

Acute
**High-level programming language design
for distributed computation**
Design rationale and language definition

12th October 2004

Peter Sewell* James J. Leifer[†] Keith Wansbrough* Mair Allen-Williams*
Francesco Zappa Nardelli[†] Pierre Habouzit[†] Viktor Vafeiadis*

*University of Cambridge [†]INRIA Rocquencourt

<http://www.cl.cam.ac.uk/users/pes20/acute>

Abstract: This paper studies key issues for distributed programming in high-level languages. We discuss the design space and describe an experimental language, Acute, which we have defined and implemented.

Acute extends an OCaml core to support distributed development, deployment, and execution, allowing type-safe interaction between separately-built programs. It is expressive enough to enable a wide variety of distributed infrastructure layers to be written as simple library code above the byte-string network and persistent store APIs, disentangling the language runtime from communication.

This requires a synthesis of novel and existing features: (1) type-safe marshalling of values between programs; (2) dynamic loading and controlled rebinding to local resources; (3) modules and abstract types with abstraction boundaries that are respected by interaction; (4) global names, generated either freshly or based on module hashes: at the type level, as runtime names for abstract types; and at the term level, as channel names and other interaction handles; (5) versions and version constraints, integrated with type identity; (6) local concurrency and thread thunkification; and (7) second-order polymorphism with a namespace construct. We deal with the interplay among these features and the core, and develop a semantic definition that tracks abstraction boundaries, global names, and hashes throughout compilation and execution, but which still admits an efficient implementation strategy.

Key-words: programming languages, distributed programming, marshalling, serialisation, abstract types, modules, rebinding, version control, type theory, ML

Contents

List of typing judgements	7
List of figures	8
1 Introduction	9
I Design Rationale	13
2 Distributed abstractions: language vs libraries	13
3 Basic type-safe distributed interaction	13
4 Dynamic linking and rebinding to local resources	14
4.1 References to local resources	14
4.2 What to ship and what to rebind	15
4.3 Evaluation strategy: the relative timing of variable instantiation and marshalling	16
4.4 The structure of marks and modules	16
4.5 Controlling when rebinding happens	17
4.6 Controlling what to rebind to	18
4.7 The relationship between modules and the filesystem	18
4.8 Module initialisation	19
4.9 Marshalling references	19
5 Naming: global module and type names	20
6 Naming: expression names	22
6.1 Establishing shared names	22
6.2 A refinement: ties	24
6.3 Polymorphic name operations	25
6.4 Implementing names	25
7 Versions and version constraints	25
8 Interplay between abstract types, rebinding and versions	27
8.1 Definite and indefinite references	27
8.2 Breaking abstractions	29
8.3 Overriding valuability checks	31
8.4 Exact matching or version flexibility?	31
8.5 Marshalling inside abstraction boundaries	32
9 Concurrency, mobility, and thunkify	32
9.1 Language-level concurrency vs OS threads	32
9.2 Interaction primitives	33
9.3 Thunkification	33
9.4 Naming and grouping	33
9.5 Thread termination	34
9.6 Nonexistent threads, mutexes, and condition variables	34
9.7 References, names, marshalling, and thunkify	35
9.8 Module initialisation, concurrency, and thunkify	35
9.9 Thunkify and blocking calls	35
9.10 Concurrency: the constructs	36
9.11 Example	37
10 Polymorphism	37
10.1 A refinement: marshal keys and name equality	38
11 Pulling it all together: examples	38

12 Related work	39
13 Conclusions and future work	40
II Semantics	43
14 Semantics overview	43
14.1 Naming	43
14.2 Typing	44
14.3 Compilation	45
14.4 Operational judgements	46
14.5 Colours and bracket dynamics	47
14.6 Marshalling and unmarshalling	48
14.7 Module field instantiation	49
14.8 Concurrency	50
15 Semantics Examples	52
15.1 Compilation: hash modules	52
15.2 Compilation: fresh modules	53
15.3 Compilation: hash module dependencies	53
15.4 Compilation: cfresh modules	54
15.5 Compilation: constructing expression names from module hashes	55
15.6 Compilation: type normalisation and marshalling within abstraction boundaries	57
15.7 Compilation: imports	59
15.8 Compilation: imports with abstract type fields	61
15.9 Compilation: breaking abstractions	62
15.10 Marshalled values	63
III Definition	71
16 Language Definition	71
16.1 Metavariables	71
16.2 Syntax	73
16.2.1 Binding	89
16.3 Typing	91
16.3.1 Typing for Source Internal and Compiled Forms	91
16.3.2 Typing for Compiled and Executing Forms	104
16.4 Typed Desugaring	109
16.5 Errors	111
16.6 Valuability helper functions	113
16.7 Compilation	115
16.8 Operational semantics	125
16.8.1 The judgements	125
16.8.2 Values	125
16.8.3 Reduction contexts and closure rules	126
16.8.4 Simple expression forms	128
16.8.5 Marshalling and unmarshalling	133
16.8.6 Module field instantiation	135
16.8.7 Name operations	138
16.8.8 Concurrency	140
16.8.9 Polymorphism	144
16.9 Type Preservation and Progress	146
16.10 Runtime type checking	146
16.11 Vacuous bracket optimization	147
16.12 Closures	148
16.12.1 Value closures	148

16.12.2 Type closures	150
IV Communication Infrastructure Example	151
V Implementation	170
17 Overview	170
18 Command line options	171
19 Concrete user source grammar	173
20 Concrete compiled-form grammar	178
21 Library interfaces	183
22 The IO module	187
Appendix: Acute syntax summary	189
References	191
Index	193

Typing judgements

$E_n \vdash \text{ok}$	global name environment E_n is well-formed	92
$E_n, E \vdash \text{ok}$	type environment E is well-formed	92
$E_n \vdash \text{nmodule}_{eqs} M : Sig_0 \text{ version } vne = Str \text{ ok}$	nmodule global name data ok	92
$E_n \vdash \text{nimport}_{eqs} M : Sig_0 \text{ version } vc \text{ like } Str \text{ ok}$	nimport global name data ok	92
$E_n \vdash h \text{ ok}$	hash h is well-formed	92
$E \vdash K \text{ ok}$	kind K is well-formed	93
$E \vdash_{eqs} K \approx K'$	kinds K and K' are equal	93
$E \vdash_{eqs} K <: K'$	kind K is a subkind of K'	93
$E \vdash_{eqs} T : K$	type T has kind K	94
$E \vdash_{eqs} T \approx T'$	types T and T' are equal	94
$E \vdash_{eqs} \text{ok}$	equation-set colour eqs is well-formed	95
$E \vdash sig \text{ ok}$	signature sig is well-formed	96
$E \vdash_{eqs} sig <: sig'$	signature sig is a subsignature of sig'	96
$E \vdash_{eqs} sig \approx sig'$	signature Sig is equal to Sig'	96
$E \vdash_{eqs} str : sig$	structure body str has signature sig	96
$E \vdash Sig \text{ ok}$	signature Sig is well-formed	97
$E \vdash_{eqs} Sig <: Sig'$	signature Sig is a subsignature of Sig'	97
$E \vdash_{eqs} Sig \approx Sig'$	under assumptions E , signature Sig is equal to Sig'	97
$E \vdash_{eqs} Str : Sig$	structure Str has signature Sig	97
$\vdash str \text{ flat}$	str as flat as possible	97
$\vdash Str \text{ flat}$	Str as flat as possible	97
$\vdash sig \text{ flat}$	sig as flat as possible	97
$\vdash Sig \text{ flat}$	Sig as flat as possible	97
$E \vdash_{eqs} M_M : Sig$	module M_M has signature Sig	98
$E \vdash_{eqs} h : Sig$	name h has signature Sig	98
$E \vdash_{eqs} e : T$	expression e has type T	99
$E \vdash_{eqs} e : T$	continued... expression e has type T	100
$E \vdash_{eqs} e : T$	Sugared source forms expression e has type T	101
$E \vdash p : T \triangleright E'$	pattern p matches type T , giving additional bindings E'	101
$E \vdash_{eqs} mtch : T \rightarrow T'$	match $mtch$ has type $T \rightarrow T'$	101
$E_n \vdash avne \text{ ok}$	atomic version number expression ok	102
$E_n \vdash vne \text{ ok}$	version number expression ok	102
$E \vdash ahvce \text{ ok}$	atomic name version constraint expression ok	102
$E \vdash avce \text{ ok}$	atomic version constraint expression ok	102
$E \vdash dvce \text{ ok}$	dotted version constraint expression ok	102
$E \vdash vce \text{ ok}$	version constraint expression ok	102
$E \vdash likespec \text{ ok}$	likespec ok	103
$E \vdash = Mo : Sig$	link is UNLINKED or has signature Sig	103
$E \vdash sourcedefinition \triangleright E'$	source definition $sourcedefinition$ gives additional module bindings E'	103
$E \vdash definition \triangleright E'$	compiled definition $definition$ gives additional module bindings E'	105
$E ; E_s \vdash s \text{ store}$	store s is well-formed	107
$E \vdash definitions \triangleright E'$	$definitions$ give module bindings E'	107
$E \vdash definitions \text{ eo ok}$	$definitions \text{ eo}$ is well-typed	107
$\vdash E_n ; \langle E_s, s, definitions, e \rangle : T$	pseudo-configuration is well-formed with result type T	107
$\vdash E_n ; \langle E_s, s, definitions, P \rangle : T$	configuration is well-formed with result type T	107
$\vdash mv \text{ ok}$	marshalled value ok	108
$E \vdash P \text{ ok}$	process P is ok	108

List of Figures

1	Typing Rules – Type Environments, Hashes	92
2	Typing Rules – Kinds	93
3	Typing Rules – Auxiliaries	93
4	Typing Rules – Types, Type Equality	94
5	Typing Rules – Equation sets	95
6	Typing Rules – Signatures, Subsignaturing (part 1)	96
7	Typing Rules – Signatures, Subsignaturing (part 2)	97
8	Typing Rules – flat predicates	97
9	Typing Rules – Signatures of module identifiers and hashes	98
10	Typing Rules – Expressions (part 1)	99
11	Typing Rules – Expressions (part 2)	100
12	Typing Rules – Sugared Forms	101
13	Typing Rules – Patterns, Matches	101
14	Typing Rules – Version number and constraint expressions	102
15	Typing Rules – Definition auxiliaries	103
16	Typing Rules – Source Definitions	103
17	Typing Rules – Compiled Definitions	105
18	Link Checking	106
19	Typing Rules – Store, Configurations	107
20	Typing Rules – Marshalled Values	108
21	Typing Rules – Processes	108

1 Introduction

Distributed computation is now pervasive, with execution, software development, and deployment spread over large networks, long timescales, and multiple administrative domains. Because of this, many distributed systems cannot be deployed or updated atomically. They are not composed of multiple instances of a single program version, but instead of many versions of many programs that need to interoperate, perhaps sharing some libraries but not others. Moreover, the intrinsic concurrency and nondeterminism of distributed systems, and the complexity of the underlying network layers, makes them particularly hard to understand and debug.

Existing programming languages, such as ML, Haskell, Java and C#, provide good support for local computation, with rich type structures and (mostly) static guarantees of type safety. When it comes to distributed computation, however, they fall short, with little support for its many system-development challenges.

This paper addresses the design of distributed languages. Our focus is on the higher-order, typed, call-by-value programming of the ML tradition: we concentrate on what must be added to ML-like languages to support typed distributed programming. We discuss the design space and describe a programming language, Acute, which we have defined and implemented.

Acute extends an OCaml core with a synthesis of several novel and existing features, broadly addressing naming and identity in the distributed setting. It is not a proposal for a full-scale language, but rather a testbed for experimentation. Our extensions are lightweight changes to ML, but suffice to enable sophisticated distributed infrastructure, e.g. substantial parts of JoCaml [JoC] or Nomadic Pict [SWP99], to be programmed as simple libraries (and its support for interaction between programs goes well beyond these). We demonstrate this with an example typed communication library, written in Acute above the byte-string TCP Sockets API, which requires and uses most of the new features.

The paper is divided into four parts. Part I is devoted to an informal presentation of the main design points from the programmer’s point of view, omitting details of the semantics. It is supported by a full definition of Acute in Part III, with the main points of the semantics explained in Part II (including the compiled code and marshalled values from the Part I examples), and by an implementation. The definition covers syntax, typing, compilation, and operational semantics. The implementation is a prototype, efficient enough to run moderate examples while remaining close to the semantics. Part IV gives the Acute code for the communication infrastructure example of §11. Part V gives a brief description of the implementation together with the current command-line options, concrete syntax and standard libraries. The definition and implementation have both been essential: the synthesis of the various features has involved many semantic subtleties. The definition is too large (on the scale of the ML definition rather than an idealised λ -calculus) to make proofs of the properties feasible with the available resources and tools. To increase confidence in both semantics and implementation, therefore, our implementation can optionally type-check the entire configuration after each reduction step.

Design rationale Part I is structured as follows, with §2–10 discussing the main design points, §11 demonstrating that Acute does indeed support typeful distributed programs with an example distributed communication infrastructure library, and §12 and §13 describing related and future work and concluding. An appendix summarises most of the Acute syntax.

§2 and §3 set the scene: we discuss the right level of abstraction for a general-purpose distributed language, arguing that it should not have a commitment to any particular form of communication. We then recall the design choices for simple type-safe marshalling, for trusted and untrusted interaction.

§4: We introduce *dynamic linking* and *rebinding* to local resources in the setting of a language with an ML-like second-class module system. There are many questions here: of how to specify which resources should be shipped with a marshalled value and which dynamically rebound; what evaluation strategy to use; when rebinding takes effect; and what to rebind to. In this Section our aim is to expose the design choices rather than identify definitive solutions. It is a necessary preliminary to our work in §§5–11. For Acute we make interim choices, reasonably simple and sufficient to bring out the typing and versioning issues involved in rebinding, which here is at the granularity of module identifiers. A running Acute program consists roughly of a sequence of `module` definitions (of ML structures), `imports` of modules with specified signatures, which may or may not be linked, and marks which indicate where rebinding can take effect; together with running processes and a shared store.

§5: Type-safe marshalling demands a notion of *type identity* that makes sense across multiple versions of differing programs. For concrete types this is conceptually straightforward, but with abstract types more care is necessary. We generate globally-meaningful *type names* either by *hashing* module definitions, taking their dependencies into

account; *freshly at compile-time*; or *freshly at run-time*. The first two enable different builds or different programs to share abstract type names, by sharing their module source code or object code respectively; the last is needed to protect the invariants of modules with effect-full initialisation.

§6: Globally-meaningful *expression-level names* are needed for type-safe interaction, e.g. for communication channel names or RPC handles. They can also be constructed as hashes or created fresh at compile time or run time; we show how these support several important idioms. The polytypic support and swap operations of Shinwell, Pitts and Gabbay’s FreshOCaml [SPG03] are included to support swizzling of local names during communication.

§7: In a single-program development process one ensures the executable is built from a coherent set of versions of its modules by controlling static linking — often by building from a single source tree. With dynamic linking and rebinding more support is required: we add *versions* and *version constraints* to modules and imports respectively. Allowing these to refer to module names gives flexibility over whether code consumers or producers have control.

§8: There is subtle interplay between versions, modules, imports, and type identity, requiring additional structure in modules and imports. A mechanism for looking through abstraction boundaries is also needed for some version-change scenarios.

§9: Local concurrency is important for distributed programming. Acute provides a minimal level of support, with threads, mutexes and condition variables. Local messaging libraries can be coded up using these, though in a production implementation they might be built-in for performance. We also provide *thunkification*, allowing a collection of threads (and mutexes and condition variables) to be captured as a thunk that can then be marshalled and communicated (or stored); this enables various constructs for mobility to be coded up.

Part I is an extended version of [SLW⁺]. The main changes are:

- addition of §4.7 on the relationship between modules and the filesystem;
- addition of §4.8 on module initialisation;
- addition of §4.9 on marshalling references;
- addition of §6.2–§6.4 on naming: name ties, polytypic name operations, and the implementation of names;
- extension of §7 on versioning;
- extension of §8.2 on breaking abstractions and `with!`;
- addition of §8.5 on marshalling inside abstraction boundaries;
- extension of §9 on concurrency, with §9.1–9.11 covering the choices for threads and `thunkify` in more detail, discussing several interactions between language features; and
- addition of §10 on polymorphism and namecase.

Semantics and Implementation The definition of compilation describes how global type- and expression-level names are constructed. Unusually, the semantics preserves the module structure throughout computation, instead of substituting it away; this is needed to express rebinding. Abstraction boundaries are also preserved, with a generalisation of the *coloured brackets* of Grossman et al [GMZ00] to the entire Acute language (except, to date, the System F constructs). This is technically delicate (and not needed for implementations, which can erase all brackets) but provides useful clarity in a setting where abstraction boundaries may be complex, with abstract types shared between programs.

The semantics preserves also the internal structure of hashes and type data associated with freshly-created names. This too can be erased in implementations, which can implement hashes and fresh names with literal bit-strings (e.g. 160-bit SHA1 hashes and pseudo-random numbers), but is needed to state type preservation and progress properties. The abstraction-preserving semantics makes these rather stronger than usual.

The Acute implementation is written in FreshOCaml, as a meta-experiment in using the Fresh features for a medium-scale program (some 25 000 lines). It is a prototype: designed to be efficient enough to run moderate examples while remaining rather close to the semantics. The runtime interprets an intermediate language which is essentially the abstract syntax extended with closures.

