bullet T \longrightarrow T' \quad B; \Gamma \vdash_\bullet T : \lambda \quad B; \Gamma \vdash_\bullet E :^P T \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{dynned } E \text{ at } T \longrightarrow \text{dynned } E \text{ at } T'} \text{ (D/econv.cong.dynned.t)}$$

$$\frac{B; \Gamma \vdash_\bullet E :^P T \quad B; \Gamma \vdash_\bullet T : \lambda \quad B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\text{dynned } E \text{ at } T]_c^{\text{DYN}} \longrightarrow \text{dynned } E \text{ at } T} \text{ (D/econv.col.dynned)}$$

The dynamisation function Our typing rules let us write any monomorphic dynamisation function $\lambda x : T. \text{dyn } x \text{ at } T$, with the type $T \rightarrow^I \text{DYN}$. They also let us write the polymorphic dynamisation function $\lambda t : \text{TYPE}^l. \lambda x : \text{Typ } t. \text{dyn } x \text{ at } \text{Typ } t$, with the type $\Pi t : \text{TYPE}^l. {}^P \text{Typ } t \rightarrow^I \text{DYN}$. Applying one of these dynamisation function produces a value of the form $\text{dynned } [V]_c^{\text{conc}^B(T)} \text{ at } \text{conc}_c^B(T)$ where c is the ambient colour.

IV.6.3 Communication inter-machines

IV.6.3.1 Introduction

The present dissertation was motivated by the need for dynamic type-checkin in distributed programs with abstract types. The present chapter has so far mainly dealt with abstract types, and we have now introduced dynamic type-checking. We now add the last ingredient: inter-machine communication.

As in chapter II, we assume the availability of some serialisation mechanism to send values between programs running on different machines. In the present discussion, we deal with networked programs, but many considerations also apply to time- rather than space-separated programs, i.e., a program writing data to persistent storage and another program later reading the data.

In order for a value sent by a machine A to be correctly received and decoded on a machine B , the programs running on the two machines must agree on their interpretations of the bit strings they exchange. We assume that all programs are written in the same language and use the same serialisation library, so that it suffices to ensure that the exchanged values do not depend on any manner of environment that is not shared by the two machines. Our semantics does contain one

¹⁹Not only would it be useless, it would also work around a putative stratification (see section V.3.1.1).

machine-dependent element: abstract types defined on one machine may not be available on another machine.

As a first step, we will make the soundness-safe assumption that abstract types defined on one machine are distinct from abstract types defined on any other machine, i.e., abstract types are globally fresh. We spent most of chapter II studying how to lift this restriction, and we will see in section IV.6.3.5 how to integrate these ideas into system \mathcal{D} .

IV.6.3.2 Communication and colours

In this section, we assume the existence of two primitives `send` and `recv` for respectively sending and receiving a value. Communication can take place on a network, via temporary storage or by any other means. More precisely, since we are working in a typed language, we will assume two type-indexed families of primitives `sendT` and `recvT`, the type T being that of transmitted values; their types are `sendT : T →I UNIT` et `recvT : UNIT →I T`. In order for communications to respect typing, the communication protocol must ensure that values sent by `sendT` will only ever be received by `recvT'` when it can be guaranteed that any value of type T also has the type T' , which we model with the constraint $T <: T'$.

One thorny issue is that T and T' live in different contexts: the sending locus and the reception locus may have different knowledge of abstract types. In our framework, this means that the ambient colour might differ between sending and reception. As we saw in section IV.5.3.1, the colour influences both the validity and the semantics of a type. This also applies to the transmitted value, which may have the type T in the sending colour c without having the type T , or indeed any type, in the reception colour c' .

One way to ensure the safety of communication is to index the primitives by a colour as well as a type, i.e., `sendcT` and `recvc'T'`, and require that the communication protocol ensure colour compatibility $c \subseteq c'$ (or rather more precisely $\Gamma \vdash_{c'} c$ transparent) as well as type compatibility $\Gamma \vdash_c T <: T'$. However adapting a communication protocol to ensure colour compatibility is not straightforward, all the less as the ambient colour of a value may change as it is passed around (whereas the type annotations T and T' are usually known statically).

One way to prevent any incompatibility from occurring is to require that the sending colour be empty, in other words that the sending type T as well as the transmitted value be *universal*. We will study how to achieve this in section IV.6.3.3.

If we wish to transmit values between arbitrary colours, they need to be protected. We encountered a similar situation in section IV.5.2.2: given a value V and a type T , we need to construct a value that is “equivalent” to V and has a type “equivalent” to T in the empty colour \bullet . The solution is to use concretisation and send $[V]_c^{\text{conc}_c^B(T)}$ instead of V , where $\text{conc}_c^B(T)$ is the type T with uses of c expanded out. The extra coloured bracket protects V by bestowing upon it any type equation that it may need. Having obtained the universal value $[V]_c^{\text{conc}_c^B(T)}$, we can safely send it to any receiver for the type $\text{conc}_c^B(T)$.

IV.6.3.3 Universals

[Sorry, this fragment has not been translated yet.]

IV.6.3.4 Nonce sharing

In section IV.6.3.2, we discussed how to send values between different colours. We glossed over the fact that a colour is only defined in a certain lexis — comparing colours defined in different lexes, or

transmitting a colour from one lexis to another, does not a priori make sense. However in a network of machines, each machine would have its own lexis.

We modelled the execution of a program (consisting of a single thread running on a given machine) by a reduction relation of the form $B \vdash E \longrightarrow_{\bullet} B' \vdash E'$ (at the top level, outside of any bracket, the ambient colour is empty). An immediate generalisation to networked programs leads us to consider a family of reductions $B_i \vdash E_i \longrightarrow_{\bullet} B'_i \vdash E'_i$ where the index i represents the machine on which the reduction takes place. In this model, communication must take into account the change of lexis from B_i to B_j as well as the colour change.

Recall that a lexis is a set of nonces (plus some information about these nonces), and each nonce that is added to the lexis is freshly created (by the rule $(\mathcal{C}/ered.seal)$ and globally unique). Two lexes B_1 and B_2 formed on different machines are therefore disjoint; it is clear that (B_1, B_2) , or indeed any lexis made by interleaving the elements of B_1 with the elements of B_2 is also a well-formed lexis, with the same information stored for each nonce as in B_1 or B_2 . In the metatheory, we can just merge lexes and model the evolution of a networked program by a reduction relation of the form

$$B \vdash E_1 \parallel \dots \parallel E_n \longrightarrow B' \vdash E'_1 \parallel \dots \parallel E'_n$$

(where $E_1 \parallel \dots \parallel E_n$ notes the parallel composition of n expressions, each running on its own machine).

This model makes considerations about mixing lexes moot as far as the metatheory is concerned. However, in practice, requiring each nonce to be broadcast as soon as it is created would be prohibitively expensive (and might be impossible in networks with complex dynamic topologies). Fortunately one can easily implement the shared lexis model by considering that each machine only has a partial copy of the global lexis at any time, and requiring every transmission of a value to also contain any information necessary to reconstruct the parts of the lexis that the value depends on (that is, the nonces contained in a value as well as their dependencies). Thus the lexis is spread around lazily. Note that although nonce creation requires the generation of a globally unique name, this does not in practice require synchronisation: it suffices that each machine have a globally unique name that can be included in the nonce, which is the case in most distributed systems.

IV.6.3.5 Static sealing and hashes

Nonces are singularised identities in the sense of section II.6.1.2, as a fresh nonce is generated whenever a new family of abstract types is created by evaluating a dynamic sealing construct $E \text{ !! } T$. In section IV.4.4.2, we presented system \mathcal{W} , which has another notion of sealing, namely static sealing $E \text{ :: } T$. Unlike dynamic sealing, static sealing creates a new family of abstract types once and for all at program compile- or initialisation-time, and thus requires an identity to be generated at the corresponding time.

In a distributed environment, there are several choices as to when to generate identities for statically sealed modules. The two main possibilities, compile- and initialisation-time, give different results.