Syntax For concreteness we summarise the most interesting constructs of Acute for types, expressions, and definitions. The full grammar is given in the Definition and summarised in an appendix. The **highlighted** forms do not

occur in source programs. Here h is a module name, hash- or freshly-generated; n is a freshly-generated name, and $[e]_{eqs}^T$ is a coloured bracket. The other constructs are explained later.

```

T ::= ... | T name | thread | h.t | n

e ::= ... | marshal e1 e2 : T | unmarshal e as T |
freshT | cfreshT | hash(MM.x)T | hash(T, e2)T' | hash(T, e2, e1)T' |
swap e1 and e2 in e3 | supportT e | thunkify | [e]eqsT
```

```

sourcedefinition ::=
module mode MM : Sig version vne = Str withspec |
import mode MM : Sig version vce likespec by resolvespec = Mo |
mark MK
```

These are added to a substantial fragment of ML. The core language of Acute consists of normal ML types and expressions: booleans, integers, strings, tuples, lists, options, recursive functions, pattern matching, references, exceptions, and invocations of OS primitives in standard libraries. It does not have standard ML-style polymorphism, as our distributed infrastructure examples need first-class existentials (e.g. to code up polymorphic channels) and first-class universals (for marshalling polymorphic functions). We therefore have explicit System F style polymorphism, and for the time being the implementation does some ad-hoc partial inference. The full grammar of types is

```

T ::= int | bool | string | unit | char | void | T1 * .. * Tn | T1 + .. + Tn | T → T' | T list | T option | T ref | exn |
MM.t | t | ∀ t. T | ∃ t. T | T name | T tie | thread | mutex | cvar | thunkifymode | thunkkey | thunklet | h.t | n
```

The module language includes top-level declarations of structures, containing expression fields and type fields, with both abstract and manifest types in signatures. Module initialisation can involve arbitrary computation.

We omit some other standard features, simply to keep the language small: user-defined type operators, constructors, and exceptions; substructures; and functors (we believe that adding first-order applicative functors would be straightforward; going beyond that would be more interesting). Some more substantial extensions are discussed in the Conclusion. To avoid syntax debate we fix on one in use, that of OCaml.

Contribution Our contribution is threefold: discussion of the design space and identification of features needed for high-level typed distributed programming, the synthesis of those features into a usable experimental language, and their detailed semantic design. We build on our previous work on global type names and dynamic rebinding [Sew01, LPSW03, BHS⁺03] which developed some of these ideas for small calculi. The main technical innovations here are: a uniform treatment of names created by hash, fresh, or compile-time fresh, both for type names and (covering the main usage scenarios) for expression names, dealing with module initialisation and dependent-record modules; explicit versions and version constraints, with their delicate interplay with imports and type equality; module-level dynamic linking and rebinding; support for thunkification; and an abstraction-preserving semantics for all the above.

Part I

Design Rationale

2 Distributed abstractions: language vs libraries

A fundamental question for a distributed language is what communication should be built in to the language runtime and what should be left to libraries. The runtime must be widely deployed and so is not easily changed, whereas additional libraries can easily be added locally. In contrast to some previous languages (e.g. Facile [TLK96], Obliq [Car95], and JoCaml [JoC]), we believe that *a general-purpose distributed programming language should not have a built-in commitment to any particular means of interaction.*

The reason for this is essentially the complexity of the distributed environment: system designers must deal with partial failure, attack, and mobility — of code, of devices, and of running computations. This complexity demands a great variety of communication and persistent store abstractions, with varying performance, security, and robustness properties. At one extreme there are systems with tightly-coupled computation over a reliable network in a single trust domain. Here it might be appropriate to use a distributed shared memory abstraction, implemented above TCP. At another extreme, interaction may be intrinsically asynchronous between mutually-untrusting runtimes, e.g. with cryptographic certificates communicated via portable persistent storage devices (smartcards or memory sticks), between machines that have no network connection. In between, there are systems that require asynchronous messaging or RMI but, depending on the network firewall structure, tunnel this over a variety of network protocols.

To attempt to build in direct support for all the required abstractions, in a single general-purpose language, would be a never-ending task. Rather, the language should be at a level of abstraction that makes distribution and communication explicit, allowing distributed abstractions to be expressed as libraries.

Acute has constructs `marshal` and `unmarshal` to convert arbitrary values to and from byte strings; they can be used above any byte-oriented persistent storage or communication APIs.

This leaves the questions of (a) how these should behave, especially for values of functional or abstract types, and (b) what other local expressiveness is required, especially in the type system, to make it possible to code the many required libraries. The rest of the paper is devoted to these.

3 Basic type-safe distributed interaction

In this section we establish our basic conventions and assumptions, beginning with the simplest possible examples of type-safe marshalling. We first consider one program that sends the result of marshalling 5 on a fixed channel:

```
I0.send( marshal "StdLib" 5 : int )
```

(ignore the "StdLib" for now) and another that receives it, adds 3 and prints the result:

```
I0.print_int(3+(unmarshal(I0.receive()) as int))
```

Compiling the two programs and then executing them in parallel results in the second printing 8. This and subsequent examples are executable Acute code. For brevity they use a simple address-less `I0` library, providing communication primitives `send:string->unit` and `receive:unit->string`. (There are two implementations of `I0`, one uses TCP via the Acute sockets API, with the loopback interface and a fixed port; the other writes and reads strings from a file with a fixed name.) Below we write the parallel execution of the two separately-built programs `p1` and `p2` separated by a `—`.

For safety, a type check is obviously needed at run-time in the second program, to ensure that the type of the marshalled value is compatible with the type at which it will be used. For example, the second program here

```
I0.send( marshal "StdLib" "five" : string )
```

—

```
IO.print_int(3+(unmarshal(IO.receive()) as int))
```

should raise an exception. Allowing interaction via an untyped medium inevitably means that some dynamic errors are possible, but they should be restricted to clearly-identifiable program points, and detected as early as possible. This error can be detected at unmarshal-time, rather than when the received value is used as an argument to `+`, so we should do that type check at unmarshal-time. (In some scenarios one may be able to exclude such errors at compile-time, e.g. when communicating on a typed channel; we return to this in §6.)

The unmarshal dynamic check might be of two strengths. We can:

- (a) include with the marshalled value an explicit representation of the type at which it was marshalled, and check at unmarshal-time that that type is equal to the type expected by the unmarshal — in the examples above, `int=int` and `string=int` respectively; or
- (b) additionally check that the marshalled value is a well-formed representation of something of that type.

In a trusted setting, where one can assume that the string was created by marshalling in a well-behaved runtime (which might be assured by network locality or by cryptographically-protected interaction with trusted partners), option (a) suffices for safety.

If, however, the string might have been created or modified by an attacker, then we should choose (b), to protect the integrity of the local runtime. This option is not always available, however: when we consider marshalled values of an abstract type, it may not be possible to check at unmarshal-time that the intended invariants of the type are satisfied. They may have never been expressed explicitly, or be truly global properties. In this case one should marshal only values of concrete types.¹

A full language should provide both, but in Acute we focus on the trusted case, with option (a), and the problems of distributed typing, naming, and rebinding it raises. Techniques for the untrusted case, including XML support and proof-carrying code, are also necessary but are largely orthogonal.

We do not discuss the design of the concrete wire format for marshalled values — the Acute semantics presupposes just a partial `raw_unmarshal` function from strings to abstract syntax of configurations, including definitions and store fragments; the prototype implementation simply uses canonical pretty-prints of abstract syntax. A production language would need an efficient and standardised internal wire format, and for some purposes (and for simple types) a canonical XML representation would be useful for interoperation. In the untrusted case XML is now widely used and good language support for (b) is clearly important.

4 Dynamic linking and rebinding to local resources

4.1 References to local resources

Marshalling closed values, such as the `5` and `"five"` above, is conceptually straightforward. The design space becomes more interesting when we consider marshalling a value that refers to some local resources. For example, the source code of a function (it may be useful to think of a large plug-in software component) might mention identifiers for:

- (1) ubiquitous standard library calls, e.g., `print_int`;
- (2) application-specific library calls with location-dependent semantics, e.g., routing functions;
- (3) application code that is not location-dependent but is known to be present at all relevant sites; and
- (4) other let-bound application values.

In (1–3) the function should be *rebound* to the local resource where and when it is unmarshalled, whereas in (4) the definitions of resources must be copied and sent along before their usages can be evaluated.

There is another possibility: a local resource could be converted into a *distributed reference* when the function is marshalled, and usages of it indirected via further network communication. In some scenarios this may be desirable, but in others it is not, where one cannot pay the performance cost for those future invocations, or cannot depend

¹One could imagine an intermediate point, checking the representation type but ignoring the invariants, but the possibility of breaking key invariants is in general as serious as the possibility of breaking the local runtime.

on future reliable communication (and do not want to make each invocation of the resource separately subject to communication failures). Most sharply, where the function is marshalled to persistent store, and unmarshalled after the original process has terminated, distributed references are nonsensical. Following the design rationale of §2, we do not support distributed references directly, aiming rather to ensure our language is expressive enough to allow libraries of ‘remotable’ resources to be written above our lower-level marshalling primitives.

4.2 What to ship and what to rebound

Which definitions fall into (2) (to be rebound) and (4) (to be shipped) must be specified by the programmer; there is usually no way for an implementation to infer the correct behaviour. How this should be expressed in the language is explored below.

On the other hand, tracking which definitions need not be shipped (3) because they are present at the receiver can be amenable to automation in some scenarios: in the case where we have good connectivity, and are communicating one-to-one rather than via multicast, the two parties can exchange fingerprints of what is required/present. If there is a repeated interchange of messages, the parties may even cache this data from one to another. We believe a good language should make it possible to encode such algorithms, but again, the variety of choices of desirable distributed behaviour leads us to believe that none should be built in. Encoding them requires some reflectivity — to inspect the set of resources required by a value, and calculate the subset of those that are not already present at the receiver. In this paper we do not go into this further, and such *negotiation* protocols are not expressible in Acute at present.

Instead, we adapt the mechanism proposed in [BHS⁺03] (from a lambda-calculus setting to an ML-style module language) to indicate which resources should be rebound and which shipped for any marshal operation. An Acute program consists roughly of a sequence of module definitions, interspersed with *marks*, followed by running processes; those module definitions, together with implicit module definitions for standard libraries, are the resources. Marks essentially name the sequence of module definitions preceding them. Marshal operations are each with respect to a mark; the modules below that mark are shipped and references to modules above that mark are rebound, to whatever local definitions may be present at the receiver. The mark "StdLib" used in §3 is declared at the end of the standard library; both this mark and library are in scope in all examples.

In the following example the sender declares a module M with a y field of type int and value 6. It then marshals and sends the value fun ()->M.y. This marshal is with respect to mark "StdLib", which lies above the definition of module M, so a copy of the M definition is marshalled up with the value fun ()->M.y. Hence, when this function is applied to () in the receiver the evaluation of M.y can use that copy, resulting in 6.

```
module M : sig val y:int end = struct let y=6 end
IO.send( marshal "StdLib" (fun ()->M.y))
—
(unmarshal (IO.receive ()) as unit -> int) ()
```

On the other hand, references to modules above the specified mark can be rebound. In the simplest case, one can rebound to an identical copy of a module that is already present on the receiver (for (3) or (1)). In the example below, the M1.y reference to M1 is rebound, whereas the first definition of M2 is copied and sent with the marshalled value. This results in () and ((6,3),4) for the two programs.

```
module M1:sig val y:int end = struct let y=6 end
mark "MK"
module M2:sig val z:int end = struct let z=3 end

IO.send( marshal "MK" (fun ()-> (M1.y,M2.z))
          : unit->int*int)
—
module M1:sig val y:int end = struct let y=6 end
module M2:sig val z:int end = struct let z=4 end
((unmarshal(IO.receive()) as unit->int*int)(),M2.z)
```

Note that we must permit running programs to contain multiple modules with the same source-code name and interface but with different definitions — here, after the unmarshal, the receiver has two versions of M2 present, one used by the unmarshalled code and the other by the original receiver code.

In more interesting examples one may want to rebound to a local definition of M1 even if it is not identical, to pick up some truly location-dependent library. The code below shows this, terminating with () and (7,3).

```

module M1:sig val y:int end = struct let y=6 end
import M1:sig val y:int end version * = M1
mark "MK"
module M2:sig val z:int end = struct let z=3 end
IO.send( marshal "MK" (fun ()-> (M1.y,M2.z))
        : unit->int*int )
—
module M1:sig val y:int end = struct let y=7 end
module M2:sig val z:int end = struct let z=4 end
(unmarshal (IO.receive ()) as unit->int*int) ()

```

The sender has two modules, M1 and M2, with M1 above the mark MK. It marshals a value `fun ()-> (M1.y,M2.z)`, that refers to both of them, with respect to that mark. This treats M2.z statically and M1.y dynamically at the marshal/unmarshal point: a copy of M2 is sent along, and on unmarshalling at the receiver the value is rebound to the local definition of M1, in which y=7. To permit this rebounding we add an explicit *import*

```
import M1 : sig val y:int end version * = M1
```

An import introduces a module identifier (the left M1) with a signature; it may or may not be linked to an earlier module or import (this one is, to the earlier M1). The `version *` overrides the default behaviour, which would constrain rebounding only to identical copies of M1. Marks are simply string constants, not binders subject to alpha equivalence, as they need to be dynamically rebound. For example, if one marshals a function that has an embedded marshal with respect to "StdLib", and then unmarshals and executes it elsewhere, one typically wants the embedded marshal to act with respect to the now-local "StdLib".

4.3 Evaluation strategy: the relative timing of variable instantiation and marshalling

A language with rebounding cannot use a standard call-by-value operational semantics, which substitutes out identifier definitions as it comes to them, as some definitions may need to be rebound later. Two alternative CBV reduction strategies were developed in [BHS⁺03] in a simple lambda-calculus setting: *redex-time*, in which one instantiates an identifier with its value only when the identifier occurs in redex-position, and *destruct-time* where instantiation may occur even later. Here, to make the semantics as intuitive as possible, we use the redex-time strategy for module references (local expression reduction remains standard CBV).

For example, the first occurrence of M.y in the first program below will be instantiated by 6 before the marshal happens, whereas the second occurrence would not appear in redex-position until a subsequent unmarshal and application of the function to (); the second occurrence is thus subject to rebounding. The results are () and (6,2).

```

module M:sig val y:int end = struct let y=6 end
import M:sig val y:int end version * = M
mark "MK"
IO.send( marshal "MK" (M.y, fun ()-> M.y)
        : int * (unit->int) )
—
module M:sig val y:int end = struct let y=2 end
let ((x:int),(f:unit->int)) =
  (unmarshal(IO.receive()) as int*(unit->int)) in
(x, f ())

```

4.4 The structure of marks and modules

A running Acute program has a linear sequence of evaluated definitions (marks, module definitions and imports) scoping in the running processes. Imports may be linked only to module definitions (or imports) that precede them in

this sequence. When a value is unmarshalled that carried additional module definitions with it, those definitions are added to the end of the sequence.

This linear structure is not ideal. There are some obvious possible alternatives, whose exploration we leave for future work. An unordered set of module definitions would allow cyclic linking; or a tree structure would allow the usual structure of nested scopes to be expressed. In a sufficiently reflective language (i.e. one that would support negotiation, as mentioned above) one could think of coding up marks, dynamically maintaining particular sets of module names. One might well want explicit control over what must *not* be shipped, e.g. due to license restrictions or security concerns.

With any mark structure one has to decide where to put module definitions carried with values being unmarshalled. A useful criterion is to ensure that *repeated* marshalling/unmarshalling, moving code between many machines, behaves well. With the linear structure, putting definitions at the end of the sequence ensures they are inside all marks, and so will be picked up by subsequent marshals. In the hierarchical or unordered cases it is less clear what to do.

A further criterion is that the user of a module should not be required to know its dependency tree — in particular, if one specifies that the module be shipped, other modules that it may have dynamically loaded should be treated sensibly.

We also have to decide what to do with marks occurring between modules being marshalled: they can either be discarded or copied and sent. In Acute we take the latter semantics, but neither is fully satisfactory: in one, shipped module code may refer to marks that are not present locally; in the other there can be unwanted mark shadowing. This is a limitation of the linear structure.

4.5 Controlling when rebinding happens

We have to choose whether or not to allow execution of partial programs, which are those in which some imports are not linked to any earlier module definition (or import). Partial programs can arise in two ways. First, they can be written as such, as in conventional programs that use dynamic linking, where a library is omitted from the statically-linked code, to be discovered and loaded at runtime. For example:

```
import M : sig val y:int end version * = unlinked
fun () -> M.y
```

Secondly, they can be generated by marshalling, when one marshals a value that depends on a module above the mark (intending to rebind it on unmarshalling). For example, the final state of the receiver in

```
module M:sig val y:int end = struct let y=6 end
import M:sig val y:int end version * = M
mark "MK"
IO.send( marshal "MK" (fun ()->M.y) : unit->int )
—
unmarshal (IO.receive ()) as unit->int
```

is roughly the program below.

```
import M : sig val y:int end version * = unlinked
fun () -> M.y
```

If we disallow execution of partial programs then, when we unmarshal, all the unlinked imports that were sent with the marshalled value must be linked in to locally-available definitions; the unmarshal should fail if this is not possible.

Alternatively, if we allow execution of partial programs, we must be prepared to deal with an $M.x$ in redex position where M is declared by an unlinked import. For any particular unmarshal, one might wish to force linking to occur at unmarshal time (to make any errors show up as early as possible) or defer it until the imported modules are actually used. The latter allows successful execution of a program where one happens not to use any functionality that requires libraries which are not present locally. Moreover, the ‘usage point’ could be expressed either explicitly (as with a call to the Unix `dlopen` dynamic loader) or implicitly, when a module field appears in redex-position.

A full language should support this per-marshall choice, but for simplicity Acute supports only one of the alternatives: it allows execution of partial programs, and no linking is forced at unmarshal time. Instead, when an unlinked $M.x$ appears in redex position we look for an M to link the import to.

4.6 Controlling what to rebind to

How to look for such an M is specified by a *resolvespec* that can (optionally) be included in the import. By default it will be looked for just in the running program, in the sequence of modules defined above the import. Sometimes, though, one may wish to search in the local filesystem (e.g. for conventional shared-object names such as `libc.so.6`), or even at a web URI. In Acute we make an ad-hoc choice of a simple *resolvespec* language: a *resolvespec* is a finite list of *atomic resolvespecs*, each of which is either `Static_Link`, `Here_Already` or a URI. Lookup attempts proceed down the list, with `Static_Link` indicating the import should already be linked, `Here_Already` prompting a search for a suitable module (with the right name, signature and version) in the running program, and a URI prompting a file to be fetched and examined for the presence of a suitable module.

There is a tension between a restricted and a general *resolvespec* language. Sometimes one may need the generality of arbitrary computation (as in Java classloaders), e.g. for the negotiation scenario above, or as in browsers that dynamically discover where to obtain a newly-required plugin. On the other hand, a restricted language makes it possible to analyse a program to discover an upper bound on the set of modules it may require — necessary if one is marshalling it to a disconnected device, say. A full language should support both, though the majority of programs might only need the analysable sublanguage.

This *resolvespec* data is added to imports, for example:

```
import M : sig val y:int end version * by
  "http://www.cl.cam.ac.uk/users/pes20/acute/M.ac"
  = unlinked
M.y + 3
```

Here the $M.y$ is in redex-position, so the runtime examines the *resolvespec* list associated with the import of M . That list has just a single element, the URI `http://www.cl.cam.ac.uk/users/pes20/acute/M.ac`. The file there will be fetched and (if it contains a definition of M with the right signature) the modules it contains will be added to the running program just before the import, which will be linked to the definition of M . The $M.y$ can then be instantiated with its value.

URI *resolvespecs* are, of course, a limited form of distributed reference.

Note that this mechanism is not an exception — after M is loaded, the $M.y$ is instantiated in its original evaluation context $_ + 3$. It could be encoded (with exceptions and affine continuations, or by encoding imports as option references) but here we focus on the user language.

One would like to be able to limit the resources that a particular unmarshal could rebind to, e.g. to sandboxed versions of libraries, to securely encapsulate untrusted code. This was possible in our earlier λ -calculus work [BHS⁺03], but to support sufficiently-flexible limits here it seems necessary to have more structure than the Acute linear sequence of marks and modules.

4.7 The relationship between modules and the filesystem

Programs are decomposed not just into modules, but into separate source files. We have to choose whether (1) source files correspond to modules (as in OCaml, where a file named `foo.ml` implicitly defines a module `Foo`), or (2) source files contain sequences of module definitions, and are logically concatenated together in the build process, or (3) both are possible. As we shall see in the following sections, to deal with version change we sometimes need to refer to the results of previous builds. For Acute we take the simplest possible structure that supports this, following (2) with files containing compilation units:

```
compilationunit ::=
  empty
  e
  sourcedefinition ;; compilationunit
  includesource sourcefilename ;; compilationunit
  includecompiled compiledfilename ;; compilationunit
```


The result of compilation is a compiled unit which is just a sequence of compiled module definitions followed by an optional expression.

```
compiledunit ::=  
  empty  
  e  
  definition ;; compiledunit
```

This means that the decomposition of a program into files does not affect its semantics, except that when code is loaded by a URI *resolvespec* an entire compiled unit is loaded.

In Acute any modules shipped with a marshalled value are loaded into the local runtime, but are not saved to local persistent store to be available to future runtime instances. One could envisage a closer integration of communication and package installation.

4.8 Module initialisation

In ML, module evaluation can involve arbitrary computation. For example, in

```
module fresh M : sig val x: int ref  val y:unit          end  
  = struct let x=ref 3      let y=IO.print_int !x end
```

the structure associates non-value expressions to both *x* and *y*; the evaluation to a structure value involves expression evaluation which has both store and IO effects. The store effect enables per-module state to be created.

This is also possible in Acute, though as we shall see in §5 it is necessary to distinguish between modules that have such initialisation effects and modules that do not. The evaluation order for a single sequential program is straightforward: a program is roughly a sequence of module definitions followed by an expression; the definitions are evaluated in that order, followed by the expression.

New module definitions can be introduced dynamically, both by unmarshalling and fetched via *resolvespecs*. The evaluation order ensures that any modules that must be marshalled have already been evaluated, and so unmarshalling only ever adds module value definitions to the program.

Consider now the definitions fetched via a *resolvespec*. As we do not have cyclic linking, these definitions must be added before the **import** that demanded them. One could allow such definitions to be compiled units of unevaluated definitions. In the sequential case this would be straightforward: simply by evaluating the extant definition list in order, any newly-added definitions would be evaluated before control returns to the program below. With concurrency, however, there may be multiple threads referring to an import that triggers the addition of new definitions, and some mechanism would be required to block linking of that import until they are fully evaluated (or, equivalently, block instantiation from each new definition until it is evaluated). This flow of control seems complex both from the programmer's point of view and to express in the semantics; we therefore do not allow non-evaluated definitions to be fetched via a *resolvespec*. We return to the interaction between module initialisation and concurrency in §9.8.

In a language with finer-grain control of linking (for the negotiation discussed in §4.2) one might want more control over initialisation, allowing clients to demand their own freshly-initialised occurrences of modules, but Acute does not support this at present.

4.9 Marshalling references

Acute contains ML-style references, so we have to deal with marshalling of values that include store locations. For example:

```
let (x:int ref) = ref %[int] 5 in IO.send( marshal "StdLib" x : int ref)  
—  
IO.print_int ( ! %[int] (unmarshal (IO.receive ()) as int ref ))
```

Here the best choice for the core language semantics seems to be for the marshalled value to include a copy of the reachable part of the store, to be disjointly added to the store of any unmarshaller. Just as in §4.1 we reject the

alternative of building in automatic conversion of local references to distributed references, as no single distributed semantics (which here should include distributed garbage collection) will be satisfactory for all applications. A full language must be rich enough to express distributed store libraries above this, of course, and perhaps also other constructs such as those of [SY97, Bou03].

Some applications would demand distributed references together with distributed garbage collection (as JoCaml provides [Fes01]). We leave investigation of this, and of the type-theoretic support it requires, to future work.

One might well add more structure to the store to support more refined marshalling. In particular, one can envisage *regions* of local and of distributed store, perhaps related to the mark structure. We leave the development of this to future work also.

5 Naming: global module and type names

We now turn to marshalling and unmarshalling of values of abstract types. In ML, and in Acute, abstract types can be introduced by modules. For example, the module

```
module EvenCounter
: sig
  type t          = struct
    type t=int
    val start:t    let start = 0
    val get:t->int  let get = fun (x:int)->x
    val up:t->t     let up = fun (x:int)->2+x
  end
end
```

provides an abstract type `EvenCounter.t` with representation type `int`; this representation type is not revealed in the signature above. The programmer might intend that all values of this type satisfy the ‘even’ invariant; they can ensure this, no matter how the module is used, simply by checking that the `start` and `up` operations preserve evenness.

Now, for values of type `EvenCounter.t`, what should the unmarshal-time dynamic type equality check of §3 be? It should ensure not just type safety w.r.t. the representation type, but also *abstraction safety* — respecting the invariants of the module. Within a single program, and for communication between programs with identical sources, one can compare such abstract types by their source-code paths, with `EvenCounter.t` having the same meaning in all copies (this is roughly what the manifest type and singleton kind static type systems of Leroy [Ler94] and Harper et al [HL94] do).

For distributed programming we need a notion of type equality that makes sense at runtime across the entire distributed system. This should respect abstraction: two abstract types with the same representation type but completely different operations will have different invariants, and should not be compatible. Moreover, we want common cases of interoperation to ‘just work’: if two programs share an (effect-free) module that defines an abstract type (and share its dependencies) but differ elsewhere, they should be able to exchange values of that type.

We see three cases, with corresponding ways of constructing globally-meaningful type names.

Case 1 For a module such as `EvenCounter` above that is effect-free (i.e. evaluation of the structure body involves no effects) we can use module *hashes* as global names for abstract types, generalising our earlier work [LPSW03]. The type `EvenCounter.t` is compiled to `h.t`, where the hash `h` is (roughly)

```
hash(
  module EvenCounter
  : sig
    type t          = struct
      type t=int
      val start:t    let start = 0
      val get:t->int  let get = fun (x:int)->x
      val up:t->t     let up = fun(x:int)->2+x
    end
  end
)
```

i.e. the hash of the module definition (in fact, of the abstract syntax of the module definition, up to alpha equivalence and type equality, together with some additional data). If one unmarshals a pair of type `EvenCounter.t` *

`EvenCounter.t` the unmarshal type equality check will compare with $h.t * h.t$. This allows interoperation to just work between programs that share the `EvenCounter` source code, without further ado.

In constructing the hash for a module M we have to take into account any dependencies it has on other modules M' , replacing any type and term references $M'.t$ and $M'.x$. In our earlier work we did so by substituting out the definitions of all manifest types and terms (replacing abstract types by their hash type name). Now, to avoid doing that term substitution in the implementation, we replace $M'.x$ by $h'.x$, where h' is the hash of the definition of M' . This gives a slightly finer, but we think more intuitive, notion of type equality. We still substitute out the definitions of manifest types from earlier modules. This is forced: in a context where $M.t$ is manifestly equal to `int`, it should not make any difference to subsequent types which is used.

Case 2 Now consider effect-full modules such as the `NCounter` module below, where evaluating the `up` expression to a value involves an IO effect.

```

module fresh NCounter
: sig          = struct
  type t      type t=int
  val start:t  let start = 0
  val get:t->int let get = fun (x:int)->x
  val up:t->t  let up =
                let step=IO.read_int() in
                fun (x:int)->step+x
end          end

```

This reads an `int` from standard input at module initialisation time, and the invariant — that all values of type `NCounter.t` are a multiple of that `int` — depends on that effect. For such effect-full modules a fresh type name should be generated each time the module is initialised, at run-time, to ensure abstraction safety.

Case 3 Returning to effect-free modules, the programmer may wish to *force* a fresh type name to be generated, to avoid accidental type equalities between different ‘runs’ of the distributed system. A fresh name could be generated each time the module is initialised, as in the second case, or each time the module is compiled. This latter possibility, as in our earlier work [Sew01], enables interoperation between programs linked against the same compiled module, while forbidding interoperation between different builds.

For abstract types associated with modules it suffices to generate hashes or fresh names h per module, using the various $h.t$ as the global type names for the abstract types of that module.

We let the programmer specify which of the three behaviours is required with a `hash`, `fresh`, or `cfresh` mode in the module definition, writing e.g. `module hash EvenCounter`. In general it would be abstraction-breaking to specify `hash` or `cfresh` for an effect-full module. To prevent this requires some kind of effect analysis, for which we use coarse but simple notions of *valuability*, following [HS00], and of *compile-time valuability*. We say a module is valuable if all of the expressions in its structure are and if its types are hash-generated. The set of valuable expressions is intermediate between the syntactic values and the expressions that a type-and-effect system could identify as effect-free, which in turn are a subset of the semantically effect-free expressions. They can include, e.g., applications of basic operators such as `2+2`, providing useful flexibility.

The compile-time valuable, or *cvaluable*, modules can also include `cfresh` but otherwise are similar to the valuable modules. The *non-valuable* modules are those that are neither valuable nor *cvaluable*. If none of the `fresh`, `hash` or `cfresh` keywords are specified then a valuable module defaults to `hash`; a *cvaluable* module defaults to `cfresh`; and a non-valuable module must be `fresh`. On occasion it seems necessary to override the valuability checks, which we make possible with `hash!` and `cfresh!` modes. This is discussed in §8.3.

Acute also provides first-class System F existentials, as the experience with `Pict` [PT00] and `Nomadic Pict` [SWP99, US01] demonstrates these are important for expressing messaging infrastructures. For these a fresh type name will be constructed at each `unpack`, to correspond with the static type system.

6 Naming: expression names

Globally-meaningful *expression-level names* are also needed, primarily as interaction handles — dispatch keys for high-level interaction constructs such as asynchronous channels, location-independent communication, reliable messaging, multicast groups, or remote procedure (or function/method) calls. For any of these an interaction involves the communication of a pair of a handle and a value. Taking asynchronous channels as a simple example, these pairs comprise a channel name and a value sent on that channel. A receiver dispatches on the handle, using it to identify a local data structure for the channel (a queue of pending messages or of blocked readers). For type safety, the handle should be associated with a type: the type of values carried by the channel. (RPC is similar except that an additional affine handle must also be communicated for the return value.)

In Acute we build in support for the generation and typing of name expressions, leaving the various and complex dynamics of interaction constructs to be coded up above marshalling and byte-string interaction. As in FreshOCaml [SPG03], for any type T we have a type

T name

of names associated with it. Values of these types (like type names) can be generated freshly at runtime, freshly at compile-time, or deterministically by hashing, with expression forms `fresh`, `cfresh`, `hash(M.x)`, `hash(T, e)`, and `hash(T, e, e)`. We detail these forms below, showing how they support several important scenarios. In each, the basic question is how one establishes a name shared between sender and receiver code such that testing equality of the name ensures the type correctness of communicated values.

The expression `fresh` evaluates to a fresh name at run-time. The expression `cfresh` evaluates to a fresh name at compile-time. It is subject to the syntactic restriction that it can only appear in a compile-time valuable context. The expression `hash(M.x)` compiles to the hash of the pair of n and the label x , where n is the (hash- or fresh-)name associated with module M , which must have an x component. The expression `hash(T, e)` evaluates e to a string and then computes the hash of that string paired with the runtime representation of T . (Recall that a string can be injectively generated from an arbitrary value by marshalling). The expression `hash(T, e2, e1)` evaluates $e1$ to a T' name and $e2$ to a string, then hashes the triple of the two and T .

Each name form generates T names that are associated with a type T . For `fresh` and `cfresh` it is the type annotation; for `hash(M.x)` it is the type of the x component of module M ; for `hash(T, e)` it is T itself; and for `hash(T, e2, e1)` it is T . Of these, `fresh` is non-valuable; `cfresh` is compile-time valuable; `hash(M.x)` has the same status as M ; and `hash(T, e)` and `hash(T, e2, e1)` have the join of the status of their component parts.

(A purer collection of hash constructs, equally expressive, would be `hash(T)`, `hash(e1, e2)` (of a name and a string) and `hash(e1, T)` (of a name and a type). We chose the set above instead as they seem to be the combinations that one would commonly wish to use.)

6.1 Establishing shared names

For clarity we focus on distributed asynchronous messaging, supposing a module `DChan` which implements a distributed `DChan.send` by sending a marshalled pair of a channel name and a value across the network.

```
module hash DChan :
  sig
    val send : forall t. t name * t -> unit
    val recv : forall t. t name * (t -> unit) -> unit
  end
```

This uses names of type T name as channel names to communicate values of type T .²

Scenario 1 The sender and receiver both arise from a single execution of a single build of a single program. The execution was initiated on machine A, and the receiver is present there, but the sender was earlier transmitted to machine B (e.g. within a marshalled lambda abstraction).

²Acute does not yet support user-definable type constructors. If it did we would define an abstract type constructor `Chan.c : Type -> Type` and have `send : forall t. t Chan.c name * t -> unit`.

Here the sender and receiver can originate from a single lexical scope and a channel name can be generated at runtime with a `fresh` expression. This might be at the expression level, e.g.

```
let (c : int name) = fresh in
with sender DChan.send %[int] (c,v) and receiver DChan.recv %[int] (c,f) for some v:int and
f:int->unit (the %[int] is an explicit type application), or a module-level binder
```

```
module M : sig    val c : int name    end
= struct let c = fresh    end
```

generating the fresh name when the `let` is evaluated or the `module` is initialised respectively. This first scenario is basically that supported by JoCaml and Nomadic Pict.

Commonly one might have a single receiver function for a name, and tie together the generation of the name and the definition of the function, with an additional `DChan` field

```
val fresh_recv : forall t. (t -> unit) -> t name
```

implemented simply as

```
Function t -> fun f ->
  let c=fresh in DChan.recv %[t] (c,f); c
```

and used as below.

```
module M : sig    val c : int name    end
= struct let c = DChan.fresh_recv %[int]
  (fun x -> IO.print_int x+1) end
```

Note that this `M` is an effect-full module, creating the name for `c` at module initialisation time.

Scenario 2 The sender and receiver are in different programs, but both are statically linked to a structure of names that was built previously, with expression `cfresh` for compile-time fresh generation.

Here one has a repository containing a compiled instance of a module such as

```
module cfresh M : sig val c : int name end
= struct let c = cfresh end
```

in a file `m.aco`, which is included by the two programs containing the sender and receiver:

```
includecompiled "m.aco"
DChan.send %[int] (M.c,v)
—
includecompiled "m.aco"
DChan.recv %[int] (M.c,f)
```

Different builds of the sender and receiver programs will be able to interact, but rebuilding `M` creates a fresh channel name for `c`, so builds of the sender using one build of `M` will not interact with builds of the receiver using another build of `M`.

This can be regarded as a more disciplined alternative to the programmer making use of an explicit off-line name (or GUID) generator and pasting the results into their source code.

Scenario 3 The sender and receiver are in different programs, but both share the source code of a module that defines the function `f` used by the receiver; the hash of that module (and the identifier `f`) is used to generate the name used for communication.

This covers the case in which the sender and receiver are different execution instances of the same program (or minor variants thereof), and one wishes typed communication to work without any (awkward) prior exchange of names via the build process or at runtime. The shared code might be

```
module hash N : sig    val f : int -> unit    end
= struct let f = fun x->IO.print_int (x+100) end
```

```

module hash M : sig   val c : int name       end
= struct let c = hash(int,"",hash(N.f) %[]) %[] end

```

in a file `nm.ac`, included by the two programs containing the sender and receiver:

```

includesource "nm.ac"
DChan.send %[int] (M.c,v)
—
includesource "nm.ac"
DChan.recv %[int] (M.c,N.f)

```

The `hash(N.f)` gives a T name where $T = \text{int} \rightarrow \text{unit}$ is the type of `N.f`; the surrounding hash coercion `hash(int,"",-)` constructs an `int name` from this.³ This involves a certain amount of boiler-plate, with separate structures of functions and of the names used to access them, but it is unclear how that could be improved.