Generating stamps at compile-time [Mac84] is one traditional way of obtaining comparable designations of abstract types. This is not suitable when the identity of a type depends on the run-time behaviour of the program, but this is never the case with our static sealing. Some module systems for distributed programs [Sew01] explicitly allow for abstract type generation at compile-time. This feature has a grave practical defect, namely the impossibility of reconstructing a program from its source alone. If two instances of the same program are deployed, they will only have compatible types if they stem from the same compilation, not if the program was distributed in source form. For this reason, we choose not to support any way to generate module identities at compile-time.

Generating new identities at program-initialisation time allows for less compatibility than at compile-time. However the behaviour is easily predictable and reproducible: any excessive generativity can be spotted in testing. Thus we propose that this is a viable semantics for static sealing.

None of the generation semantics described so far allows sharing abstract types between independently compiled instances of the same program (let alone independently deployed instances of a program component). Yet most cases where static sealing is used — often to enforce data structure invariants — correspond to cases where structural module identities, i.e., *hashhashes* are desired (see section II.3). It is therefore natural to designate statically sealed modules by their hash. The identities are defined by a purely mathematical computation and therefore reproducible at will. As in HAT (see section III.2.7.5), we can see hashes as unifying separate definitions of the “same” module on different machines.

We shall not describe hash formation for system \mathcal{W} formally here. This construction requires that the statically sealed module be lifted from its local potentially-generative context as described in section IV.4.4.5. Note that in system \mathcal{W} , unlike in HAT, the identity of a type is an arbitrarily-sized term which may mention more than one hash. For example, if f is a statically sealed module with the signature $\Pi x : T_0. \Sigma t : \text{TYPE}. T_1$ that is applied to a dynamically sealed module which was given the nonce a , the identity of the type field in the resulting module is $\pi_1(h a)$ where h is the hash of the functor.

IV.7 Conclusion

Summary In the present chapter we presented a description language called TOPHAT for a module system for an ML-like language. The main features of this language are:

- structures and functors, whose types are respectively dependent sums and dependent products;
- a way to test the equivalence of two modules, and propagate knowledge of such an equivalence, using singleton signatures;
- abstract types can be defined by sealing a module, and an effect system determines which expressions remain comparable;
- an reduction abstraction-preserving, thanks to coloured brackets;
- a dynamic type-checking construct that does not depend on the program context.

Soundness The most basic requirement for a type system is that it for the proposed execution mechanism. Appendix B contains a soundness proof for TOPHAT, classically formulated as two theorems: type preservation by reduction (?? ()) and progress of well-typed expressions (?? ()).

Decidability of type-checking Another expected property of a type system for a programming language is decidability, i.e., we would like an algorithm for deciding whether a given expression has a given type²⁰. In particular, we would need to decide when two types are equivalent. Decision procedures exist for weaker type systems, in particular the one proposed by Dreyer, Crary and Harper [DCH03]. However their algorithm does not easily generalise to our system, and we regretfully leave the question open.

²⁰**Type inference** would in fact be desirable. However inference is known to be undecidable in much weaker type systems such as system F. With the type annotations that we require in the syntax, in particular on function arguments, type reconstruction might not be substantially harder than verification.

Chapter V

Conclusion

V.1 Summary

[Sorry, this fragment has not been translated yet.]

V.2 Related work

V.2.1 Theoretic considerations

[[CM88], [OTCP90]]

V.2.2 Programming languages

[Modula-3, Java, .NET, Objective Caml]

V.2.3 Acute and HashCaml

[Sorry, this fragment has not been translated yet.]

V.2.4 Alice ML

Rossberg's work is to my knowledge the only other in-depth treatment of the main topic of this dissertation. Interestingly, my and his independent study of the problem led us towards the same tools.

Rossberg's first step [Ros03] was to use coloured brackets [ZGM99] to keep track of abstract types at run-time and obtain an abstraction-preserving reduction relation, in a manner similar to our HAT [LPSW03]. Our theories differ in that Rossberg's approach is purely generative: abstract types created on different machines are incompatible.

Rossberg also studied the generalisation from simple modules to a full-fledged ML module calculus [Ros07]. His implementation builds on Alice ML [PSL]. Rossberg defines the $\lambda_{SA}^{\omega}\psi$ -calculus, which models the core of Alice ML. This calculus includes a construct that defines an abstract type (§10.5–10.7), and he shows that this is equivalent to ML-like module sealing (specifically, see my section IV.4.4.2). An abstract type is identified by a type variable α with an *abstraction kind* (§11.3). An abstraction kind $A(\tau)$ is similar to a singleton kind $S(\tau)$ but only the singleton kind allows implicit conversion between α and τ . Abstraction-kinded type variables play the same role as my lexis-stored nonces. The type system of $\lambda_{SA}^{\omega}\psi$ -calculus allows explicit conversions between

an abstract type and its representation type anywhere in the program, whereas I materialise such conversions with coloured brackets.

Rossberg proves an *opacity* property (§12.9): a program that does not contain any explicit conversion between an abstract type and its representation is parametric with respect to said representation. Rossberg also proposes a mechanism to seal functors (§13), allowing for both applicative and generative functors. Given the complexity of both systems, I leave to future work a comparison between the expressivity of $\lambda_{\text{SA}}^{\omega}$ -calculus with functors and that of TOPHAT.

V.3 Future work

V.3.1 Improvements to the theory

V.3.1.1 Stratification

[indexing TYPE with a universe]

V.3.1.2 One or two language levels?

[Sorry, this fragment has not been translated yet.]

V.3.1.3 Effect analysis

[Sorry, this fragment has not been translated yet.]

V.3.1.4 Colours and brackets

[Sorry, this fragment has not been translated yet.]

V.3.1.5 Decidability of type-checking

[Sorry, this fragment has not been translated yet.]

V.3.1.6 Parametricity

[Sorry, this fragment has not been translated yet.]

V.3.2 Supplementary features

V.3.2.1 Field names and width subsignaturing

[Sorry, this fragment has not been translated yet.]

V.3.2.2 Towards a programming language

[polymorphism; recursion; libraries]

V.3.2.3 Generic programming

[Sorry, this fragment has not been translated yet.]

V.3.2.4 Security

[Sorry, this fragment has not been translated yet.]

V.3.3 Implementation

V.3.3.1 Hash computation

[Sorry, this fragment has not been translated yet.]

V.3.3.2 Typing `TOPHAT`

[Sorry, this fragment has not been translated yet.]

V.3.3.3 Integration into Objective Caml: the module system

Adding named structure fields and width subtyping to `TOPHAT` as described in section V.3.2.1 yields a language that covers all the features of the module calculus of Objective Caml $[L^+]$. But is our language compatible, i.e., is it a conservative extension of Objective Caml?

The answer is a qualified “no”. There are programs that Objective Caml accepts and we reject, because Objective Caml treats every functor as applicative, even if its body contains side effects. This is unacceptable in `TOPHAT` as applications of applicative functors must be able to be statically evaluated. One way to improve compatibility would be to introduce a notion of separation (in the sense of separability [Dre05] as discussed in section IV.4.2.3). It is however debatable whether this is desirable: treating a functor whose application has side effects as applicative does not break structural typing but does not fully respect abstraction. We prefer to treat any functor whose body has side effects as generative because when applicativity is desired, the body is usually pure ([Dre05, RRS]). For example, all the functors in the standard library of Objective Caml have a pure body (mainly consisting of type definitions and immediate functions, as well as a few data structure values).

The existing sealing of Objective Caml should be considered a static sealing (see section IV.4.4.2). It would be desirable to add a dynamic sealing construct. Another necessary extension is syntax to mark a functor as generative (i.e., a purity annotation on functor types), in order for all signatures to be expressible in the source language.

In addition to examining the module language, we need to check for incompatibilities with the core language. We discussed polymorphism in section V.3.2.2. We can freely extend `TOPHAT` with impure constructs; a safe choice is to make almost all core expressions impure. The main requirements with respect to purity are that projecting a field of a module and immediate functions must be considered pure. In fact, Objective Caml (like any implementation of Standard ML) already performs a suitable purity analysis, in order to check the value restriction for polymorphism [Wri95, Gar04].

V.3.4 Applications of dynamic typing

[Sorry, this fragment has not been translated yet.]

V.3.4.1 The JoCaml name server

The JoCaml “name server” was one of the main motivations of this work. The JoCaml language [FLFS07, MM01] is statically typed, including communications [FLMR97]. However this result only

applies inside a single program instance: when two separate instances communicate, the fact that the value sent by one instance has the type expected by the other instance cannot result solely from adherence to a protocol that only allows for verification inside a single instance, which static typechecking is.

The recommended programming methodology for JoCaml keeps unsafe interactions to a minimum: one instance publishes a communication channel of an agreed-upon type, and other instances can send values (including other channels) over this initial channel, all communications but the initial reception of the public channel being type-safe. The JoCaml standard library provides a `Ns` module to assist in equipping deployed programs with a name server. This name server is a particular program instance which acts as a database for communication channels (the names in question). Participating instances can publish their entry points by uploading them to the name server. A program instance that wants to join the network can query the name server to obtain a channel to send data on. The only type-checking that must take place at run-time is that performed by new participants as they check that the data returned by the name server matches their typing expectations (the actual verification may be performed by the name server itself; in any case the name server must retain typing information for the values that it stores).

Appendix A

Formal definition of TOPHAT

This appendix is a précis of the language TOPHAT, which is identical to \mathcal{D} of chapter IV.

$E ::=$	expression
$x \mid y \mid t \mid \dots$	variables
$()$	unit value
$\text{false} \mid \text{true}$	boolean (generically bv)
$0 \mid 1 \mid \dots$	integer (generically \underline{n})
$\langle T \rangle$	type field
(E_1, E_2)	pair
$\pi_i E$	projection ($i \in \{1, 2\}$)
$\lambda x : T. E$	lambda-abstraction
$E_1 E_2$	application
$\text{let } x = E_0 \text{ in } E : T$	local binding
$E !!_c T$	sealed and coloured module
$[E]_c^T$	coloured bracket
$\text{dyn } E \text{ at } T$	dynamic
$\text{dynned } E \text{ at } T$	universal dynamic
$\text{undyn } E \text{ at } T \text{ else } E'$	dynamic type verification
$T ::=$	type
UNIT	unit
BOOL	booleans
INT	integers
$\text{Typ } E$	projection from a type field
$\Sigma x : T_1. T_2$	dependent sum (also written $T_1 * T_2$ when $x \notin \text{fv } T_2$)
$\Pi x : T_0. {}^\gamma T_1$	dependent product (also written $T_1 \rightarrow^\gamma T_2$ when $x \notin \text{fv } T_1$)
$S(E)$	singleton
TYPE^K	abstract type field
$\langle A \rangle$	abstract type
DYN	dynamically typed values
$K ::=$	kind
$\{$	monomorphic (fully specified)
$*$	polymorphic (partially specified)

$A ::=$	module component
a	nonce
$A E$	application
$\pi_i A$	projection ($i \in \{1, 2\}$)
$\gamma ::=$	effect
P	pure
I	impure
$\xi ::=$	primary colour
a	nonce
x	variable
$c ::=$	colour
\bullet	empty colour (also written $\{\}$)
$\{a_1, \dots, a_k, x_1, \dots, x_k\}$	finite set of primary colours
$B ::=$	lexis
nil	empty
$B, a = E :_{c_0} T$	nonce a with implemented by E with the signature T
$\Gamma ::=$	environnement
nil	empty
$\Gamma, x :_c T$	binding of the variable x
$J ::=$	local judgement right-hand side
ok	environment correction
$T : K$	type kinding (generalising $T : *$)
$T \longrightarrow T'$	typing conversion
$T \equiv T'$	convertibility equivalence on types
$E \longrightarrow E'$	expression conversion
$E \equiv E'$	convertibility equivalence on expressions
$T_1 <: T_2$	subtyping
$c_0 \text{ transparent}$	colour transparency
$A \triangleright E : T$	component revelation
$A \longrightarrow A'$	component conversion
$A \equiv A'$	convertibility equivalence on components
$E :_{\gamma} T$	expression typing
$V ::=$	quasi-value
$() \mid \text{bv} \mid \underline{n}$	constant
$\langle T \rangle$	type field
(V_1, V_2)	pair
$\lambda x : T. E$	lambda-abstraction
$[V]_{c'}^{\langle A \rangle}$	potentially abstraction-making coloured bracket
$\text{dynned } V^\bullet \text{ at } T$	universal dynamic

$V^c ::=$	$() \mid \text{bv} \mid \underline{n}$ $\langle T \rangle$ (V_1^c, V_2^c) $\lambda x : T. E$ $[V^{c'}]_{c'}^{(A^{V^{c \cap c'}})}$ dynned V^\bullet at T	value in c constant type field pair lambda-abstraction coloured bracket, if $A^{V^{c \cap c'}}$ is abstract in c but concrete in c' universal dynamic
$A^V ::=$	a $A^V V$ $\pi_i A^V$	component value nonce application to a quasi-value projection ($i \in \{1, 2\}$)
$A^{V^c} ::=$	a $A^{V^c} V^c$ $\pi_i A^{V^c}$	abstract component in c nonce, if opaque in c application of a functor to a value projection ($i \in \{1, 2\}$)
$C_{c'}^c ::=$	$E_1 _$ $_ V_2$ $(_, E_2)$ $(V_1, _)$ $\pi_i _$ $\text{let } x = _ \text{ in } E : T$ $_ !!_{c_1} T$ $[_]_{c'}^T$ $[V^{c_1}]_{c_1}^{\text{Typ}} _$ $\text{dyn } _ \text{ at } T$ $\text{dynned } _ \text{ at } T$ $\text{undyn } _ \text{ at } T \text{ else } E'$	evaluation context with inner colour c' and outer colour c function argument applied function first component of a pair second component of a pair projection ($i \in \{1, 2\}$) local bound sealing coloured bracket type field on a bracket, when $c' = c \cap c_1$ dynamic universal dynamic, when $c' = \bullet$ dynamic type verification

$\mathbf{self}^{\text{BT}}(A) = \text{BT}$ $\mathbf{self}^{\Sigma x : T_1. T_2}(A) = \Sigma x : \mathbf{self}^{T_1}(\pi_1 A). \mathbf{self}^{T_2}(\pi_2 A)$ $\mathbf{self}^{\Pi x : T_0. P T_1}(A) = \Pi x : T_0. P(\mathbf{self}^{T_1}(A x))$ $\mathbf{self}^{\Pi x : T_0. I T_1}(A) = \Pi x : T_0. I T_1$ $\mathbf{self}^{S(E')}(A) = S(E)$ $\mathbf{self}^{\text{TYPE}^K}(A) = S(\langle\langle A \rangle\rangle)$	if BT is a base type (UNIT, BOOL, INT, DYN)
$\mathbf{conc}_c^B(\langle A_1 \rangle) = \text{Typ reveal}^B(A_1)$ $\mathbf{conc}_c^B(\langle A_1 \rangle) = \langle A_1 \rangle$ $\mathbf{conc}_c^B([E]_{c'}^T) = [E]_{c'}^{\mathbf{conc}_{c \cap c'}^B(T)}$ (other cases by simple induction)	if $\text{underl}(A_1) \in c$ if $\text{underl}(A_1) \notin c$

$$\begin{aligned}
 \{x \leftarrow_{c_0} E_0\}x &= E_0 \\
 \{x \leftarrow_{c_0} E_0\}y &= y && \text{si } y \neq x \\
 \{x \leftarrow_{c_0} E_0\}[E]_c^T &= [\{x \leftarrow_{c_0} E_0\}E]_{\{x \leftarrow_{c_0} E_0\}c}^{\{x \leftarrow_{c_0} E_0\}T} \\
 \{x \leftarrow_{c_0} E_0\}c &= (c \setminus \{x\}) \cup c_0 && \text{si } x \in c \\
 \{x \leftarrow_{c_0} E_0\}c &= c && \text{si } x \notin c
 \end{aligned}$$