It might be preferable to regard the hash coercion as a family of polymorphic operators, indexed by pairs of type constructors $\Lambda \vec{t}. T_1$ and $\Lambda \vec{t}. T_2$ (of the same arity), of type $\forall \vec{t}. T_1 \text{ name} \rightarrow T_2 \text{ name}$.

Scenario 4 The sender and receiver are in different programs, sharing no source code except a type and a string; the hash of the pair of those is used to generate the name used for communication.

```

let c = hash(int,"foo") %[] in
DChan.send %[int] (c,v)
—
let c = hash(int,"foo") %[] in
DChan.recv %[int] (c,f)

```

This idiom requires the minimum shared information between the two programs. It can be seen as a disciplined, typed, form of the use of untyped “traders” to establish interaction media between separate distributed programs.

Scenario 5 The sender and receiver have established by some means a single typed shared name `c`, but need to construct many shared names for different communication channels. The hash coercion can be used for this also, constructing new typed names from old names, new types, and arbitrary strings. Whether this will be a common idiom is unclear, but it is easy to provide and seems interesting to explore. For example:

```

let c1 = hash(int,"one",c)
let c2 = hash(int,"two",c)
let c3 = hash(bool,"",c)
DChan.send %[int] (c1,v1); DChan.send %[int] (c2,v2); DChan.send %[bool] (c3,v3);
—
let c1 = hash(int,"one",c)
let c2 = hash(int,"two",c)
let c3 = hash(bool,"",c)
DChan.recv %[int] (c1,f1); DChan.recv %[int] (c2,f2); DChan.recv %[bool] (c3,f3);

```

Whether this will be a common idiom is unclear, but it is easy to provide and seems interesting to explore.

6.2 A refinement: ties

Scenario 3 of §6.1 above used a `hash(N.f)` as part of the construction of a name `M.c` used to access the `N.f` function remotely, linking the name and function together with a call `DChan.recv (M.c,N.f)`. It may be desirable to provide stronger language support for establishing this linkage, making it harder to accidentally use an unrelated name and function pair. For this, we propose a built-in abstract type

`T tie`

of those pairs, with an expression form `M@x` that constructs the pair of `hash(M.x)` and the value of `M.x` (of type T `tie` where $M.x : T$), and projections from the abstraction type `name_of_tie` and `val_of_tie`.

³Such coercions support `Chan.c` type constructors too, e.g. to construct an `int Chan.c name` from an `(int->unit) name`.

6.3 Polytypic name operations

We include the basic polytypic FreshOCaml expressions for manipulating names:

```
swap e1 and e2 in e3
e1 freshfor e2
support %[T] e
```

Here `swap` interchanges two names in an arbitrary value, `freshfor` determines whether a name does not occur free in an arbitrary value, and `support` calculates the set of names that do occur free in an arbitrary value (returning them as a duplicate-free list, at present).

We anticipate using these operations in the implementation of distributed communication abstractions. For example, when working with certain kinds of distributed channel one must send routing information along with every value, describing how any distributed channels mentioned in that value can be accessed.

We do not include the FreshOCaml name abstraction and pattern matching constructs just for simplicity — we foresee no difficulty in adding them.

In contrast to FreshOCaml, when one has values that mention store locations, the polytypic operations have effect over the reachable part of the Acute heap. This seems forced if we are to both (a) implement distributed abstractions, as above, and (b) exchange values of imperative data type implementations.

For constructing efficient datastructures over names, such as finite maps, we provide access to the underlying order relation, with a comparison between any two names of the same type.

```
compare_name %[T] : T name -> T name -> int
```

This cannot be preserved by name swapping, obviously, and so it would be an error to use it under any name abstraction, and in any other place subject to swapping. Nonetheless, the performance cost of not including it is so great we believe it is required. To ameliorate the problem slightly one might add a type

```
T nonswap
```

with a single constructor `Nonswap` that can be used to protect structures that depend on the ordering, with `swap` either stopping recursing or raising an exception if it encounters the `Nonswap` constructor. For the time being, however, `T nonswap` is not included in Acute.

6.4 Implementing names

In the implementation, all names are represented as fixed-length bit-strings (e.g. from 2^{160}) — both module-level and expression-level names, generated both by hashes and freshly. The representations of fresh names are generated randomly. More specifically: we do not want to require that the implementation generates each individual name randomly, as that would be too costly — we regard it as acceptable to generate a random start point at the initialisation of each compilation and the initialisation of each language runtime instance, and thereafter use a cheap pseudo-random number function for compile-time fresh and run-time fresh (the successor function would lead to poor behaviour in hash tables). This means that a low-level attacker would often be able to tell whether two names originated from the same point, and that (for making real nonces etc) a more aggressively random `fresh` would be required.

There is a possible optimisation which could be worthwhile if many names are used only locally: the bit-string representations could be generated lazily, when they are first marshalled, with a finite map associating local representations (just pointers) to the external names which have been exported or imported. This could be garbage-collected as normal. Whether the optimisation would gain very much is unclear, so we propose not to implement it now (but bear in mind that local channel communication should be made very cheap).

In order to implement the polytypic name operations the expression-level names must be implemented with explicit types.

7 Versions and version constraints

In a single-executable development process, one ensures the executable is built from a coherent set of versions of its component modules by choosing what to link together — in simple cases, by working with a single code directory tree.

In the distributed world, one could do the same: take sufficient care about which modules one links and/or rebinds to. Without any additional support, however, this is an error-prone approach, liable to end up with semantically-incoherent sets of versions of components interoperating. Typechecking can provide some basic sanity guarantees, but cannot capture these semantic differences.

One alternative is to permit rebinding only to identical copies of modules that the code was initially linked to. Usually, though, more flexibility will be required — to permit rebinding to modules with “small” or “backwards-compatible” changes to their semantics, and to pick up intentionally location-dependent modules. It is impractical to specify the semantics that one depends upon in interfaces (in general, theorem proving would be required at link time, though there are intermediate behavioural type systems). We therefore we introduce *versions* as crude approximations to semantic module specifications. We need a language of versions, which will be attached to modules, a language of version constraints, which will be attached to imports, a satisfaction relation, checked at static and dynamic link times, and an implication relation between constraints, for chains of imports.

Now, how expressive should these languages be? Analogously to the situation for *resolvespecs*, there is a tension between allowing arbitrary computation in defining the relations and supporting compile-time analysis. Ultimately, it seems desirable to make the basic module primitives parametric on abstract types of versions and constraints — in a particular distributed code environment, one may want a particular local choice for the languages. For Acute once again we choose not the most general alternative, but instead one which should be expressive enough for many examples, and which exposes some key design points.

Scenario 1 It is common to use version numbers which are supplied by the programmer, with no checked relationship to the code. As an arbitrary starting point, we take version numbers to be nonempty lists of natural numbers, and version constraints to be similar lists possibly ending in a wildcard `*` or an interval; satisfaction is what one would expect, with a `*` matching any (possibly empty) suffix. Many minor enhancements are possible and straightforward. Arbitrarily, we enhance version constraints with closed, left-open and right-open intervals, e.g. `1.5-7`, `1.8.-7`, and `2.4.7-`. These are certainly not exactly what one wants (they cannot express, for example, the set of all versions greater than `2.3.1`) but are indicative. The *meanings* of these numbers and constraints is dependent on some social process: within a single distributed development environment one needs a shared understanding that new versions of a module will be given new version numbers commensurate with their semantic changes.

Scenario 2 To support tighter version control than this, we can make use of the global module names (hash- or freshly-generated) introduced in §5: equality testing of these names is an implementable check for module semantic identity. We let version numbers include `myname` and version constraints include module identifiers `M` (those in scope, obviously). In each case the compiler or runtime writes in the appropriate module name. This supports a useful idiom in which code producers declare their exact identity as the least-significant component of their version number, and consumers can choose whether or not to be that particular. For example, a module `M` might specify it is version `2.3.myname`, compiled to `2.3.0xA564C8F3`; an import in that scope might require `2.3.M`, compiled to `2.3.0xA564C8F3`, or simply `2.3.*`; both would match it.

A key point is the balance of power between code producers and code consumers. The above leaves the code producer in control, who can “lie” about which version a module is — instead of writing `myname` they might write a name from a previous build. This is desirable if they know there are clients out there with an exact-name constraint but also know that their semantic change from that previous build will not break any of the clients.

Scenario 3 Finally, to give the code consumer more control, we allow constraints not only on the version field of a module but also on its actual name (which is unforgeable within the language). Typically one would have a definition of the desired version available in the filesystem (in Acute bringing it into scope as `M` with an `include`) and write `name=M`. (These exact-name constraints are also used to construct default `imports` when marshalling). One could also cut-and-paste a name in explicitly: `name=0xA564C8F3`. To guarantee that only mutually-tested collections of modules will be executed together, e.g. when writing safety-critical software, this would be the desired idiom everywhere, perhaps with development-environment support.

The current Acute version numbers and constraints, including all the above, are as follows.

<code>avne ::=</code>	Atomic version number expression
<code>n</code>	natural number literal
<code>N</code>	numeric hash literal
<code>myname</code>	name of this module


```

vne ::=      Version number expression
  avne | avne.vne

avce ::=      Atomic version constraint expression
  n           natural number literal
  N           numeric hash literal
  M           name of module M

dvce ::=      Dotted version constraint
  avce | n-n' | -n | n- | * | avce.dvce

vce ::=      Version constraint
  dvce        dotted version constraint
  name = M    exact-name version constraint

```

Version number and constraint expressions appear in modules and imports as below.

```

definition ::= ...
  module M:Sig version vne = Str ...
  | import M:Sig version vce ...
    by resolvespec = Mo

```

In constructing hashes for modules we also take into account their version expressions, to prevent any accidental equalities. That version expression can mention `myname`, and, as we do not wish to introduce recursive hashes, the hash must be calculated before compilation replaces `myname` with the hash.

It turns out that one needs exact-name version constraints not just for user-specified tight version constraints, as in the idiom above, but also during marshalling, when one may have to generate imports for module bindings that cross a mark. Exact-name constraints seem to be the only reasonable default to use there.

One might wish to extend the version language further with conjunctive version number expressions and disjunctive constraints. One might also wish to support cryptographic signatures on version numbers. Both would affect the balance of power between code producer and consumer, and further experience is needed to discover what is most usable.

Finally, we have had to choose whether version numbers are hereditary or not. A hereditary version number for a module `M` would include the version numbers of all the modules it depends on (and the version constraints of all the imports it uses), whereas a non-hereditary version number is just a single entity, as in the grammar above. The hereditary option clearly provides more data to users of `M`, but, concomitantly, requires those users to understand the dependency structure — which usually one would like a module system to insulate them from. If one really needs hereditary numbers, perhaps the best solution would be to support version number expressions that can calculate a number for `M` in terms of the numbers of its immediate dependencies, e.g. adding tuples and `version(M)` expressions to the `avne` grammar.

Just as for *withspecs* one might need rich development-environment support. Local specifications of version constraints, spread over the imports in the source files of a large software system, could be very inconvenient. One might want to refer to the version numbers of a source-control system such as CVS, for example.

8 Interplay between abstract types, rebinding and versions

8.1 Definite and indefinite references

With conventional static linking, module references such as `M.t` are *definite*, in the terminology of [HP]: in any fully-linked executable there is just a single such `M`, though (with separate compilation) it may be unknown at compile-time

which module definition for M it will be linked to. In contrast, the possibility of rebinding makes some references *indefinite* — during a single distributed execution, they may be bound to different modules.

For example, consider a module that declares an abstract type that depends on the term fields of some other module:

```
module M : sig    val f:int->int        end
              = struct let f=fun(x:int)->x+2 end
module EvenCounter
: sig          = struct
  type t      type t=int
  val start:t let start = 0
  val get:t->int let get = fun (x:int)->x
  val up:t->t   let up = fun (x:int)->M.f x
end           end
```

In the absence of any rebinding, the runtime type name for the abstract type `EvenCounter.t` would be the hash of the `EvenCounter` abstract syntax with `M.f` replaced by `h.f`, where `h` is the hash of the abstract syntax of `M`. This dependence on the `M` operations guarantees type- and abstraction-preservation.

Now, however, if there is a mark between the two module definitions, a marshal can cut and rebind to any other module with the same signature, perhaps breaking the invariant of `EvenCounter.t` that its values are always even. The `M.f` module reference below is indefinite.

```
module M : sig    val f:int->int        end
              = struct let f=fun (x:int)->x+2 end
import M : sig val f:int->int end version * = M
mark "MK"
module EvenCounter
: sig          = struct
  type t      type t=int
  val start:t let start = 0
  val get:t->int let get = fun (x:int)->x
  val up:t->t   let up = fun (x:int)->M.f x
end           end
IO.send(marshal "MK" (fun ()->EvenCounter.get
(EvenCounter.up EvenCounter.start)):unit->int)
—
module M : sig    val f:int->int        end
              = struct let f=fun (x:int)->x+3 end
(unmarshal (IO.receive ()) as unit->int) ()
```

To prevent this kind of error one can use a more restrictive version constraint in the import of `M` that `EvenCounter` uses, either by using an exact-name constraint `name=M` to allow linking only to definitions of `M` that are identical to the definition in the sender, or by using some conventional numbering. If no import is given explicitly, an exact-name constraint is assumed.

The version constraint should be understood as an assertion by the code author that whatever `EvenCounter` is linked with, so long as it satisfies that constraint (and also has an appropriate signature, and is obtained following any *resolvespec* present), the intended invariants of `EvenCounter.t` will be preserved.

Now, what should the global type name for `EvenCounter.t` be here? Note that the original author might not have had any `M` module to hand, and even if they did (as above), that module is not privileged in any way: `EvenCounter` may be rebound during computation to other `M` matching the signature and version constraint. In generating the hash for `EvenCounter`, analogously to our replacement of definite references `M'.x` by the hash of the definition of `M'`, we replace indefinite import-bound references such as `M.f` by the hash of the *import*. This names the set of all `M` implementations that match that signature and version constraint.

In the case above this hash would be roughly

```
hash(import M:sig val f:int->int end version * )
```

and where one imports a module with an abstract type field

```
import M : sig type t val x:t end
  version 2.4.7- ...
```

the hash h =

```
hash(import M : sig type t val x:t end
  version 2.4.7- ...)
```

provides a global name $h.t$ for that type.

In the `EvenCounter` example, the imported module had no abstract type fields. Where they do, for type soundness we have to restrict the modules that the import can be linked to, to ensure that they all have the same representation types for these abstract type fields. We do so by requiring imports with abstract type fields to have a *likespec* (in place of the ... above), giving that information. A compiled *likespec* is essentially a structure with a type field for each of the abstract type fields of the import.

At first sight this is quite unpleasant, requiring the importers of a module to ‘know’ representation types which one might expect should be hidden. With indefinite references to modules with abstract types, however, some such mechanism seems to be forced, otherwise no rebinding is possible. Moreover, in practice one would often have available a version of the imported library from which the information can be drawn. For example, one might be importing a graphics library that exists in many versions, but for which all versions that share a major version number also have common representations of the abstract types of point, window, etc. A typical import might have the form

```
import Graphics:sig type t end version 2.3.*
  like Graphics2_0
```

(with more types and operations) where `Graphics2.0` is the name of a graphics module implementation, which is present at the development site, and which can be used by the compiler to construct a structure with a representation for each of the abstract types of the signature.

While the abstraction boundaries are not as rigid as in ML, this should provide a workable idiom for dealing with large modular evolving systems, supporting rebinding but also allowing one to restrict type representation information to particular layers. The only alternative seems to be to make all types fully concrete at interfaces where rebinding may occur.

To correctly deal with abstract types defined by modules following an import, which use abstract type fields of the imported module in their representation types, compiled *likespecs* must be included in the hashes of imports.

On the other hand, we choose not to include *resolvespecs* in import hashes. This is debatable — the argument against including them is that it is useful to be able to change the location of code without affecting types, and so without breaking interoperation (e.g. to have a local code mirror, to change a web code repository to avoid a denial-of-service attack etc.).

Note that the indefinite character of our imports makes them quite different from module imports that are resolved by static linking, where typing can simply use module paths to name any abstract types and no *likespec* machinery is required. Both mechanisms are needed.

8.2 Breaking abstractions

In ongoing software evolution, implementations of an abstract type may need to be changed, to fix bugs or add functionality, while values of that type exist on other machines or in a persistent store. It is often impractical to simultaneously upgrade all machines to a new implementation version.

A simple case is that in which the representation of the abstract type is unchanged and where the programmer asserts that the two versions have compatible invariants, so it is legitimate to exchange values in both directions. This may be the case even if the two are not identical, e.g. for an efficiency improvement or bug fix. Here there should be some mechanism for forcing the old and new types to be identical, breaking the normal abstraction barrier.

In [Sew01, LPSW03] we proposed a *strong coercion* with! to do so, and Acute includes a variant of this. By analogy with ML sharing specifications, we allow a module definition to have a *withspec*, a list of equalities between abstract types and representations of modules constructed earlier (often this will be of previous builds of the same module).

```
definition ::= ...
```

```

    module M : Sig version vne = Str withspec
withspec      ::= empty | with! withspecbody
withspecbody ::= empty | M.t=T,withspecbody

```

The compiler checks the representation type of these $M.t$ are equal to the types specified (respecting any internal abstraction boundaries); if they are, the type equalities can be used in typechecking this definition.

For example, suppose the `EvenCounter` module definition of §5 was compiled to a file `p11_even.aco` and is widely deployed in a distributed system, and that later one needs a revised `EvenCounter` module, adding an operation or fixing a bug without making an incompatible type. A new module with an added down operation can be written as follows.

```

includecompiled "p11_even.aco"
module EvenCounter
: sig
    type t = EvenCounter.t
    val start:t
    val get:t->int
    val up:t->t
    val down:t->t
end
= struct
    type t=int
    let start = 0
    let get = fun (x:int)->x
    let up = fun (x:int)->2+x
    let down = fun (x:int)-> x-2
end
with! EvenCounter.t = int

```

In the interface here the type t of the new module is manifestly equal to the abstract type t of the previously-built module, and the `with!` enables the type equality between that abstract type and `int` to be used when typing the new module. The new type is compiled to be manifestly equal to (the internal hash-name of) the old type. (For this example, where the previous `EvenCounter` had a hash-generated type, one could include the source of the previous module rather than the compiled file, but if it were `cfresh`-generated the compiled file is obviously needed.)