(other cases follow the usual notion of capture-avoiding substitution)

$$\begin{aligned}
 \mathbf{underl}(a) &= a \\
 \mathbf{underl}(A E) &= \mathbf{underl}(A) \\
 \mathbf{underl}(\pi_i A) &= \mathbf{underl}(A) \\
 \mathbf{reveal}^B(a) &= E && \text{where } a = E :_c T \in B \\
 \mathbf{reveal}^B(A E) &= (\mathbf{reveal}^B(A)) E \\
 \mathbf{reveal}^B(\pi_i A) &= \pi_i (\mathbf{reveal}^B(A))
 \end{aligned}$$

$$\begin{aligned}
 &\frac{B; \text{nil} \vdash_{c_0} E :^P T \quad \text{when } a \notin \mathbf{dom} B}{B, a = E :_{c_0} T; \text{nil} \vdash_{\bullet} \text{ok}} (\mathbf{envok.a}) && \frac{B; \Gamma \vdash_{c'} \text{ok} \quad \text{when } a = E :_{c_0} T \in B \wedge c_0 \subseteq c'}{B; \Gamma \vdash_{c' \cup \{a\}} \text{ok}} (\mathbf{envok.c.a}) \\
 &\frac{B; \Gamma \vdash_{c'} \text{ok} \quad \text{when } x :_{c_0} T \in \Gamma}{B; \Gamma \vdash_{c' \cup \{x\}} \text{ok}} (\mathbf{envok.c.x}) && \frac{}{\text{nil}; \text{nil} \vdash_{\bullet} \text{ok}} (\mathbf{envok.nil}) && \frac{B; \Gamma \vdash_c T : * \quad \text{when } x \notin \mathbf{dom} \Gamma}{B; \Gamma, x :_c T \vdash_{\bullet} \text{ok}} (\mathbf{envok.x}) \\
 &\frac{B; \Gamma \vdash_c \text{ok} \quad \text{when } \Gamma = (\Gamma_0, x :_{c_0} T, \Gamma_1) \wedge x \in c}{B; \Gamma \vdash_c \xi \text{ transparent}} (\mathbf{vis.env}) && \frac{B; \Gamma \vdash_c \text{ok} \quad \text{when } \xi \in c}{B; \Gamma \vdash_c \xi \text{ transparent}} (\mathbf{vis.in}) && \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \bullet \text{ transparent}} (\mathbf{vis.o}) \\
 &&& \frac{B; \Gamma \vdash_c c_1 \text{ transparent} \quad B; \Gamma \vdash_c c_2 \text{ transparent}}{B; \Gamma \vdash_c c_1 \cup c_2 \text{ transparent}} (\mathbf{vis.union})
 \end{aligned}$$

$$\begin{aligned}
 &\frac{B; \Gamma \vdash_c c_0 \text{ transparent} \quad \text{when } a = E :_{c_0} T \in B}{B; \Gamma \vdash_c a \triangleright E : T} (\mathbf{ac.a}) && \frac{B; \Gamma \vdash_c A \triangleright E : \Pi x : T_0. {}^P T_1 \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c A E_0 \triangleright E E_0 : \{x \leftarrow_c E_0\} T_1} (\mathbf{ac.app}) \\
 &\frac{B; \Gamma \vdash_c A \triangleright E : \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_1 A \triangleright \pi_1 E : T_1} (\mathbf{ac.proj.1}) && \frac{B; \Gamma \vdash_c E_1 :^P S(\pi_1 E) \quad B; \Gamma \vdash_c A \triangleright E : \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_2 A \triangleright \pi_2 E : \{x \leftarrow_c E_1\} T_2} (\mathbf{ac.proj.2})
 \end{aligned}$$

$$\begin{aligned}
 &\frac{B; \Gamma \vdash_c A \triangleright E : \text{TYPE}^K}{B; \Gamma \vdash_c (A) : K} (\mathbf{tok.abs}) && \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{BOOL} : \lambda} (\mathbf{tok.base.bool}) && \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{DYN} : K} (\mathbf{tok.base.dyn}) \\
 &\frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{INT} : \lambda} (\mathbf{tok.base.int}) && \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{UNIT} : \lambda} (\mathbf{tok.base.unit}) && \frac{B; \Gamma \vdash_c E :^P \text{TYPE}^K}{B; \Gamma \vdash_c \text{Typ } E : K} (\mathbf{tok.field})
 \end{aligned}$$

$$\begin{array}{c}
\frac{B; \Gamma \vdash_c T' : K'}{B; \Gamma, x :_c T' \vdash_{c \cup \{x\}} T'' : K''} \text{(tok.fun.I)} \quad \frac{B; \Gamma \vdash_c T' : K'}{B; \Gamma, x :_c T' \vdash_{c \cup \{x\}} T'' : K''} \text{(tok.fun.P)} \\
\frac{B; \Gamma \vdash_c T' : K'}{B; \Gamma \vdash_c \Pi x : T'. T'' : K''} \text{(tok.pair)} \quad \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{TYPE}^K : *} \text{(tok.type)} \quad \frac{B; \Gamma \vdash_c E :^P T}{B; \Gamma \vdash_c S(E) : \lambda} \text{(tok.sing)} \\
\frac{B; \Gamma \vdash_c T : K'}{\text{when } K' \leq K} \text{(tok.sub)} \\
\frac{B; \Gamma \vdash_c T : K'}{B; \Gamma \vdash_c T : K} \\
\frac{B; \Gamma \vdash_c E_0 \longrightarrow E'_0 \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c A \triangleright E : \Pi x : T_0. T_1} \text{(aconv.cong.app.arg)} \quad \frac{B; \Gamma \vdash_c A \longrightarrow A' \quad B; \Gamma \vdash_c A \triangleright E : \Pi x : T_0. T_1}{B; \Gamma \vdash_c E_0 :^P T_0} \text{(aconv.cong.app.fun)} \\
\frac{B; \Gamma \vdash_c A \longrightarrow A' \quad B; \Gamma \vdash_c A \triangleright E : \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_i A \longrightarrow \pi_i A} \text{(aconv.cong.proj)} \\
\frac{B; \Gamma \vdash_c A_1 \longrightarrow A_2}{B; \Gamma \vdash_c A_1 \equiv A_2} \text{(aeq.conv)} \quad \frac{B; \Gamma \vdash_c A \triangleright E : T}{B; \Gamma \vdash_c A \equiv A} \text{(aeq.refl)} \quad \frac{B; \Gamma \vdash_c A_2 \equiv A_1}{B; \Gamma \vdash_c A_1 \equiv A_2} \text{(aeq.sym)} \quad \frac{B; \Gamma \vdash_c A_1 \equiv A_2 \quad B; \Gamma \vdash_c A_2 \equiv A_3}{B; \Gamma \vdash_c A_1 \equiv A_3} \text{(aeq.trans)} \\
\frac{B; \Gamma \vdash_c A \longrightarrow A' \quad B; \Gamma \vdash_c A \triangleright E : \text{TYPE}^K}{B; \Gamma \vdash_c \langle A \rangle \longrightarrow \langle A' \rangle} \text{(tconv.cong.abs)} \quad \frac{B; \Gamma \vdash_c E \longrightarrow E' \quad B; \Gamma \vdash_c E :^P \text{TYPE}^*}{B; \Gamma \vdash_c \text{Typ } E \longrightarrow \text{Typ } E'} \text{(tconv.cong.field)} \\
\frac{B; \Gamma \vdash_c T_0 \longrightarrow T'_0 \quad B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} T_1 : *}{B; \Gamma \vdash_c \Pi x : T_0. T_1 \longrightarrow \Pi x : T'_0. T_1} \text{(tconv.cong.fun.arg)} \\
\frac{B; \Gamma \vdash_c T_0 : * \quad B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} T_1 \longrightarrow T'_1}{B; \Gamma \vdash_c \Pi x : T_0. T_1 \longrightarrow \Pi x : T_0. T'_1} \text{(tconv.cong.fun.ret)} \\
\frac{B; \Gamma \vdash_c T_1 \longrightarrow T'_1 \quad B; \Gamma, x :_c T_1 \vdash_{c \cup \{x\}} T_2 : *}{B; \Gamma \vdash_c \Sigma x : T_1. T_2 \longrightarrow \Sigma x : T'_1. T_2} \text{(tconv.cong.pair.1)} \\
\frac{B; \Gamma, x :_c T_1 \vdash_{c \cup \{x\}} T_2 \longrightarrow T'_2 \quad B; \Gamma \vdash_c T_1 : *}{B; \Gamma \vdash_c \Sigma x : T_1. T_2 \longrightarrow \Sigma x : T_1. T'_2} \text{(tconv.cong.pair.2)} \quad \frac{B; \Gamma \vdash_c E \longrightarrow E'}{B; \Gamma \vdash_c S(E) \longrightarrow S(E')} \text{(tconv.cong.sing)} \\
\frac{B; \Gamma \vdash_c A \triangleright E : \text{TYPE}^K \quad B; \Gamma \vdash_c \text{underl}(A) \text{ transparent}}{B; \Gamma \vdash_c \langle A \rangle \longrightarrow \text{Typ } E} \text{(tconv.abs)} \quad \frac{B; \Gamma \vdash_c T : *}{B; \Gamma \vdash_c \text{Typ } \langle T \rangle \longrightarrow T} \text{(tconv.field)} \\
\frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c S(()) \longrightarrow \text{UNIT}} \text{(tconv.unit)}
\end{array}$$

$$\begin{array}{c}
 \frac{B; \Gamma \vdash_c T_1 \longrightarrow T_2}{B; \Gamma \vdash_c T_1 \equiv T_2} \text{(teq.conv)} \quad \frac{B; \Gamma \vdash_c T : *}{B; \Gamma \vdash_c T \equiv T} \text{(teq.refl)} \quad \frac{B; \Gamma \vdash_c T_2 \equiv T_1}{B; \Gamma \vdash_c T_1 \equiv T_2} \text{(teq.sym)} \quad \frac{B; \Gamma \vdash_c T_1 \equiv T_2 \quad B; \Gamma \vdash_c T_2 \equiv T_3}{B; \Gamma \vdash_c T_1 \equiv T_3} \text{(teq.trans)} \\
 \\
 \frac{B; \Gamma \vdash_c E \longrightarrow E' \quad B; \Gamma \vdash_c E :^P T_0 \quad B; \Gamma \vdash_c E_1 :^P \Pi x : T_0. T_1}{B; \Gamma \vdash_c E_1 E \longrightarrow E_1 E'} \text{(econv.cong.app.arg)} \quad \frac{B; \Gamma \vdash_c E \longrightarrow E' \quad B; \Gamma \vdash_c E :^P \Pi x : T_0. T_1}{B; \Gamma \vdash_c E E_0 \longrightarrow E' E_0} \text{(econv.cong.app.fun)} \\
 \\
 \frac{B; \Gamma \vdash_{c'} E \longrightarrow E' \quad B; \Gamma \vdash_{c'} E :^P T \quad B; \Gamma \vdash_{c \cap c'} T : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E]_{c'}^{T'} \longrightarrow [E']_{c'}^{T'}} \text{(econv.cong.col.e)} \quad \frac{B; \Gamma \vdash_{c'} E \longrightarrow E' \quad B; \Gamma \vdash_{c \cap c'} T_1 \longrightarrow T_2 \quad B; \Gamma \vdash_{c \cap c'} T_1 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E]_{c'}^{T_1} \longrightarrow [E]_{c'}^{T_2}} \text{(econv.cong.col.t)} \\
 \\
 \frac{B; \Gamma \vdash_{\bullet} E \longrightarrow E' \quad B; \Gamma \vdash_{\bullet} T : \lambda \quad B; \Gamma \vdash_{\bullet} E :^P T \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{dynned } E \text{ at } T \longrightarrow \text{dynned } E' \text{ at } T} \text{(econv.cong.dynned.e)} \\
 \\
 \frac{B; \Gamma \vdash_{\bullet} T \longrightarrow T' \quad B; \Gamma \vdash_{\bullet} T : \lambda \quad B; \Gamma \vdash_{\bullet} E :^P T \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{dynned } E \text{ at } T \longrightarrow \text{dynned } E \text{ at } T'} \text{(econv.cong.dynned.t)} \quad \frac{B; \Gamma \vdash_c T \longrightarrow T'}{B; \Gamma \vdash_c \langle T \rangle \longrightarrow \langle T' \rangle} \text{(econv.cong.field)} \\
 \\
 \frac{B; \Gamma \vdash_c T_0 \longrightarrow T'_0 \quad B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} E_1 :^Y T_1}{B; \Gamma \vdash_c (\lambda x : T_0. E_1) \longrightarrow (\lambda x : T'_0. E_1)} \text{(econv.cong.fun.arg)} \\
 \\
 \frac{B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} E \longrightarrow E' \quad B; \Gamma, x :_c T_0, y :_{c \cup \{x\}} S(E) \vdash_{c \cup \{y\} \cup \{x\}} E_1 :^Y T_1}{B; \Gamma \vdash_c (\lambda x : T_0. \{y \leftarrow_{c \cup \{x\}} E\} E_1) \longrightarrow (\lambda x : T_0. \{y \leftarrow_{c \cup \{x\}} E'\} E_1)} \text{(econv.cong.fun.body)} \\
 \\
 \frac{B; \Gamma, x :_c T_0 \vdash_{c'} E_1 :^Y T_1 \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 \longrightarrow T'_1 \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 : \lambda}{B; \Gamma \vdash_c (\lambda x : T_0. E_1 !!_{c'} T_1) \longrightarrow (\lambda x : T_0. E_1 !!_{c'} T'_1)} \text{(econv.cong.fun.seal)} \\
 \\
 \frac{B; \Gamma \vdash_c E \longrightarrow E' \quad B; \Gamma \vdash_c E_2 :^P T_2}{B; \Gamma \vdash_c (E, E_2) \longrightarrow (E', E_2)} \text{(econv.cong.pair.1)} \quad \frac{B; \Gamma \vdash_c E \longrightarrow E' \quad B; \Gamma \vdash_c E_1 :^P T_1}{B; \Gamma \vdash_c (E_1, E) \longrightarrow (E_1, E')} \text{(econv.cong.pair.2)} \\
 \\
 \frac{B; \Gamma \vdash_c E \longrightarrow E' \quad B; \Gamma \vdash_c E :^P \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_i E \longrightarrow \pi_i E'} \text{(econv.cong.proj)} \quad \frac{B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} E_1 :^P T_1 \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c (\lambda x : T_0. E_1) E_0 \longrightarrow \{x \leftarrow_c E_0\} E_1} \text{(econv.app)} \\
 \\
 \frac{B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\text{bv}]_{c'}^{\text{BOOL}} \longrightarrow \text{bv}} \text{(econv.col.base.bool)} \quad \frac{B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\underline{n}]_{c'}^{\text{INT}} \longrightarrow \underline{n}} \text{(econv.col.base.int)}
 \end{array}$$