The *withspec* is, in effect, a declaration by the programmer that the old and new implementations respect the same important invariants — here, that values of the representation type will always be even. In general they will not respect exactly the same invariants. For example, here the new module allows negative `ints`, but the programmer implicitly asserts that the clients of the old module will not be broken by this.

It would not suffice to check only that the new module respects at least the important invariants of the old, as if the types are made identical then values produced by either module can be acted upon by operations of the other.

In the more complex case where the old and new invariants are not compatible, or where the two representation types differ, the programmer will have to write an upgrade function. The same strong coercion can be used to make this possible, with a module that contains two types, one coerced to each. An example is given in [LPSW03].

There are several design options for *withspecs*. In our earlier proposals `with!` coerced an abstract type of the module being defined to be equal to an earlier abstract type. Here instead the `with!` simply introduces a type equality to the typechecking environment; manifest types in the signature of the new module can be used to make the type field of the compiled signature equal to the old. This simplifies the semantics slightly and may be conceptually clearer. We allow the *withspec* type equalities to be used both for typechecking the body of the new module and for checking that it does have the interface specified. One might instead only allow them to be used for the latter; it is unclear whether this would always be expressive enough. The programmer has to specify the representation type in a *withspec* explicitly. This is fine for small examples, e.g. the `int` above, but if the representation type is complex then it would be preferable to simply write `with! M.t`. That requires a somewhat more intricate semantics (as typechecking of modules with *withspecs* then depends on the representation types of earlier modules) and so we omit it for the time being. Finally, one might well want development-environment support, allowing collections of modules to be ‘pinned’ to the types in a particular earlier build without having to edit each module to add a *withspec* and make the types manifestly equal to the earlier ones.

8.3 Overriding valuability checks

The semantics for abstract type names outlined in §5 ensures that two instances of an effect-full module give rise to distinct abstract types. In general this is the only correct behaviour, as (as explained there) they may have very different invariants. In practice, however, one may often want to permit rebinding to modules which have some internal state. For example, in the communication library described in §11 the `Distributed_channel` module stores a `Tcp_string_messaging.handle` option which is set by calls to `Distributed_channel.init : Tcp.port -> unit`. One has to keep this as module state rather than threading a handle through the `Distributed_channel` interface calls so that those calls can be correctly rebound if (say) one marshals a function mentioning them. Despite the initialisation effect (evaluating `ref None`) we need the module name for `Distributed_channel` to be hash-generated, not fresh-generated, so that the abstract types in the interface are the same in different instance, so that rebinding can take place. The desired behaviour really is for the conceptually-distinct abstract types of different instances to be compatible. This could be expressed either

1. with module annotations `hash!` and `cfresh!`, which override the valuability check but otherwise are like `hash` and `cfresh`; or
2. with an expression form `ignore_effect(e)`, transparent at runtime but concealing arbitrary effects as far as valuability goes.

We choose the former, to make the coercion clearer in the module source and to avoid polluting the expression grammar, but the latter has the advantage of localising the coercion to where it is really needed.

8.4 Exact matching or version flexibility?

In §6 we focussed on name-based dispatch. An alternative idiom for remote invocation simply makes use of the dynamic rebinding facilities provided in `Acute`, e.g. as in the code below where a thunk mentioning `N.f` is shipped from one machine to another.

```
module N: sig val f: int->unit end
  = struct let f=fun x-> IO.print_int (x+1) end
mark "MARK-N"
IO.send (marshal "MARK-N" ((fun ()->N.f), 9))
—
module N: sig val f: int->unit end
  = struct let f=fun x-> IO.print_int (x+1) end
mark "MARK-N"
let (g, (y: int))=unmarshal(IO.receive()) in g () y
```

As the `marshal` is with respect to a mark ("MARK-N") below the definition of `N`, the pair of the thunk and `v` will be shipped together with an unlinked import for `N`; when the unmarshalled thunk is applied that import will become linked to the local definition of `N` on the receiver machine.

In the code as written the import will have an exact-name version constraint, but this could be liberalised by writing an explicit import in the sender, with an arbitrary version constraint.

This is quite different from the name-based dispatch of §6, where a simple name equality is checked for each communication. Here, a full link-ok check is involved, checking a subsignature relationship and a version constraint. It is therefore much more costly, but also allows much more flexible linking.

Another difference between the two schemes is that with name-based dispatch the receiver can express access-control checks by testing name equality, whereas here one would need to test equality of arbitrary incoming functions (against `fun ()->N.f` thunks), which we do not admit.

A common idiom may be to establish a shared structure of names by dynamic linking (including a version check) at the start of a lengthy interaction and thereafter to use name-based dispatch. `Acute` does not yet provide the low-level linking machinery needed for explicitly sending such a structure (see the discussion of negotiation elsewhere), so we do not explore this further here.

8.5 Marshalling inside abstraction boundaries

If one has a module defining an abstract type, and within that module marshals a value of that type, one has to choose whether it is marshalled abstractly or concretely. For example, in

```
module EvenCounter
: sig
  type t
  val start:t
  val get:t->int
  val up:t->t
  val send : t -> unit
  val recv : unit -> t
end
= struct
  type t=int
  let start = 0
  let get = fun (x:int)->x
  let up = fun (x:int)->2+x
  let send = fun (x:t) -> IO.send( marshal "StdLib" x : t)
  let recv = fun () -> (unmarshal(IO.receive()) as t)
end
EvenCounter.send (EvenCounter.start)
```

is the communicated value compatible with `int` or with `EvenCounter.t`? For Acute we take the former option: all types (in the absence of polymorphism) are fully normalised with respect to the ambient type equations before execution. Running the above in parallel with

```
IO.print_int(3+(unmarshal(IO.receive()) as int))
```

will therefore succeed.

One might well want more source-language control here, allowing the programmer to specify that such a `marshal` should be at the abstract type, but we leave this for future work (but cf. the comment on page ??). In general, with nested modules and with `with!` specifications, there may be a complex type equation set structure to select from.

9 Concurrency, mobility, and thunkify

Distributed programming requires support for local concurrency: some form of threads and constructs for interaction between them.

9.1 Language-level concurrency vs OS threads

The first question here is whether to fix a direct relationship to the underlying OS threads or take language-level threads to be conceptually distinct, which might or might not be implemented with one OS thread each. The former has the advantages of a simple relationship with the OS scheduler (which may provide rich facilities, e.g. for QoS, that some programs need) and the potential to exploit multiple processors. It has the disadvantages of different concurrency models on different OSs, and of a nontrivial relationship between threading and the language garbage collector. The latter gives the language implementor much more freedom. In particular, to support lightweight concurrency (as in Erlang, Pict, JoCaml etc.), in which many parallel components simply send a message or two, it is desirable for parallel composition to not require the (costly) construction of a new OS thread. For Acute we adopt language-level concurrency.

9.2 Interaction primitives

There are two main styles of interaction between threads: shared memory and message passing. The latter is a better fit to large-scale distributed programming and, we believe, often leads to more transparent code. The former, however, is needed when dealing with large mutable datastructures, and suits the imperative nature of ML/OCaml programming. In large programs we expect both to be required. In Acute we initially provide shared-memory interaction, as OCaml does: references can be accessed from multiple threads, with atomic dereferencing and assignment, and mutexes and condition variables can be used for synchronization. These enable certain forms of message-passing interaction to be expressed as library modules, which suffices for the time being. In future we expect to build in support for message-passing. Indeed, some forms require direct language support (or a preprocessor-based implementation), e.g. Join patterns with their multi-way binding construct.

9.3 Thunkification

We want to make it possible to checkpoint and move running computations — for fault-tolerance, for working with intermittently-connected devices, and for system management. Several calculi and languages (JoCaml, Nomadic Pict, Ambients, etc.) provided a linear migration construct, which moved a computation between locations.

It now appears more useful to support marshalling of computations, which can then be communicated, checkpointed etc. using whatever communication and persistent store constructs are in use. Taking a step further, as we have marshalling of arbitrary values, marshalling of computations requires only the addition of a primitive for converting a running computation into a value. We call this *thunkification*. Checkpointing a computation can then be implemented by thunkifying it, marshalling the resulting value, and writing it to disk. Migration can be implemented by thunkification, marshalling, and communication. Note that these are not in general linear operations — if a computation has been checkpointed to disk it may be restarted multiple times.

There are many possible forms of thunkification. The simplest is to be both subjective and synchronous: executing `thunkify` in a single thread gives a thunk of that thread, essentially capturing the (single-thread) continuation of the `thunkify`. Typically, though, the computation which one wishes to thunkify will be composed of a group of threads. The programmer would then have to manually ensure that all the threads synchronize and then thunkify themselves, and collect together the results. This would be very heavy, requiring substantial rewriting of applications to make them amenable to checkpointing or migration. Accordingly, we think it preferable to have an objective and asynchronous `thunkify`, freezing a group of threads irrespective of their current behaviour.

A group of threads may be intertwined with interaction primitives (i.e. mutexes and condition variables) used for internal communication and synchronization. Accordingly, `thunkify` should also be applicable to those interaction primitives.

Thunkification is destructive, removing the threads, mutexes and condition variables that are thunkified.

Thunkification of a group must be atomic. To see the inadequacy of a `thunkify` that operates only on a single thread, consider thunkifying a pair of threads, the first of which is performing a thread operation (e.g. `kill`) on the second. If the second is thunkified before the first then the `kill` will fail, whereas with an atomic multi-thread `thunkify` it will always succeed, either before the `thunkify` happens or after the group is unthunkified later.

9.4 Naming and grouping

Threads must be structured in some fashion. The simplest option, taken by many process calculi, is to have a running system be a flat parallel composition of anonymous threads. In contrast, operating system threads are typically named, with names provided by the system at thread creation time; these names may be reused over time and between runtimes.

For Acute some naming structure is required, to allow threads to be manipulated (thunkified, killed, etc.). We see two main possibilities:

1. globally-unique names, created freshly by the system at thread creation time; or
2. locally-unique names, provided by the programmer at thread creation time, with an exception if they are already in use on this runtime.

The other two possibilities are not useful or not implementable: if names are being created freshly by the system they might as well be globally unique, with the same representation as we use for other names; if names are being provided by the programmer then it is not in general possible to check if they are in use on any runtime.

We expect (1) to be the most commonly desired semantics. Nonetheless, in Acute we choose (2). Firstly, given (2) the programmer can implement (1) simply by providing a fresh name at each thread creation point. The difference between the two shows up when one moves a group of threads, which internally record and manipulate the thread names of the group, from one machine to another. With (1) they necessarily receive new names at the destination, so to maintain correctness all records of their old names must be permuted with the new — which may be awkward if there are external records of these names. With (2), if this movement is known to be linear then the original names can be reused without further ado.

The same two possibilities exist for the naming of interaction primitives for synchronization and communication between threads, i.e. (at present) mutexes and condition variables, and we make the same choice of (2) for them.

Many distributed process calculi have exploited a hierarchical group structure over processes, with boundaries delimiting units of migration, units of failure, synchronization regions, secure encapsulation boundaries, and administrative domains. There is a basic tension between the need for communication across boundaries and the need for encapsulation and control over untrusted components, giving rise to a complex design space which is not well-understood. The tutorial [Sew00] gives a very preliminary overview. How this tension should be resolved and what group structure should be provided as primitive is a very interesting question for future work. We conjecture that groups for migration and synchronization units can be expressed rather easily in Acute with flat parallel compositions of named threads, and that is what the language currently provides.

Any group structure should — presumably — also structure the interaction primitives (mutexes, channels, etc.) but here there are additional complications, as these are necessarily going to be used for interaction across a boundary, so the interactands may be split apart by thunkification.

A further motivation for richer group structure comes from performance requirements. When programming in a message-passing style (as in the π -calculus and in the derived languages JoCaml, Pict, and Nomadic Pict) one may have many threads which contain only a single asynchronous output. For performance it may be necessary to optimise these, not always creating thread names and scheduler entries for them. If threads can discover their own names, e.g. by a

```
self : unit -> thread name
```

primitive, then this optimisation is nontrivial: a thread which outputs the value of an expression involving `self` must have been created with a name, whereas outputs of other values need not. This led us to explore grouping structures of named groups containing anonymous threads. Ultimately we rejected them, returning to the flat parallel compositions of named threads, as they seemed excessively complex and it seemed likely that a rather simple static analysis would be able to identify most non-self outputs.

9.5 Thread termination

Acute threads do not return values, and their termination cannot be synchronized upon. We have no strong opinion about these choices, making them for simplicity for the time being. Thread termination is observable indirectly, as `thunkify` and `kill` raise exceptions if called on non-existent threads.

9.6 Nonexistent threads, mutexes, and condition variables

In conventional single-machine programming it is straightforward to ensure that any mutexes and condition variables used must already exist — in OCaml, for example, the type system guarantees this. In Acute, however, this is no longer possible.

Firstly, mutex names may be marshalled (either alone or in a function such as `function () -> unlock m`) and then unmarshalled on another machine. In the absence of thunkification it is debatable whether this is useful: one might imagine forbidding such examples, either with a dynamic check at marshal-time or a rich type system that identifies non-marshalleable types. With thunkification, however, one may certainly need to marshal a thunkified group of threads together with their internal mutexes. Secondly, thunkification can remove a mutex, leaving active threads

that refer to it. This scenario seems inescapable: if one moves some threads, they typically are going to have been interacting, in some fashion, with other threads at the source.

Accordingly, the mutex and condition variable operations may fail dynamically, giving `Nonexistent_mutex` and `Nonexistent_cvar` exceptions. One would expect high-level communication libraries, e.g. of distributed communication channels and migration, to ensure such errors never occur.

9.7 References, names, marshallng, and thunkify

Semantically, it is tempting to treat store locations as another variety of name, similar to thread and mutex names. In Acute we do not make this identification as the cost seems under-motivated. A naive implementation, indirecting all access via a name lookup, would obviously be absurd. Even an optimisation, using local pointers but keeping a name with every store value, would be rather expensive — in a typical program there are many more store locations than mutexes or threads (it would be necessary to keep a name for each explicitly as garbage collection can relocate pointers but the name order must be preserved).

Further, the dynamic semantics is rather different: marshallng copies the reachable fragment of the store, whereas names are simply marshalled as the values that they are. Thunkifying threads and mutexes is destructive, removing them from the running system. Copying the reachable fragment of the store ensures that dereferencing and assignment can never fail dynamically (which we think would be unacceptable) whereas the implicit marshallng of entire threads seems unlikely to be desirable. Further practical experience is required to assess these choices.

9.8 Module initialisation, concurrency, and thunkify

Without module initialisation all threads are simply executing an expression. With initialisation, however, at least one thread might be executing a sequence of definitions (followed by an expression), evaluating expressions on the right-hand-side of structures in programs as below.

```
module fresh M : sig   val x: int ref   val y:unit           end
                    = struct let x=ref 3   let y=IO.print_int !x end
M.x := 7
```

These expressions may spawn other threads, which may interact (via the store, mutexes etc.) with the first.

In fact, as discussed in §4.8, no uninitialised definitions can be dynamically added to the system, so it is an invariant that at most one thread is executing in definitions (though the semantics actually allows definitions in all threads, for uniformity).

The initial thread has no other special status.

Now, what should `thunkify` do if invoked on such a thread? Acute has a second-class module system, so there is (unfortunately) no way to represent a suspended module-level computation in the expression language. The `thunkify` must therefore either abort or block until module initialisation is complete. For the time being we take the former choice, raising a `Thunkify_thread_in_definition` exception.

9.9 Thunkify and blocking calls

With any form of thread migration or (more generally) with our thunkification one has to deal with threads that are blocked in system calls. There are two possibilities:

1. have the `thunkify` block until the target thread returns, thunkifying its state just after the return; or
2. have the `thunkify` return immediately, thunkifying the state of the target thread with a raise of a `Thunkify_EINTR` exception replacing the blocked call, and discarding the eventual return value of the call. This is analogous to the Unix `EINTR` error, returned when a system call is interrupted by a signal, which applications must be prepared to deal with.

Both are desirable, in different circumstances, and so we allow a per-thread choice. Note that this applies only to blocking (or “slow”) system calls such as `read()`, not to the many non-blocking system calls which return quickly. The language semantics must distinguish the two classes.

Taking this further, it is unpleasant for the system interface to be special in this way. For example, suppose one has a user library module that provides a wrapper around the system interface; one might want to identify some of the user module entry points as blocking and have similar `thunkify` behaviour. This would be conceptually straightforward if the functions provided by the module are all first-order and cannot be partially applied, in which case there is a straightforward notion of a thread executing ‘in’ the module. `thunkify` could behave as (2) as far as the calling thread is concerned and raise an asynchronous exception in the user library code. We believe this kind of mechanism is desirable, but have not explored it in detail.

9.10 Concurrency: the constructs

Putting these choices together, we have types

```
thread
mutex
cvar
thunkifymode
thunkkey
```

The first three types are empty; they are introduced to form types `thread name`, `mutex name`, and `cvar name`. A `thunkifymode` is either `Interrupting` or `Blocking`; type `thunkkey` has three constructors, `Thread`, `Mutex` and `CVar`, each taking a name of the associated type; the first takes also a `thunkifymode`.

We have operations for threads, mutexes, condition variables and thunkification as below.

```
create_thread : thread name -> (T->unit) -> T -> unit
self : unit -> thread name
kill : thread name -> unit

create_mutex : mutex name -> unit
lock : mutex name-> unit
try_lock : mutex name -> bool
unlock : mutex name -> unit

create_cvar : cvar name -> unit
wait : cvar name -> mutex name -> unit
signal : cvar name -> unit
broadcast : cvar name -> unit

thunkify : thunkkey list -> thunkkey list -> unit

exit : int -> T
```

In addition, we have a control operator

```
e1 ||| e2
```

that spawns its first argument, as syntactic sugar for

```
create_thread fresh (function () -> e1); e2
```

Here `thunkify` takes a list of `thunkkeys` specifying which threads, mutexes and condition variables to thunkify; it returns a function which takes a list of the same shape specifying the names to give these entities and then atomically re-creates them.