$$\begin{array}{c}
\frac{B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [()]_{c'}^{\text{UNIT}} \longrightarrow ()} \text{(econv.col.base.unit)} \\
\frac{B; \Gamma \vdash_{\bullet} E :^P T \quad B; \Gamma \vdash_{\bullet} T : \lambda \quad B; \Gamma \vdash_{c'} \text{ok} \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\text{dynned } E \text{ at } T]_{c'}^{\text{DYN}} \longrightarrow \text{dynned } E \text{ at } T} \text{(econv.col.dynned)} \\
\frac{B; \Gamma \vdash_{c'} T_0 <: T_2 \quad B; \Gamma, x :_{c'} T_2 \vdash_{c' \cup \{x\}} E :^I T_1 \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\lambda x : T_2. E]_{c'}^{\Pi x : T_0. I T_1} \longrightarrow \lambda x : T_0. E !!_{c' \cup \{x\}} T_1} \text{(econv.col.fun.I)} \\
\frac{B; \Gamma \vdash_{c'} T_0 <: T_2 \quad B; \Gamma, x :_{c'} T_2 \vdash_{c' \cup \{x\}} E :^P T_1 \quad B; \Gamma, x :_{c \cap c'} T_0 \vdash_{(c \cap c') \cup \{x\}} T_1 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [\lambda x : T_2. E]_{c'}^{\Pi x : T_0. P T_1} \longrightarrow \lambda x : T_0. [E]_{c' \cup \{x\}}^{T_1}} \text{(econv.col.fun.P)} \\
\frac{B; \Gamma \vdash_{c_2} E :^P T_2 \quad B; \Gamma \vdash_{c_1 \cap c_2} T_2 : \lambda \quad B; \Gamma \vdash_{c_1} T_2 <: T_1 \quad B; \Gamma \vdash_{c \cap c_1} T_1 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [[E]_{c_2}^{T_2}]_{c_1}^{T_1} \longrightarrow [E]_{c_1 \cup c_2}^{T_1}} \text{(econv.col.merge)} \\
\frac{B; \Gamma \vdash_{c'} E_1 :^P T_1 \quad B; \Gamma \vdash_{c \cap c'} T_1 : \lambda \quad B; \Gamma, x :_{c'} T_1 \vdash_{c' \cup \{x\}} E_2 :^P T_2 \quad B; \Gamma \vdash_{c'} E_2 :^P \{x \leftarrow_{c'} [E_1]_{c'}^{T_1}\} T_2 \quad B; \Gamma, x :_{c \cap c'} T_1 \vdash_{(c \cap c') \cup \{x\}} T_2 : \lambda \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [(E_1, E_2)]_{c'}^{\Sigma x : T_1. T_2} \longrightarrow ([E_1]_{c'}^{T_1}, [E_2]_{c'}^{\{x \leftarrow_{c'} [E_1]_{c'}^{T_1}\} T_2})} \text{(econv.col.pair)} \\
\frac{B; \Gamma \vdash_{c'} E' :^P S(E) \quad B; \Gamma \vdash_{c \cap c'} E :^P T \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E']_{c'}^{S(E)} \longrightarrow E} \text{(econv.col.sing)} \quad \frac{B; \Gamma \vdash_c E_1 :^P T_1 \quad B; \Gamma \vdash_c E_2 :^P T_2}{B; \Gamma \vdash_c \pi_i (E_1, E_2) \longrightarrow E_i} \text{(econv.proj)} \\
\frac{B; \Gamma \vdash_c E :^P \text{TYPE}^*}{B; \Gamma \vdash_c E \longrightarrow \langle \text{Type } E \rangle} \text{(econv.eta.field)} \quad \frac{B; \Gamma \vdash_c E :^P \Pi x : T_0. \gamma T_1}{B; \Gamma \vdash_c E \longrightarrow (\lambda x : T_0. E x)} \text{(econv.eta.fun)} \\
\frac{B; \Gamma \vdash_c E :^P \Sigma x : T_1. T_2}{B; \Gamma \vdash_c E \longrightarrow (\pi_1 E, \pi_2 E)} \text{(econv.eta.pair)}
\end{array}$$

$$\begin{array}{c}
 \frac{B; \Gamma \vdash_c E_1 \longrightarrow E_2}{B; \Gamma \vdash_c E_1 \equiv E_2} \text{(eeq.conv)} \quad \frac{B; \Gamma \vdash_c E :^P T}{B; \Gamma \vdash_c E \equiv E} \text{(eeq.refl)} \quad \frac{B; \Gamma \vdash_c E_2 \equiv E_1}{B; \Gamma \vdash_c E_1 \equiv E_2} \text{(eeq.sym)} \quad \frac{B; \Gamma \vdash_c E_1 \equiv E_2 \quad B; \Gamma \vdash_c E_2 \equiv E_3}{B; \Gamma \vdash_c E_1 \equiv E_3} \text{(eeq.trans)} \\
 \\
 \frac{B; \Gamma \vdash_c T'_0 <: T_0 \quad B; \Gamma, x :_c T'_0 \vdash_{c \cup \{x\}} T_1 <: T'_1 \quad B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} T_1 : *}{\text{when } \gamma \sqsubseteq \gamma' \text{ (tsub.cong.fun)}} \quad \frac{B; \Gamma \vdash_c T_1 <: T'_1 \quad B; \Gamma, x :_c T_1 \vdash_{c \cup \{x\}} T_2 <: T'_2 \quad B; \Gamma, x :_c T'_1 \vdash_{c \cup \{x\}} T'_2 : *}{\text{(tsub.cong.pair)}} \\
 \frac{B; \Gamma \vdash_c \text{ok}}{\text{when } K_1 \leq K_2 \text{ (tsub.cong.type)}} \quad \frac{B; \Gamma \vdash_c T \equiv T'}{B; \Gamma \vdash_c T <: T'} \text{(tsub.eq)} \quad \frac{B; \Gamma \vdash_c T <: T'}{B; \Gamma \vdash_c T <: T''} \text{(tsub.trans)} \\
 \frac{B; \Gamma \vdash_c E :^P T}{B; \Gamma \vdash_c S(E) <: T} \text{(tsub.sing)} \\
 \\
 \frac{B; \Gamma \vdash_c E_1 :^{\gamma_1} \Pi x : T_0. \gamma_2 T \quad B; \Gamma \vdash_c E_0 :^P T_0}{B; \Gamma \vdash_c E_1 E_0 :^{\gamma_1 \sqcup \gamma_2} \{x \leftarrow_c E_0\} T} \text{(et.app)} \quad \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{bv} :^P \text{BOOL}} \text{(et.base.bool)} \quad \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \underline{n} :^P \text{INT}} \text{(et.base.int)} \\
 \\
 \frac{B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c () :^P \text{UNIT}} \text{(et.base.unit)} \quad \frac{B; \Gamma \vdash_{c'} E :^{\gamma} T \quad B; \Gamma \vdash_{c \cap c'} T : \iota \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c [E]_{c'}^T :^{\gamma} T} \text{(et.col)} \quad \frac{B; \Gamma \vdash_c E :^{\gamma} T \quad B; \Gamma \vdash_c T : \iota}{B; \Gamma \vdash_c \text{dyn } E \text{ at } T :^{\dagger} \text{DYN}} \text{(et.dyn)} \\
 \\
 \frac{B; \Gamma \bullet E :^P T \quad B; \Gamma \bullet T : \iota \quad B; \Gamma \vdash_c \text{ok}}{B; \Gamma \vdash_c \text{dynned } E \text{ at } T :^P \text{DYN}} \text{(et.dynned)} \quad \frac{B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} E :^{\gamma} T_1}{B; \Gamma \vdash_c \lambda x : T_0. E :^P \Pi x : T_0. \gamma T_1} \text{(et.fun)} \\
 \\
 \frac{B; \Gamma \vdash_c E_0 :^{\dagger} T_0 \quad B; \Gamma, x :_c T_0 \vdash_{c \cup \{x\}} E :^{\dagger} T \quad B; \Gamma \vdash_c T : *}{B; \Gamma \vdash_c (\text{let } x = E_0 \text{ in } E : T) :^{\dagger} T} \text{(et.let)} \quad \frac{B; \Gamma \vdash_c E_1 :^{\gamma} T_1 \quad B; \Gamma \vdash_c E_2 :^{\gamma} T_2}{B; \Gamma \vdash_c (E_1, E_2) :^{\gamma} T_1 * T_2} \text{(et.pair)} \\
 \\
 \frac{B; \Gamma \vdash_c E :^{\gamma} \Sigma x : T_1. T_2}{B; \Gamma \vdash_c \pi_1 E :^{\gamma} T_1} \text{(et.proj.1)} \quad \frac{B; \Gamma \vdash_c E :^P \Sigma x : T_1. T_2 \quad B; \Gamma \vdash_c E_1 :^P S(\pi_1 E)}{B; \Gamma \vdash_c \pi_2 E :^P \{x \leftarrow_c E_1\} T_2} \text{(et.proj.2)} \quad \frac{B; \Gamma \vdash_{c \cup c'} E :^{\gamma} T \quad B; \Gamma \vdash_c T : *}{B; \Gamma \vdash_c (E !!_{c'} T) :^{\dagger} T} \text{(et.seal)} \\
 \\
 \frac{B; \Gamma \vdash_c T : K}{B; \Gamma \vdash_c \langle T \rangle :^P \text{TYPE}^K} \text{(et.type)} \quad \frac{B; \Gamma \vdash_c E :^{\gamma} \text{DYN} \quad B; \Gamma \vdash_c E' :^{\gamma} T}{B; \Gamma \vdash_c \text{undyn } E \text{ at } T \text{ else } E' :^{\dagger} T} \text{(et.undyn)} \quad \frac{B; \Gamma \vdash_c x \text{ transparent}}{\text{when } x :_{c'} T \in \Gamma} \text{(et.x)} \\
 \\
 \frac{B; \Gamma \vdash_c E :^P T}{B; \Gamma \vdash_c E :^P S(E)} \text{(et.sing)} \quad \frac{B; \Gamma \vdash_c E :^{\gamma} T \quad B; \Gamma \vdash_c T <: T' \quad \text{when } \gamma \sqsubseteq \gamma'}{B; \Gamma \vdash_c E :^{\gamma'} T'} \text{(et.sub)}
 \end{array}$$

$$(\lambda x : T. E) V^c \longrightarrow_c \{x \leftarrow_c V^c\} E \quad (\text{ered.app})$$

$$[\text{bv}]_{c'}^{\text{BOOL}} \longrightarrow_c \text{bv} \quad (\text{ered.col.base.bool})$$

$$[\underline{n}]_{c'}^{\text{INT}} \longrightarrow_c \underline{n} \quad (\text{ered.col.base.int})$$

$$[()]_{c'}^{\text{UNIT}} \longrightarrow_c () \quad (\text{ered.col.base.unit})$$

$$[\text{dynned } V^\bullet \text{ at } T]_{c'}^{\text{DYN}} \longrightarrow_c \text{dynned } V^\bullet \text{ at } T \quad (\text{ered.col.dynned})$$

$$[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. I T_1} \longrightarrow_c \lambda x : T_0. (E !!_{c' \cup \{x\}} T_1) \quad (\text{ered.col.fun.I})$$

$$[\lambda x : T_2. E]_{c'}^{\Pi x : T_0. P T_1} \longrightarrow_c \lambda x : T_0. [E]_{c' \cup \{x\}}^{T_1} \quad (\text{ered.col.fun.P})$$

$$[[V^{c_2}]_{c_2}^{(A_2)}]_{c_1}^{(A_1)} \longrightarrow_c [V^{c_2}]_{c_1 \cup c_2}^{(A_1)} \quad (\text{ered.col.merge})$$

if A_1 et A_2 are both opaque in c_1 but A_2 is concrete in c_2

$$[[V_1^{c'}, V_2^{c'}]_{c'}^{\Sigma x : T_1. T_2}] \longrightarrow_c ([V_1^{c'}]_{c'}^{T_1}, [V_2^{c'}]_{c'}^{\{x \leftarrow_c [V_1^{c'}]_{c'}^{T_1}\} T_2}) \quad (\text{ered.col.pair})$$

$$[V^{c'}]_{c'}^{\text{S}(E)} \longrightarrow_c E \quad (\text{ered.col.sing})$$

$$B \vdash [V^{c'}]_{c'}^{(A)} \longrightarrow_c B \vdash [V^{c'}]_{c'}^{\text{Typ reveal}^B(A)} \quad (\text{ered.colAbs})$$

if $\text{underI}(A) \in c \cap c'$

$$[V^{c'}]_{c'}^{\text{Typ}(T)} \longrightarrow_c [V^{c'}]_{c'}^T \quad (\text{ered.colTyp})$$

$$\frac{E \longrightarrow_{c'} E'}{C_{c'}^{c'} \cdot E \longrightarrow_c C_{c'}^{c'} \cdot E'} \quad (\text{ered.context})$$

$$\text{dyn } V^c \text{ at } T \longrightarrow_c \text{dynned } [V^c]_c^{\text{conc}^B(T)} \text{ at } \text{conc}_c^B(T) \quad (\text{ered.dyn})$$

$$\text{let } x = V^c \text{ in } E : T \longrightarrow_c \{x \leftarrow_c V^c\} E \quad (\text{ered.let})$$

$$\pi_i (V_1^c, V_2^c) \longrightarrow_c V_i \quad (\text{ered.proj})$$

$$B \vdash V^{c \cup c'} !!_{c'} T \longrightarrow_c B, a = V^{c \cup c'} :_{c \cup c'} T \vdash [V^{c \cup c'}]_{c \cup c' \cup \{a\}}^{\text{self}^T(a)} \quad (\text{ered.seal})$$

where a is fresh (i.e., $a \notin \text{dom } B$)

$$B \vdash \text{undyn} (\text{dyn } V^c \text{ at } T) \text{ at } T' \text{ else } E' \longrightarrow_c B \vdash \begin{cases} V^c & \text{if } B; \text{nil} \vdash_c T <: T' \\ E' & \text{otherwise} \end{cases} \quad (\text{ered.undyn})$$

Bibliography

- [AM91] Andrew Appel and David B. MacQueen. Standard ML of new jersey. In J. Maluszynski and M. Wirsing, editors, *Programming Language Implementation and Logic Programming, Proceedings of the 3rd Intn'l Symposium*, volume 528 of *LNCS*, pages 1–13. Springer Verlag, 1991.
- [CM88] L. Cardelli and D. MacQueen. Persistence and type abstraction. In Malcolm P. Atkinson, Peter Buneman, and Ronald Morrison, editors, *Data Types and Persistence. Edited Papers from the Proceedings of the First Workshop on Persistent Objects, Appin, Scotland, August 1985*, Topics in Information Systems, pages 31–41. Springer, 1988. Since revised.
- [DCH03] Derek Dreyer, Karl Crary, and Robert Harper. A type system for higher-order modules. In *POPL '03: Proceedings of the 30th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 236–249, New York, NY, USA, 2003. ACM Press.
- [Dre02] Derek Dreyer. Moscow ML's higher-order modules are unsound, December 2002. Message on the TYPES forum. Online at <http://www.seas.upenn.edu/~sweirich/types/archive/1999-2003/msg01136.html>.
- [Dre05] Derek Dreyer. *Understanding and Evolving the ML Module System*. PhD thesis, Carnegie Mellon University, 2005.
- [FLFS07] Cédric Fournet, Fabrice Le Fessant, and Alan Schmitt. *The JoCaml language (beta release)*, 2007.
- [FLMR97] Cédric Fournet, Cosimo Laneve, Luc Maranget, and Didier Rémy. Implicit typing à la ml for the join-calculus. In *CONCUR '97: Proceedings of the 1997 8th International Conference on Concurrency Theory*, pages 196–212. Springer-Verlag, July 1997.
- [Gar04] Jacques Garrigue. Relaxing the value restriction. In *FLOPS '04: Proceedings of the 7th International Symposium on Functional and Logic Programming*, volume 2998 of *Lecture Notes in Computer Science*, pages 196–213. Springer-Verlag, April 2004.
- [Gog05] Healdene Goguen. A syntactic approach to eta equality in type theory. In *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 75–84, 2005.
- [HL94] Robert Harper and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *POPL '94: Proceedings of the 21st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 123–137, New York, NY, USA, 1994. ACM Press.