9.11 Example

Below is a simple use of `thunkify`, capturing the state of a single running thread and an (unused) mutex.

```
let rec delay x = if x=0 then () else delay (x-1) in
let rec f x = IO.print_int x; IO.print_newline (); f (x+1) in
let t1 = fresh in
let m1 = fresh in
let _ = create_thread t1 f 0 in
let _ = create_mutex m1 in
let _ = delay 15 in
let v = thunkify ((Thread (t1,Blocking))::(Mutex m1)::[]) in
IO.send( marshal "StdLib" v : thunkkey list -> unit )
—
let rec delay x = if x=0 then () else delay (x-1) in
let exit_soon = create_thread fresh (fun () -> delay 15 ; exit 0) () in
let v = (unmarshal(IO.receive())) as thunkkey list -> unit in
v ((Thread (fresh,Blocking))::(Mutex fresh)::[])
```

When run the first program prints 0 1 2 3 4 and the second 5 6 7 8. The marshalled value, containing the thunk, is shown in §15.10.

10 Polymorphism

Ultimately, both subtype and parametric polymorphism should be included. Many version changes involve subtyping, e.g. the addition of fields to a manifest record type argument of a remote function; it should be possible to make these transparent to the callers. Parametric polymorphism is of course needed in some form for ML-style programming. In the distributed setting it seems to be particularly useful to have first-class universals, allowing polymorphic functions to be communicated, and first-class existentials. The latter support an idiom, common in Pict and Nomadic Pict, in which one packages a channel name and a value that can be sent on that channel, as a value of type $\exists t.t \text{ name} * t$. This lets one express communication infrastructure libraries that can uniformly forward messages of arbitrary types.

There are two substantial difficulties here. Firstly, type inference is challenging for such combinations of subtyping and parametric polymorphism. A partial type inference algorithm will be required, and it must be pragmatically satisfactory — inferring enough annotations, and unsurprising to the programmer. This is the subject of recent research on *local type inference* [PT98, HP99] and *coloured local type inference* [OZZ01]. Without subtyping, the MLF of Le Botlan and Rémy [LBR03] allows full System F but can infer types for all ML-typable programs.

Secondly, the interaction between subtyping and hash types requires further work — one can imagine, for instance, that a subhash order derived from subtype and subversion relationships needs to be dynamically propagated.

In Acute we sidestep both of these issues for the time being, making an interim choice that suffices for writing non-trivial examples, e.g. of polymorphic communication infrastructure modules. Acute has no subtyping. The basic scheme is monomorphic, but with type inference. The definition of the internal language has explicit type annotations, on pattern variables and on built-in constructors such as `[]` and `None`. In the external language these annotations can all be inferred by a unification-based algorithm. To this we add first class System F universals and existentials, with types `forall t.T` and `exists t.T` and explicit type abstractions, applications, packs and unpacks, with expression forms

```
Function t -> e
e %[T]
{T, e} as T'
let {t, x} = e1 in e2
```

There is no automatic generalisation, and the subsignature relation remains, as in the monomorphic case, without generalisation. We also have no user-definable type constructors. The expression forms could easily be more tightly integrated with the other pattern matching and function forms.

Traditional ML implementations can erase all types before execution. In contrast, an Acute runtime needs type representations at marshal and unmarshal points, to execute the expressions `marshal e : T` and `unmarshal e as T`. (These types can often be inferred). Type representations are also needed at `fresh`, `cfresh` and `hash(...)` points. Our prototype implementation keeps all type information, throughout execution, so that we can do runtime typechecking between reduction steps. A production implementation would probably do a flow analysis to determine where types are required, adding type representation parameters to functions as needed. The only operations that a production implementation needs to do on these type representations are (1) compare them for syntactic equality, (2) construct them when a polymorphic function is applied to its type parameter, and (3) take hashes of them. It is therefore not necessary to keep all the type structure. Indeed, one could (with a small probabilistic reduction in safety) work with hashes of types at runtime. Alternatively, if one keeps the structure it would be possible to add some form of runtime type analysis [Wei02] at little extra cost, at least for non-abstract types.

10.1 A refinement: marshal keys and name equality

In the implementation of distributed communication libraries one may often be communicating values of types such as `exists t. t name * T` (with the `t` potentially occurring in `T`) where the `t name` is used as a demultiplexing/dispatch key at the receiver.

To statically type the receiver code an enhanced conditional or matching form is needed: having compared that `t name` with the locally-stored name associated with (say) a channel data structure, typing the `true` branch must be in an environment where the two are known to be of the same type.

The enhanced form could be either an explicit type equality test or a name equality test. At present we do not see a strong argument either way. A type equality test is perhaps cleaner, but would lead to runtime type information being required at more program points. A general name equality test, if `e1=e2` then `e3` else `e4`, where `e1` and `e2` are of arbitrary `T1 name` and `T2 name` types, is the most obvious alternative, but this requires a slightly intricate treatment of multiple type equalities in the semantics. For the time being we combine name equality testing with existential unpacks, with

```
namecase e1 with {t, (x1, x2)} when x1=e
  -> e2
otherwise -> e3
```

where `e1:exists t. t name * T`, the `e:T' name` is evaluated first and used to build an equality pattern, and in the `e2` branch it is known that `t=T'`. Obviously such existentials are not uniformly parametric in Acute.

If one is communicating values of type `exists t. t name * t`, and is demultiplexing on the `t name`, the explicit type in the marshalled value (and the unmarshal-time type equality check) could be omitted; name equality gives an equally strong guarantee. If communicating many small values the performance gain of this could be worth direct language support for such ‘marshal keys’.

11 Pulling it all together: examples

To date, we have written many small examples in Acute (for automated testing), and three larger programs. The first two are `blockhead` and `minesweeper` games that mostly exercise local computation; the latter uses marshalling to save and restore the game state. The third is a communication infrastructure library which shows how most of the Acute features are needed and used. It has the following modules:

`Tcp_connection_management` maintains TCP connections to TCP addresses (IP address/port pairs), creating them on demand. `Tcp_string_messaging` uses that to provide asynchronous messaging of strings to TCP addresses. These are both hash modules, with abstract types of handles; they spawn daemons to deal with incoming communications.

Separately, a module `Local_channel` provides local (within a runtime) asynchronous messaging, again with an abstract type of channel management handles and with polymorphic `send:forall t. t name * t -> unit` and `recv:forall t. t name*(t->unit) -> unit` (to register a handler). Channel states are stored as existential packages of lists of pending messages or receptors; the `namecase` operation is used to manipulate them. Mutexes are needed for protection.

`Distributed_channel` pulls these together, with `send:forall t.string->(Tcp.addr*t name)->t->unit` (and a similar `recv`) for distributed asynchronous messaging to TCP addresses. The string names the mark to marshal with respect to. For a local address this simply uses `Local_channel`. For a remote address the `send` marshals its `t` argument and uses `Tcp_string_messaging`; the `recv` unmarshals and generates a local asynchronous output. This deals with the non-mobile case — active receivers cannot be moved from one runtime to another. However, code that uses this module, e.g. functions that invoke `send` and `recv`, can be marshalled and shipped between runtimes; the module initialisation state includes the TCP messaging handles and so rebinding to different instances of `send` and `recv` works correctly. A simple RFI module implements remote function invocation above distributed channels.

Clients of this libraries can use any of the various ways of creating shared typed names discussed in §6 and §8.4. Moreover, the use of first-class marks means that clients have the same flexible control over the marshalling that goes on as direct users of `marshal`.

Going further, a `Nomadic_pi` module supports mobility of running computations, with named *groups* of threads, each with a local channel manager, that can migrate between machines. Migration uses `thunkify` to capture the group’s channel and thread state. Threads within a group can interact via local channels; groups can interact with a location-dependent `send_remote` that sends a message to a channel of a group assumed to be at a particular TCP address.

The location-independent messaging algorithms of JoCaml or high-level Nomadic Pict should be easy to express above this (the former requiring the polytypic support and swap operations to manipulate the free channel names of a communicated value).

12 Related work

Acute builds on our earlier work: compile-time fresh generation of abstract type names and channel names [Sew01]; hash-generation of effect-free abstract type names [LPSW03]; and dynamic rebinding [BHS⁺03]. There is extensive related work on module systems, dynamic binding, dynamic type tests, and distributed process calculi. For most of this we refer the reader to the discussion in those papers, confining our attention here to some of the most relevant distributed programming language developments.

Early work on adding local concurrency to ML resulted in Concurrent ML [Rep99] and the initial Facile, both based on the SML/NJ implementation. Facile was later extended with rich support for distributed execution, including a notion of *location* and computation mobility [TLK96]. dML [OK93] was another distributed extension of ML, implementable by translation into remote procedure calls without requiring communication at higher types. Erlang [AVWW96] supports concurrency, messaging and distribution, but without static typing.

The Pict experiment [PT00] investigated how one could base a usable programming language purely on local concurrency, with a π -calculus core instead of primitive functions or objects. The Distributed Join Calculus [FGL⁺96] and subsequent JoCaml implementation [JoC] modified the π primitives with a view to distribution, and added location hierarchies and location migration. The runtime involved a complex forwarding-pointer distributed infrastructure to ensure that, in the absence of failure, communication was location-independent. (Polyphonic C[#] [BCF02] adds the Join Calculus local concurrency primitives to a class-based language.) Other work in the 1990s was also aimed at providing distribution transparency, notably Obliq [Car95], with network-transparent remote object references above Modula3’s network objects.

Distribution transparency, while perhaps desirable in tightly-coupled reliable networks, cannot be provided in systems that are unreliable or span administrative boundaries. Work on Nomadic Pict [SWP99, US01] adopted a lower level of abstraction, showing how a wide variety of distributed infrastructure algorithms, including one similar to that of the JoCaml implementation, could be expressed in a high-level language; one was proved correct. The low level of abstraction means the core language can have a clean and easily-understood failure semantics; the work is a step towards the argument of §2.

A distinct line of work has focussed on typing the entire distributed system to prevent resource access failures, for $D\pi$ [HRY04] and with modal types [MCHP04]. Even where this is possible, however, programmers must still deal with low-level network failure.

Work on Alice [Ali03, Ros03] is perhaps closest to ours, with ML modules, support for marshalling (‘pickling’) arbitrary values, and run-time fresh generation of abstract type names.

Many of the language designs cited above address distributed *execution*, with type-safe interaction within a single program that forks across the network, but there has been little work on distributed *development*, on typed interaction *between* programs⁴, or on version change.

Both Java and .NET have some versioning support, though neither is integrated with the type system. Java serialisation, used in RMI, includes *serialVersionUID*s for classes of any serialised objects. These default to (roughly) hashes of the method names and types, not including the implementation. Class authors can override them with hashes of previous versions. Linking for Java, and in particular binary compatibility, has been studied by Drossopoulou et al. [DEW99]. The .NET framework supports versioning of *assemblies* [dot03]. Sharable assemblies must have *strong names*, which include a public key, file hashes, and a *major.minor.build.revision* version. Compile-time assembly references can be modified before use by XML policy files of the application, code publisher, and machine administrator; the semantics is complex.

Explicit versioning is common in package management, however. For example, both RedHat and Debian packages can contain version constraints on their dependencies, with numeric inequalities and capability-set membership. ELF shared objects express certain version constraints using pathname and symlink conventions. Vesta [ves] provides a rich configuration language.

As discussed in §3 Acute addresses the case in which complex values must be communicated and the interacting runtimes are not malicious. Much other work applies to the untrusted case, with various forms of proof-carrying code and wire-format XML typing which we cannot discuss here.

13 Conclusions and future work

We have addressed key issues in the design of high-level programming languages for distributed computation, discussing the language design space and presenting the Acute language. Acute is a synthesis of an OCaml core with several novel features: dynamic rebinding, global fresh and hash-based type and term naming, versions, type- and abstraction-safe marshalling, etc. It is an experimental language, not a proposal for a full production language, but (as demonstrated by our examples) it shows much of what is needed for higher-order typed distributed computation.

The new constructs should also admit an efficient implementation. The two main points are the tracking of runtime type information, and the implementation of redex-time reduction and rebinding. For the first, note that an implementation does not need to have types for all runtime values, but only (hashes of) the types that reach marshal and unmarshal points. The second would be a smooth extension of OCaml's existing CBV implementation: OCaml currently maintains each field reference $M.x$ as a pointer until it is in redex position, when it is then dereferenced. Since field references inside a thunk remain as pointers, they could easily be rebound with only modest changes to the run-time. Of course compile-time inlining optimisations between parts of code separated by a mark would no longer be possible.

A great deal of future work remains. In the short term, more practical experience in programming in Acute is needed, and there are unresolved semantic issues in the interaction between explicit polymorphism, coloured brackets, and marshalling. Straightforward extensions would ease programming: user definable type operators and recursive datatypes, first-order functors, and richer version languages. A more efficient implementation runtime may be needed for larger examples. Improved tool support for the semantics would be of great value, for meta-typechecking, for conformance testing, and for proofs of soundness.

More fundamentally:

- We must study more refined low-level linking, for negotiation and for access control (escaping the linear mark/module structure). This may demand recursive modules.
- The Acute operational semantics is rather complex, as is the definition of compilation. In part this seems inevitable — the semantics deals with dynamic linking, marshalling, concurrency, thunkify, and coloured brackets, all of which are dynamically intricate (and few of which are covered by existing large-scale definitions). Additionally, our focus has been on a direct semantics of the user language, rather than a combination of a simpler core and a translation, and Acute has evolved through several phases. It should be possible to make the compilation semantics less algorithmic by appealing explicitly to type canonicalisation. The operational semantics for

⁴Several, including JoCaml and Nomadic Pict, have ad-hoc 'traders' for establishing initial connections between programs.

a language with lower-level linking might well be simpler than that presented here, factoring out the algorithmic issues of *resolvespecs*, for example.

- Subtyping is needed for many version-change scenarios, perhaps with corresponding subhash relations. As mentioned in §10, the proper integration of this with polymorphism is challenging, as is the question of what subtype information needs to be propagated at run-time.
- The Acute constructs for local concurrency are very low level, and it is unclear what should be added. Join patterns, CML-style events, π -style channels, and explicit automata; all are useful idioms.
- Some distributed abstractions, such as libraries of distributed references with distributed garbage collection, may challenge the type system.
- The constructs we have presented should be integrated with support for untrusted interaction.

A combination of what has been presented in Acute with solutions to these problems would support a wide range of distributed programming well.

Acknowledgements We acknowledge a Royal Society University Research Fellowship (Sewell), a St Catharine's College Heller Research Fellowship (Wansbrough), an INRIA grant "post-doc à l'étranger" (Zappa Nardelli), EPSRC grant GRN24872, EC FET-GC project IST-2001-33234 PEPITO, and APPSEM 2. We thank Vilhelm Sjöberg and Christian Steinrücken for their implementation work on Acute's predecessors and Gilles Peskine for discussions on references and coloured brackets. Andrew Appel, Matthew Fairbairn, Jean-Jacques Lévy, Luc Maranget, Gilles Peskine, and Alisdair Wren provided comments.

Part II

Semantics

14 Semantics overview

The Acute definition, given in §16, describes syntax, typing, typed desugaring, errors from compilation and execution, compilation, and operational semantics. It also states type preservation and progress conjectures and gives semantic descriptions of two optimisations: for closures and for removing ‘vacuous brackets’. This section outlines the main points of the semantics. It should be read in conjunction with the definition.

The definition involves several related languages:

1. The *concrete source* language is the language that programmers type, e.g. `function (x,y) -> x + y + M.z`. This is concrete — a set of character sequences.
2. The *sugared source internal* language is generated by parsing, scope resolution and type inference; for example `function (x : int, y : int) -> (+) x ((+) y M_M.z)`. This is an abstract grammar, up to alpha equivalence. The x , y and M are internal identifiers, subject to alpha equivalence; the z and M are external identifiers, which are not. (In fact operators are eta-expanded to ensure they are fully applied.)
3. The *source internal* language is generated by desugaring, for example `function (u : int * int) -> match u with ((x : int), (y : int)) -> (+) x ((+) y M_M.z)`.
4. The *compiled* language is generated by compilation, which here computes global type names for hashed abstract types, carries out *withspec* and *likespec* checks, etc. The operational semantics is defined over elements of the compiled language.

Note that the compiled language contains both compiled form and source internal form components. Specifically, a compiled program consists of compiled form definitions and/or source internal form **module fresh** definitions, and an optional compiled form expression.

The main definition is of the union of the grammars for the *source internal* and *compiled* languages.

14.1 Naming

The language makes heavy use of names: at the expression level (names for communication channels, RPC handles etc.), at the type level (for abstract type names), and at the module level (names associated with modules and with imports are used both to construct abstract type names and in version constraints and version expressions).

Names, of each variety, can be generated either from module or import hashes (deterministically), or by taking (psuedo-)random numbers, at either compile-time or run-time. In an implementation these names will all be represented uniformly, e.g. as 160-bit numbers.

Both hash-generation and random-generation allow names to be safely associated with type information across the global distributed system. If one wishes to establish a shared name (expression or type) across programs, it can either be hash-generated from shared source of a module or be compile-time fresh generated and the resulting `.aco` file included by both programs. Other names, on the other hand, must be run-time generated (for names of dynamically-created channels, and of generative types that depend on computational effects).

Using hashes and random name generation means that the correct operation of programs is only probabilistically guaranteed. The name representation must be chosen to be of enough bits to make the probability of accidental collision acceptably low (e.g. lower than the rate of hardware or cosmic-ray errors).

While a production implementation would represent names purely as 160-bit numbers, in order to define typability for states reachable by computation more structure is required. The semantics is therefore expressed in terms of *structured hashes* `hash(...)` and *abstract names* `n`; the metavariable h ranges over both. Structured hashes, e.g.

$\text{hash}(\text{hmodule}_{eqs} M : \text{Sig}_0 \text{ version } vne = \text{Str})$, are formal representations of hash values that preserve all their internal structure. For abstract names n , which have no internal structure, the semantics maintains a global type environment E_n mapping all those that have been created so far to their respective module, kind or type data. Our prototype implementation can work either with structured hashes (and maintain an E_n) or with literal numeric hashes (and discard the E_n). The former allows optional run-time typechecking, of the entire configuration on a machine after every reduction step, which is a valuable tool for debugging both the language definition and the implementation. It also allows the less costly option of unmarshal-time and resolve-time typechecking.

We do not work up to alpha-equivalence of the global abstract names in E_n , instead choosing fresh names non-deterministically from those that have not been used so far. Our E_n really is a global environment, affected (in the semantics but not the implementation) by *all* running machines. This contrasts with the usual π -calculus approach of extruding binders as necessary. We make this choice to avoid having to consider alpha-equivalence of marshalled values, which are simply byte-strings, but which can be unmarshalled to values containing names.

A further subtlety arises in the version expression and constraint languages. Here it is desirable to let the programmer paste in literal hashes, and there is no way for the language to ensure that these literals all arise as the hashes of well-formed modules.

14.2 Typing

(§16.3, page 91) Much of the type system is standard, using singleton kinds to express manifest and abstract types in modules [HL94, Ler94], and with a subsignature relation based on the subkind relation $\text{EQ}(T) <: \text{TYPE}$ allowing manifest type information to be forgotten.

In contrast to most previous work on abstract types and module systems the semantics constructs global names for abstract types, at compile-time or run-time, instead of erasing all types or substituting abstractions away. A source internal language type $M_M.t$ (the t type field of module M_M) is compiled or reduced to a global type name $h.t$, where h is a hash or fresh abstract name. This is a dynamic analogue of the type-theoretic *selfification* rules in singleton-kind systems.

To establish greater confidence in the internal coherence of the semantics we preserve abstraction boundaries throughout execution, adapting and extending *coloured brackets* [GMZ00, LPSW03] to delimit subexpressions in which sets eqs of type equalities $h.t \approx T$ between abstract types $h.t$ and their representations can be used. Additionally, most type judgements, and the operational relations, are with respect to such sets of equalities eqs , reflecting which abstractions one is within. To type a coloured bracket $[e]_{eqs'}^T$, with type T and in an ambient colour eqs , one must have e of type T in colour eqs' .

The most interesting typing rules are for modules, imports, and hashes and abstract names. These latter two behave very like module identifiers, with rules for selfification and for constructing types $h.t$ and terms $h.x$ (the latter occurring only within other hashes, not in executable code). The type rules for the compiled language check that such h are used correctly, referring to their internal structure or the global type environment E_n respectively for hashes or abstract names. An implementation does not need this information, however — in particular, it is not required for the unmarshal-time type equality check.

Typing source internal language modules and imports is much as one would expect. Typing their compiled forms is more interesting, capturing a number of properties that are established by compilation.

Marshalling and unmarshalling are straightforward as far as their static typing goes, converting between arbitrary T and string.

In any given environment E and colour eqs , each semantic type may be represented by any member of an equivalence class of syntactic types defined by the relation $E \vdash_{eqs} T \approx T'$. Compilation ensures that the syntactic type chosen is always the *canonical type* from the relevant equivalence class. The canonical type is the one that is most concrete: it is the normal form under the rewrites $\{X.t \rightsquigarrow T \mid (X.t \approx T) \in eqs\}, M.t \rightsquigarrow T \mid M : \text{Sig} \in E \wedge t : \text{EQ}(T) \in \text{Sig}, \text{ and } t \rightsquigarrow T \mid t : \text{EQ}(T) \in E$. This is important because in certain circumstances the syntactic representative chosen for a semantic type is significant.

14.3 Compilation

(§16.7, page 115) Compilation involves several activities (which are recursively intertwined in the definition):

- preprocessing (i.e., replacing `includesource sourcefilename` and `includecompiled compiledfilename` by the file bodies)
- desugaring
- type-checking
- traversing module definitions calculating (and using) the names to use for global type name of abstract type
- calculating fresh names for `cfresh` modules and expressions
- checking asserted *withspec* equations are correct and that any module linking is legitimate. These are not checked by the type system as they depend on knowledge of the representation types of earlier abstract types, which is not recorded in type environments.

Formally, compilation is a relation from a name environment E_n , a *sourcefilename*, and a filesystem Φ to either a tuple of a source type environment E_0 , a compiled type environment E_1 , and a *compiledunit*, or an error.

Note that *compiledunit* includes a name environment E_n : this environment contains `cfresh` names created during compilation. This name environment has no implementation significance: its sole purpose is to allow included compiled units to be appropriately typechecked and the configuration produced by compilation to be typechecked. These two checks are both necessary for runtime typechecking, but not otherwise.

Note that compilation is not a function because the choice of name environment in the *compiledunit* is nondeterministic. This nondeterminism is common in many of the helper “functions” throughout, thus we take them all to be relations. For convenience, though, we write them as functions of their inputs, and use \rightsquigarrow rather than $=$ to relate the “input arguments” to the “results”.

Compilation has the form

$$\text{compile}_{\Phi}(\text{sourcefilename})E_n \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$$

defined to be

$$\text{compile}_{\Phi}^{\text{empty}} E_n E_{\text{const}} E_{\text{const}}(\text{includesource sourcefilename} \ ; \ ; \ \text{empty}) \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$$

where the latter relation

$$\text{compile}_{\Phi}^{\text{definitions}} E_n E_0 E_1(\text{compilationunit}) \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$$

is defined inductively on the *compilationunit*. Here *sourcefilenames* is the filenames we’ve been through (used to detect cyclic includes), *definitions* is the accumulated compiled definitions, E_n is the accumulated name environment (all names created during compilation will be disjoint from $\text{dom}(E_n)$), E_0 is the accumulated source type environment (including E_{const} at the start), E_1 is the accumulated compiled type environment (including E_{const} at the start), and *compilationunit* is what we have left to do.

The behaviour of compilation on a module (or import, similarly) depends on whether it is annotated `hash`, `cfresh` or `fresh` (which will generally depend on whether it is valuable, *c*valuable or non-valuable). We first describe the `hash` case, with steps corresponding to those in §16.7. First and secondly, all types are normalised as far as possible, replacing any types $M'_{M'.t}$ defined in earlier modules by either the corresponding $h'.t$ (if they are abstract) or the corresponding T (if they are manifest). References to earlier type fields in this module are also flattened where possible: in the structure all type definitions are substituted away; in the signature only manifest type fields can be substituted away. Thirdly, any *withspec* is checked, and the resulting set of type equations, normalised, is recorded. Fourthly, the hash of this module can be constructed, first replacing any other-module expression dependencies $M'_{M'.x}$ by the corresponding $h'.x$. Fifthly, that hash is used to selfify the remaining abstract type fields of the signature,

replacing **type** $t_t : \text{TYPE}$ by **type** $t_t : \text{EQ}(h.t)$. Sixthly, the version number expression of the module is evaluated, replacing **myhash** by the hash h . The result has the form

$$\mathbf{cmodule}_{h; eqs; Sig_0} M_M : Sig_1 \mathbf{version} \quad vn = Str$$

where h is this module's hash, eqs are any extra equations added by the *withspec*, Sig_0 is the normalised but non-selfified signature, Sig_1 is the normalised and selfified signature, vn is the version number, and Str is the normalised structure.

The body of a hash thus does not exactly match either the source module or the compiled module. It cannot be the source module as it must be type-normalised, so that hash equality respects provable type equality. It cannot be the compiled module as that would require recursive hashes — the selfification during compilation uses the hash. (One could introduce a formal recursive hash, but it seems more intuitive not too.)

Compilation of a hash-import is broadly similar, with a *likespec* rather than a *withspec*, resulting in a form

$$\mathbf{cimport}_{h; Sig_0} M_M : Sig_1 \mathbf{version} \quad vc \mathbf{like} \quad Str \mathbf{by} \quad resolvespec = Mo$$

In the **cfresh** cases compilation constructs an h for the module nondeterministically instead of by hashing, taking any n not in the domain of the ambient E_n . Expression-level **cfresh** names are constructed similarly, and compilation is otherwise similar.

In the **fresh** case the h for the module is constructed nondeterministically at the start of its execution, whereupon it can be used to selfify and normalise types similarly.

14.4 Operational judgements

(§16.8.1, 16.8.2, and 16.8.3, page 125) The runtime configurations of a single machine have the form

$$\langle E_s, s, definitions, P \rangle$$

where E_s is the store typing (not required in a production implementation), s is the store, $definitions$ is the sequence of module definitions (all of which are definition values), and P is a multiset of named running threads, mutexes, and condition variables.

$$\begin{aligned} P & ::= 0 \\ & \quad P_1 \mid P_2 \\ & \quad \mathbf{n} : definitions \ e \\ & \quad \mathbf{n} : \mathbf{MX}(\underline{b}) \\ & \quad \mathbf{n} : \mathbf{CV} \end{aligned}$$

The main operational judgements are as below. The first two of these are the main judgements; the other four are auxiliaries introduced so that most reduction axioms need only mention the relevant parts of a configuration. We sometimes call a tuple $\langle E_s, s, definitions, e \rangle$ a *pseudo-configuration*.

- $E_n ; \langle E_s, s, definitions, P \rangle \xrightarrow{\mathbf{n}:\ell} E_n' ; \langle E'_s, s', definitions', P' \rangle$ Process reduction.
- $E_n ; \langle E_s, s, definitions, P \rangle \rightarrow \mathbf{TERM}$ Program termination.
- $E_n ; \langle E_s, s, definitions, e \rangle \xrightarrow{\ell}_{eqs} E_n' ; \langle E'_s, s', definitions', e' \rangle$ Expression reduction.
- $e \xrightarrow{\ell}_{eqs} e'$
- $E_n ; e \xrightarrow{\ell}_{eqs} E_n' ; e'$

- $P \xrightarrow{\ell} P'$

where

$\ell ::=$	empty	internal reduction step
	$x^n v_1^\emptyset \dots v_n^\emptyset$ for $x^n \in \text{dom}(E_{\text{const}}) \wedge \text{os}(x^n)$	invocation of OS call
	$\text{Ok}(v^\emptyset)$	return from OS call
	$\text{Ex}(v^\emptyset)$	return from OS call
	$\text{GetURI}(URI)$	request for code at URI
	$\text{DeliverURI}(\text{definitions})$	resulting code
	CannotFindURI	nothing found at URI

We write $\xrightarrow{\text{empty}}$ simply as \rightarrow .

The class of values is parameterised by colours, with v^{eqs} ranging over the values at colour eqs . The dynamic semantics is expressed with evaluation contexts as follows.

- C_{eqs} is a single-level evaluation context at colour eqs . These are largely standard, for example $_ e$ and $v^{eqs} _$.
- $C_{eqs_2}^{eqs_1}$ is a colour-changing single-level evaluation context, at colour eqs_1 but with a hole at colour eqs_2 . The main case of these is the coloured brackets, $[_]_{eqs_2}^T$, but there are several cases where we need to construct a value at colour \emptyset , e.g. to store or to pass to a primitive operator, so this grammar includes e.g. $l :='_T _$ for $eqs_2 = \emptyset$.
- CC_{eqs} and $CC_{eqs_2}^{eqs_1}$ are multi-level evaluation contexts — simple compositions of the above.

$$CC_{eqs} ::= _ \quad CC_{eqs_2}^{eqs_1} ::= _$$

$$CC_{eqs} ::= _ \quad CC_{eqs_2}^{eqs_1} ::= _$$

- SC_{eqs} is a structure evaluation context, allowing computation in the first non-value expression field of a structure.
- TC_{eqs} is a thread evaluation context. For a thread with body just a single expression e , computation can take place there; for a thread with a body $\text{definitions } e$ where the head of definitions is a **cmodule**, computation can take place in the first non-value expression field of the structure.
- TCC_{eqs} is a composition $TC_{eqs_2} \cdot CC_{eqs_2}^{eqs_1}$, allowing computation within the expression in a TC_{eqs} hole.

14.5 Colours and bracket dynamics

The semantics preserves abstraction boundaries boundaries, generalising the *coloured brackets* of Grossman et al [GMZ00]. (At present this covers the entire Acute language except the System F polymorphism constructs.)

Coloured brackets make explicit the type equalities which are in scope for any subexpression. There is a bracket expression form

$$[e]_{eqs}^T$$

for type equations

$$eqs ::= \emptyset \mid M_M.t \approx T \mid h.t \approx T \mid eqs, eqs$$

giving the representation types of abstract types (source-language projections from a module identifier $M_M.t$ and compiled-language projections from a module name $h.t$). From the outside $[e]_{eqs}^T$ is of type T ; the type equations eqs can be used in typechecking e . We use “colour” and “type equations” interchangeably.

Brackets are not needed in a production implementation (our implementation can work with them or without them), and they are not strictly speaking necessary for the semantics — with the exception of the work of Grossman et al, and

of our previous [LPSW03] and Rossberg’s [Ros03], most operational semantics for existential types and for module systems forgets abstraction boundaries as it comes to them, e.g. with this rule for opening an existential package

$$\text{let } \{t, x\} = (\{T, e\} \text{ as } T') \text{ in } e_2 \rightarrow \{T/t, e/x\}e_2$$

or by substituting out module definitions. Maintaining abstraction boundaries requires some complexity in the semantics, but we think it well worth while. Type preservation for an abstraction-preserving semantics is intuitively a much stronger property than for a standard semantics, and so a better check of internal consistency; and making type equations explicit in both the type system and runtime provides conceptual clarity.

Brackets are not a user source language construct. They are introduced primarily when instantiating a module field reference $M_{M.x}$ from a module M_M that introduced some abstract types (see *Module field instantiation – module case, via import sequence*, §16.8.6, page 136). For a simple example, consider the `EvenCounter` of §5, with fields `start : EvenCounter.t` and `up : EvenCounter.t->EvenCounter.t`. Expressions `EvenCounter.start` and `EvenCounter.up` will be instantiated, when they appear in redex-position, to $[0]_{h.t=int}^{h.t}$ and $[\text{fun } (x:int) \rightarrow 2+x]_{h.t=int}^{h.t \rightarrow h.t}$ respectively. Here h is the hash-generated module name of `EvenCounter` as in §5.

Bracket semantics could be expressed either with a structural congruence or with reductions. We choose the latter, to support our prototype implementation. The basic points are the definition of values (§16.8.2, page 125) and the bracket-pushing reductions of §16.8.4, page 130. The latter push brackets through values in cases where the outermost value structure and the outermost type structure of the bracket type coincide, e.g.

$$[v_1^{eqs'} :: v_2^{eqs'}]_{eqs'}^T \text{ list} \rightarrow_{eqs} [v_1^{eqs'}]_{eqs'}^T :: [v_2^{eqs'}]_{eqs'}^T \text{ list}$$

Bracket type revelation permits use of the ambient type equations to reveal an abstract bracket type, and bracket elimination removes redundant nested brackets.

The semantics must also suitably-bracket expressions used in substitutions to ensure they retain their original type equations. One sees this in the rule for pushing brackets through lambdas and in the reduction axioms for function application and recursive functions.

At several points it is necessary to take a value at some equations eqs and construct a value that makes sense at the empty set of equations \emptyset , e.g. when marshalling a value, passing a value to a primitive operator or an OS call, etc.

The treatment of store locations and names is discussed in §??.

14.6 Marshalling and unmarshalling

(§16.8.5, page 133) A marshalled value is a byte-string representation of an mv , containing data as below.

$$mv ::= \text{marshalled}(E_n, E_s, s, \text{definitions}, e, T) \quad \text{Marshalled value}$$

Here e is the core value being shipped, T its type, s a store, E_s a store typing, definitions is a sequence of module definitions, and E_n is a name environment.

The E_n and E_s would not be shipped in a production implementation, but are needed to state type preservation and for runtime typechecking of reachable states. They are shipped in our implementation only if literal hashes are not being used.

As with the other syntactic objects, marshalled values are taken up to alpha equivalence. Here: the name environment E_n binds in everything to the right and internally contains no cycles; the store environment E_s binds in everything to the right and may contain internal cycles; the store s and the definitions bind to the right and may mutually refer to each other; the s may contain internal cycles.

To characterise the wire format, we simply suppose a fixed partial function `raw_unmarshal` from strings to marshalled values that includes all marshalled values in its range. The semantics for marshalling constructs an mv and then nondeterministically allows any string that is mapped to that mv . This permits small variations in the wire format (which a characterisation in terms of a function from marshalled values to strings would not). We use actual strings for wire-format marshalled values, instead of (say) adding a language type marshalled with elements of the form mv , so that the semantics can capture the behaviour of programs that do string operations — for example, extracting marshalled values from TCP byte streams.

The dynamic semantics for **marshal** $e_1 e_2 : T$ first evaluates e_1 to a string MK — the mark at which module bindings will be cut. It then evaluates e_2 to a value in the ambient colour eqs , and then to a value in the empty colour \emptyset , giving a redex **marshalz** MK $v^\emptyset : T$ in a configuration of the form

$$E_n ; \langle E_s, s, definitions, P \mid n : TCC_{eqs}.\mathbf{marshalz} \text{ MK } v^\emptyset : T \rangle$$

Suppose

$$\begin{aligned} definitions &= definitions_1 ;; \mathbf{mark} \text{ MK} ;; definitions_2 \\ \mathbf{mark} \text{ MK} &\notin definitions_2 \end{aligned}$$

In outline, what we do is prune $definitions_2$, omitting any modules that are not needed by the marshalled value, and on the way calculating which modules from $definitions_1$ are referred to. We then go through $definitions_1$ constructing an import for each of those. The constructed imports either have an exact-name constraint and **HERE_ALREADY** resolvespec, for a cut **cmodule** binding, or with the original version constraint and resolvespec, for a cut **cimport** binding. Note that this does not involve any definitions of executing threads, so the $definitions'$ that are shipped are guaranteed to be definition values. The shipped $definitions'$ includes (copies of) all the marks passed through in $definitions_2$, but not of the **mark** MK being marshalled with respect to. The marshalled value also includes a copy of the reachable part of the store: the value v^\emptyset may contain store locations. They may contain other store locations, but also module identifiers (under lambdas) from $definitions_1$ and $definitions_2$ which must be taken into account. Moreover, as $definitions$ may be the result of module initialisation, it too may contain store locations.