- [HMM90] Robert Harper, John C. Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *POPL '90: Proceedings of the 17th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 341–354, New York, NY, USA, 1990. ACM Press.
- [Klo80] Jan Willem Klop. *Combinatory reduction systems*. PhD thesis, Mathematisch Centrum, Amsterdam, 1980.
- [L⁺] Xavier Leroy et al. *The Objective Caml system*.
- [Ler] Xavier Leroy. Private communication.
- [Ler94] Xavier Leroy. Manifest types, modules, and separate compilation. In *POPL '94: Proceedings of the 21st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 109–122, New York, NY, USA, 1994. ACM Press.
- [Ler95] Xavier Leroy. Applicative functors and fully transparent higher-order modules. In *POPL '95: Proceedings of the 22nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 142–153, New York, NY, USA, 1995. ACM Press.
- [Lil97] Mark Lillibridge. *Translucent Sums: A Foundation for Higher-Order Module Systems*. PhD thesis, Carnegie Mellon University, May 1997.
- [LPSW03] James J. Leifer, Gilles Peskine, Peter Sewell, and Keith Wansbrough. Global abstraction-safe marshalling with hash types. In *ICFP '03: Proceedings of the eighth ACM SIGPLAN international conference on Functional programming*, pages 87–98, New York, NY, USA, 2003. ACM Press.
- [Mac84] David MacQueen. Modules for Standard ML. In *LFP '84: Proceedings of the 1984 ACM Symposium on LISP and functional programming*, pages 198–207, New York, NY, USA, 1984. ACM Press.
- [MM01] Louis Mandel and Luc Maranget. *The JoCaml language*, 2001.
- [OTCP90] Atsushi Ohori, Ivan Tabkha, Richard Connor, and Paul Philbrow. Persistence and type abstraction revisited. In *Implementing Persistent Object Bases, Principles and Practice, Proceedings of the Fourth International Workshop on Persistent Objects, 23-27 September 1990, Martha's Vineyard, MA, USA*, pages 141–153. Morgan Kaufmann, 1990.
- [Pie05] Benjamin C. Pierce, editor. *Advanced Topics in Types and Programming Languages*. MIT Press, 2005.
- [PS00] Benjamin C. Pierce and Eijiro Sumii. Relating cryptography and polymorphism. Manuscript, July 2000.
- [PSL] Programming System Lab, Saarland University. *The Alice ML Language*.
- [Ros] Andreas Rossberg. SML vs. Ocaml. Online at <http://www.ps.uni-sb.de/~rossberg/SMLvsOcaml.html>.

-
- [Ros03] Andreas Rossberg. Generativity and dynamic opacity for abstract types. In *PPDP '03: Proceedings of the 5th ACM SIGPLAN international conference on Principles and practice of declarative programming*, pages 241–252, New York, NY, USA, 2003. ACM Press.
- [Ros07] Andreas Rossberg. *Typed Open Programming — A higher-order, typed approach to dynamic modularity and distribution*. PhD thesis, Universität des Saarlandes, January 2007.
- [RRS] Sergei Romanenko, Claudio Russo, and Peter Sestoft. *Moscow ML Language Overview*.
- [Rus98] Claudio Russo. *Types For Modules*. PhD thesis, University of Edinburgh, 1998.
- [Sew01] Peter Sewell. Modules, abstract types, and distributed versioning. In *POPL '01: Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 236–247, New York, NY, USA, 2001. ACM Press.
- [Sha99] Zhong Shao. Transparent modules with fully syntactic signatures. In *ICFP '99: Proceedings of the fourth ACM SIGPLAN international conference on Functional programming*, pages 220–232, New York, NY, USA, 1999. ACM Press.
- [SP04] Eijiro Sumii and Benjamin C. Pierce. A bisimulation for dynamic sealing. In *POPL '04: Proceedings of the 31st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 161–172, New York, NY, USA, 2004. ACM Press.
- [Sun] Sun Microsystems, Inc. *Java API Specifications*.
- [Wad89] Philip Wadler. Theorems for free! In *FPCA '89: Proceedings of the fourth international conference on Functional programming languages and computer architecture*, pages 347–359, New York, NY, USA, 1989. ACM Press.
- [Wri95] Andrew K. Wright. Simple imperative polymorphism. *Lisp Symb. Comput.*, 8(4):343–355, 1995.
- [ZGM99] Steve Zdancewic, Dan Grossman, and Greg Morrisett. Principals in programming languages: a syntactic proof technique. In *ICFP '99: Proceedings of the fourth ACM SIGPLAN international conference on Functional programming*, pages 197–207, New York, NY, USA, 1999. ACM Press.

Index

- abstract
 - type, 8
- abstract component, 53
- abstract type, *see* type
- abstraction kind, 65
- alpha-conversion, **13**, 46
- ambient colour, **42**
- apax, *see* nonce
- applicative
 - functor, *see* functor
- ascription, **34**
- avoidance problem, 15

- closed
 - term, **13**
- colour
 - primary —, **42**, **51**
 - transparency, 57
 - variables in —, 44
- coloured bracket, **39**, 41, 51
 - absolute —, **46**
 - additive, **46**
 - pushing, 44, 48
 - universal —, 46
- coloured brackets
 - pushing, 42, 54
- comparable (module), **25**
- completely specified, *see* monomorphic
- component type, **40**
- component value, **53**
- concretisation, **44**, 47
- conversion, 20
 - type —, **22**
- convertibility, **22**

- decidability, 64
- dependency
 - of a module identity, **43**
- distributed
 - system, 7

- domain, **13**
- dynamic
 - sealing, *see* sealing

- effect, 26
- empty colour, **42**
- environment, **13**
- equitypable, **32**
- equivalence
 - convertibility, *see* convertibility
- evaluation context, **16**
- extensionality, 24

- free variable, **13**
- fully specified, *see* monomorphic
- function
 - polymorphic —, *see* polymorphic
- functor
 - applicative, **30**
 - generative, **30**
 - transparent, **30**

- generative
 - functor, *see* functor
- generics, 50

- hash, 63

- impure, 26
 - sealing, *see* sealing
- incompletely specified, *see* polymorphic
- inseparable
 - sealing, *see* sealing

- judgement
 - local typing —, **14**

- kind, **48**, 51
 - abstraction —, *see* abstraction kind

- lexis, **38**

- local binding, 29
- marshal, 61
- marshaling, 7
- minimal
 - sealing, *see* sealing
- module, 13
- module component, 40, 51
- module identity, 37, *see* nonce
- monomorphic, 44, 48

- name server, 67
- nonce, 37, 63
 - transparency, 57

- opacity, 66
- opaque, 42

- parametricity, 21, 66
- partially specified, *see* polymorphic
- phase separation, 25
- pickling, *see* marshaling
- polymorphic, 48
 - function, 49
 - type parameter, 50
 - value, 50
- projectible, 25
- pure, 26
 - module, 25

- quasi-values, 53

- reduction, 16
- reveal, 40
- revelation, 52

- sandbox, 46
- sealing, 27
 - dynamic —, 31, 63, 65
 - dynamic —dynamic, 33
 - dynamic —static, 33
 - dynamic —strong, 33
 - dynamic —weak, 33
 - impure —, 33
 - inseparable —, 33
 - minimal —, 34
 - separable —, 33
 - static —, 31, 63
 - strong —, 30
 - weak —, 30

- selfification, 21, 38, 48
- separable, 25, 67
 - sealing, *see* sealing
- serialisation, *see* marshaling
- signature, 13
- singleton, 30
 - higher-order, 18
 - of an expression, 17
- singularised identity, 37, 63
- sound, 64
- stamp, 37, 63
- stamp book, *see* lexis
- static
 - sealing, *see* sealing
- strengthening, 38
- strong
 - sealing, *see* sealing
- substitution, 13, 46

- transparent, 39, 42, 52, 57
 - functor, *see* functor
 - variable, 45
- type
 - abstract—, 51
 - arrow —, 13
 - coercion, 38
 - function, 13
 - inference, 64
 - product —, 13

- underlying nonce, 40, 52
- universal, 62
 - term, 44
 - coloured bracket, *see* coloured bracket

- value, 15
 - polymorphic —, *see* polymorphic
- value restriction, 67
- variable
 - transparency, 57

- weak
 - sealing, *see* sealing
- weakening
 - colour, 43, 45, 46