Unmarshalling of a string \underline{s} , in a configuration of the form

$$E_n ; \langle E_s, s, definitions, P \mid TCC_{eqs}.\mathbf{unmarshal} \ \underline{s} \text{ as } T \rangle$$

takes the raw_unmarshal image of \underline{s} , say $\mathbf{marshalled}(E_n', E_{s'}, s', definitions', v'^\emptyset, T')$, adds the store fragment s' to the current store s (disjointly), adds the $definitions'$ to the end of $definitions$ (avoiding clashes with alpha equivalence), and merges E_n with any new names from E_n' . Note this depends on the fact that $definitions'$ are fully evaluated.

Existing marks will thus be shadowed by marks in $definitions'$, which is sometimes desirable but not always. This is a defect of the linear mark/module structure.

14.7 Module field instantiation

(§16.8.6, page 135) This specifies the runtime semantics for resolution of module field references, describing what happens when an $M_M.x$ appears in redex position. In general we have to chase through a (possibly-empty) sequence of linked imports until we arrive at either a module definition or an unlinked import. In the former case we instantiate the $M_M.x$ with its value from the module definition. In the latter, we work through the *resolvespec* attached to the import $M'_{M'}$ that is unlinked. Each atomic resolve spec is dealt with in turn, as follows:

- **STATIC_LINK** – fail, raising an exception;
- **HERE_ALREADY** – look in the preceding modules for one that matches the signature and version constraint. If there is one, link this import to it;
- **URI** – try to load a *compiledunit* from the *URI*. If we find one containing a module that matches the external name, signature and version constraint, and has $eo = \text{empty}$, add it to the configuration's *definitions* just before the import, and link the import to it.

In the latter two cases, if there is a failure we try the subsequent atomic resolve specs, raising an exception if there are no more. Success leaves the $M_M.x$ again in redex position, where it can now be instantiated as in the former case.

Note that no additional linking is done, either to or of newly-loaded modules. Some user control of this would be desirable.

Resolution may involve IO, to pull a file containing compiled definitions from the web or filesystem. The semantics expresses this with labelled transitions $n : \text{GetURI}(URI)$ for making a request for a

URI, \mathbf{n} : DeliverURI($E_n', definitions'$) for receiving a name environment E_n' and $definitions'$, and \mathbf{n} : CannotFindURI if no file is found at the URI . The intermediate state is stored in the term, as a `resolve_blocked`($M_M.x, M'_{M'}, resolvespec$) for the blocked state and a `resolve`($M_M.x, M'_{M'}, resolvespec$) for a state which is just about to make a request. The action must be split into send and receive events as the receive may be blocked arbitrarily, and the semantics must make sense with language threads are added (note that the reduction closure rules add thread ids to the transitions of axioms of this section).

There is a choice as to how we generate coloured brackets when passing through multiple imports: (1) we could make a single *eqs* set containing all the equalities we need; or, (2) we could have a nested sequence of brackets. Choice (1) might require the bracket pushing rules compare *eqs* sets by inclusion (or logical implication). The latter doesn't suffer from this problem, and so we choose (2).

On instantiating an $M_M.x$ via a chain of imports, where M_M is bound by a `cimport` which is linked to a `cimport` which ... is linked to a `cmodule`, the equation set is the union of the *weqs* of that `cmodule`, the equations of the signature/structure boundary of that `cmodule`, and the equaltions of the signature/likespec boundary of the initial `cimport`. The intermediate imports are not relevant.

There is a technical choice relating to the semantics of instantiation, of module initialisation, and of rebinding. For a module with internal expression dependencies, e.g.

```
module M : sig val x:int val f:int->int end =
  struct
    let x = 3
    let f = fun (y:int) -> x + y
  end
M.f 10
```

we can either (1) substitute $\{3/x\}$ through the body of f in the structure at compilation or module-initialisation time (if x were bound to an effect-full computation it would have to be the latter), or (2) leave the body of f with a free occurrence of x . For (1) module field instantiation is straightforward, as when $M.f$ is in redex position (as here) it can be replaced by the expression-identifier-closed fun $(y:int) \rightarrow 3+y$. For (2), instantiation of $M.f$ would have to rewrite the x on the fly, either (a) to $M.x$ or (b) to 3. Option (2a), instantiating the $M.f$ to fun $(y:int) \rightarrow M.x+y$, allows more rebinding than (1) or (2b), as M might be rebound before the $M.x$ itself appears in redex position. If one is instantiating via an import, however, and if width subsignaturing were added to the language, it seems that one could not give a satisfactory semantics for (2a). The rewrite would have to use the module identifier, not the import identifier, and hence rebinding could often lead to link errors — it would not be enough to supply an implementation of the import one was working with as other fields of the module might be required.

14.8 Concurrency

(§16.8.8, page 140) The semantics for thread creation, termination, `self`, and `kill` are technically straightforward, written as reduction axioms for the judgement $P \rightarrow P'$ (but note that as some of the axioms need to check the set of *all* locally-used thread names, these transitions are not closed under parallel composition).

Our configurations keep the states of threads, mutexes and condition variables in a single multiset; each is named with a global name (which might be hash-, fresh- or cfresh-generated). This is notationally smoother than the alternative of having separate configuration components for each.

For mutexes, POSIX describes three semantics (“kinds” of mutex): *fast*, *recursive*, and *error checking*. The *fast* semantics blocks when a thread attempts to lock a locked mutex. Note that this leads to deadlock if a thread attempts to lock a mutex it has already locked. In the LinuxThreads implementation, any thread is allowed to unlock a locked mutex. However: “This is non-portable behaviour and must not be relied upon.” — POSIX is quite clear that it is assumed the *owner* is unlocking the mutex. Unlocking an unlocked mutex has no effect (this is also non-portable: in POSIX it is undefined). The *recursive* semantics maintains the owner and a lock count in the mutex; if a thread locks a mutex it already holds, this succeeds immediately and the count is incremented; on unlock, the count is decremented (again, LinuxThreads non-portably does not check the owner here). (POSIX specifies that unlocking an unlocked mutex should fail with EPERM, but LinuxThreads’ man page suggests that unlocking an unlocked mutex (non-portably) has no effect). The *error checking* semantics is like the fast semantics, except that locking a mutex already held by the calling thread results in an immediate EDEADLK error, and attempting to unlock a mutex not owned

by the caller results in an immediate EPERM error. Unlocking an unlocked mutex results in an immediate EPERM error. POSIX does not specify which semantics is the default. On LinuxThreads the default is *fast*, and OCaml on our Linux install appears by default to have the *fast* POSIX semantics. It is this semantics (approximately) that is expressed in our semantics. We are not committed to this, but it is fine for now.

Application programmers using threads, mutexes and condition variables depend on some fairness properties. The semantics does not express these at present.

The semantics for **thunkify** uses an auxiliary function `Thunkify` to atomically construct a thunk encapsulating the state of the threads, mutexes and condition variables being thunkified. When that thunk is applied (to a `thunkkey` list giving the names at which to reify the various parts) it uses the auxiliary `Unthunkify` to (atomically) build a process and place it in parallel with that of the running configuration.

If a blocking `thunkify` is waiting, it does not at present have a ‘lock’ of any kind on the things it is trying to thunkify, though that might be desirable. Here, there just are no transitions for such a `thunkify`, or indeed a `thunkify` with a thread in a fast system call. Races between overlapping **thunkifys** are thus possible, and the liveness properties even of a single **thunkify** are very weak.

Note that **thunkify** fails when applied to a thread which contains some uninitialised definitions. One could instead have it block until the initialising thread is finished. (As thunks are simply part of the expression language, and Acute modules are second-class, allowing thunkification of module-initialising threads would entail substantial changes to the language — they simply cannot be expressed in the syntax as it stands, and the evaluation order for their usages is problematic.)

15 Semantics Examples

This Section illustrates some aspects of the semantics with examples generated by our implementation. At present it covers just the compilation and marshalled values of the examples earlier in the paper. They are rendered in a typewriter variant of the grammar used in the semantics, close to the concrete source language. In addition, the pretty-printer collects together occurrences of module hashes and abstract names, introducing metavariables h and n . Internal identifiers are rendered with numeric subscripts, e.g. x_0 . Internal identifiers of modules and imports that are not printed are rendered as $:M?:.$

15.1 Compilation: hash modules

The result of compiling module `EvenCounter` from §5, page 20, is below. Scope resolution has introduced internal identifiers M_0 , t_0 , $start_0$, x_0 etc. Compilation has calculated a module name $hO_EvenCounter$ as a hash of an `hmodule` form, containing external module identifier, signature, version expression, and structure. This hash is taken up to alpha equivalence by choosing canonical strings for bound identifiers and up to type equality by substituting out earlier module names for identifiers and substituting out internal type dependencies. (The hash body shown is pretty-printed in a different mode to that used to build the actual hash to make it more readable, with identifiers based on the source language strings.) Both the symbolic and literal hash forms are shown. The compiled `cmodule` `EvenCounter` has two signatures, one in which source abstract types are still abstract and one in which they have been selfified using the module name and substituted through, e.g. the type $t[t_0] : Eq(hO_EvenCounter.t)$ and `val start[start_0] : hO_EvenCounter.t`. The version of the compiled module has defaulted to its hash-generated name.

```
cmodule EvenCounter[M0] hO_EvenCounter : {}
  sig
    type t[t0] : Type
    val start[start0] : t0
    val get[get0] : t0 -> int
    val up[up0] : t0 -> t0
  end (valuable, valuable)
  sig
    type t[t0] : Eq(hO_EvenCounter.t)
    val start[start0] : hO_EvenCounter.t
    val get[get0] : hO_EvenCounter.t -> int
    val up[up0] : hO_EvenCounter.t -> hO_EvenCounter.t
  end
  version hO_EvenCounter
= struct
  type t[t0] = int
  let start[start0] = 0
  let get[get0] = function (x0 : int) -> x0
  let up[up0] = function (x0 : int) -> 2 + x0
end

where
  hO_EvenCounter = hash(hmodule EvenCounter : {}
    sig
      type t[t0] : Type
      val start[start0] : t0
      val get[get0] : t0 -> int
      val up[up0] : t0 -> t0
    end
    version myname
  = struct
    type t[t0] = int
    let start[start0] = 0
```

```

        let get[get0] = function (x0 : int) -> x0
        let up[up0] = function (x0 : int) -> 2 + x0
    end)
= 0#E09083A42C03366FA0698C81E0063682

```

15.2 Compilation: fresh modules

The module `NCounter` from §5, page 21, compiles to:

```

module fresh NCounter [M0]
: sig
  type t[t0] : Type
  val start[start0] : t0
  val get[get0] : t0 -> int
  val up[up0] : t0 -> t0
end
version myname
= struct
  type t[t0] = int
  let start[start0] = 0
  let get[get0] = function (x0 : int) -> x0
  let up[up0] =
    match Pervasives[Lib_Pervasives].read_int () with (step0 : int) ->
      function (x0 : int) -> step0 + x0
end

```

The first execution step of this involves generating a fresh name for the module and hashifying it, after which the `read_int` performs IO.

15.3 Compilation: hash module dependencies

The result of compiling modules `M` and `EvenCounter` from §8, page 28, is below. Two hashes are constructed to use as the names of the two modules, `h0_M` and `h1_EvenCounter`. Note that the `up` field of the `cmodule EvenCounter` structure refers to `M[M0].f x0`, whereas the `up` field of the `hmodule EvenCounter` in the body of its hash refers to `h0_M.f x0`, using the earlier hash.

```

cmodule M[M0] h0_M : {}
  sig
    val f[f0] : int -> int
  end (valuable, valuable)
  sig
    val f[f0] : int -> int
  end
  version h0_M
= struct
  let f[f0] = function (x0 : int) -> x0 + 2
end

```

where

```

h0_M = hash(hmodule M : {}
  sig
    val f[f0] : int -> int
  end)

```

```

        end
        version myname
    = struct
        let f[f0] = function (x0 : int) -> x0 + 2
        end)
    = 0#FBCF6A65CCD4F06635C5188503EA9B72

cmodule EvenCounter[M0] h1_EvenCounter : {}
sig
    type t[t0] : Type
    val start[start0] : t0
    val get[get0] : t0 -> int
    val up[up0] : t0 -> t0
end (valuable, valuable)
sig
    type t[t0] : Eq(h1_EvenCounter.t)
    val start[start0] : h1_EvenCounter.t
    val get[get0] : h1_EvenCounter.t -> int
    val up[up0] : h1_EvenCounter.t -> h1_EvenCounter.t
end
version h1_EvenCounter
= struct
    type t[t0] = int
    let start[start0] = 0
    let get[get0] = function (x0 : int) -> x0
    let up[up0] = function (x0 : int) -> M[M0].f x0
end

where
    h1_EvenCounter = hash(hmodule EvenCounter : {}
        sig
            type t[t0] : Type
            val start[start0] : t0
            val get[get0] : t0 -> int
            val up[up0] : t0 -> t0
        end
        version myname
    = struct
        type t[t0] = int
        let start[start0] = 0
        let get[get0] = function (x0 : int) -> x0
        let up[up0] = function (x0 : int) -> h0_M.f x0
    end)
    = 0#F5EF4DE7D2DCB9E8D56EE8AAD19AE3E9

```

15.4 Compilation: cfresh modules

The cfresh code from Scenario 2, page 23, compiles to:

```

cmodule M[M0] h0 : {}
sig
    val c[c0] : int name
end (cvaluable, cvaluable)
sig

```

```

    val c[c0] : int name
  end
  version h0
= struct
  let c[c0] = name_value(n1 %[int])
  end

where
  h0 = n0 = 0#2D3C130675CA1701BB285B45679B27BD

where
  n0 = 0$2D3C130675CA1701BB285B45679B27BD

  n1 = 0$5C334F890F66E794B27733D88A8228A7 %[int]

```

15.5 Compilation: constructing expression names from module hashes

The result of compiling the shared code from Scenario 3, page 23, is below. Note the hash *h0_N* involves the intension of *N.f*, and this appears within the *c* field of the *cmodule* *M* at the end. (Skip over the intervening hashes of standard library modules *h1_IO*, *h2_Pervasives*, and *h3_Persist*.)

```

cmodule N[M0] h0N : {}
  sig
    val f[f0] : int -> unit
  end (valuable, valuable)
  sig
    val f[f0] : int -> unit
  end
  version h0N
= struct
  let f[f0] = function (x0 : int) -> IO[Lib_IO].print_int (x0 + 100)
  end

where
  h0N = hash(hmodule N : {})
    sig
      val f[f0] : int -> unit
    end
    version myname
  = struct
    let f[f0] = function (x0 : int) -> h1_IO.print_int (x0 + 100)
    end)
  = 0#75ABE6A8126FA4F96A02789EAC83E487

where
  h1_IO = hash(hmodule IO : {})
    sig
      val print_int[print_int0] : int -> unit
      val print_string[print_string0] : string -> unit
      val print_newline[print_newline0] : unit -> unit
      val send[send0] : string -> unit
      val receive[receive0] : unit -> string
    end
    version myname

```

```

= struct
  let print_int[print_int0] = function (x0 : int) -> h2_Pervasives.print_int x0
  let print_string[print_string0] =
    function (s0 : string) -> h2_Pervasives.print_string s0
  let print_newline[print_newline0] =
    function (ds0 : unit) -> match ds0 with () -> h2_Pervasives.print_newline ()
  let send[send0] = function (data0 : string) -> h3_Persist.write data0
  let receive[receive0] =
    function (ds0 : unit) -> match ds0 with () -> h3_Persist.read ()
  end)
= 0#2808905E9138A8AA18FF6FF8E169EDED

```

where

```

h2_Pervasives = hash(hmodule Pervasives : {})
  sig
    val string_of_int[string_of_int0] : int -> string
    val int_of_string[int_of_string0] : string -> int
    val print_string[print_string0] : string -> unit
    val print_int[print_int0] : int -> unit
    val print_endline[print_endline0] : string -> unit
    val print_newline[print_newline0] : unit -> unit
  end
  version myname
= struct
  let string_of_int[string_of_int0] =
    function (ds0 : int) -> %"Apervasives_string_of_int" ds0
  let int_of_string[int_of_string0] =
    function (ds0 : string) -> %"Apervasives_int_of_string" ds0
  let print_string[print_string0] =
    function (ds0 : string) -> %"Apervasives_print_string" ds0
  let print_int[print_int0] =
    function (ds0 : int) -> %"Apervasives_print_int" ds0
  let print_endline[print_endline0] =
    function (ds0 : string) -> %"Apervasives_print_endline" ds0
  let print_newline[print_newline0] =
    function (ds0 : unit) -> %"Apervasives_print_newline" ds0
  end)
= 0#4A5FE3EC8D80DFA70AD367461DD525AA
h3_Persist = hash(hmodule Persist : {})
  sig
    val write[write0] : string -> unit
    val read[read0] : unit -> string
    val write2[write20] : string -> unit
    val read2[read20] : unit -> string
  end
  version myname
= struct
  let write[write0] = function (ds0 : string) -> %"Persist_write" ds0
  let read[read0] = function (ds0 : unit) -> %"Persist_read" ds0
  let write2[write20] = function (ds0 : string) -> %"Persist_write2" ds0
  let read2[read20] = function (ds0 : unit) -> %"Persist_read2" ds0
  end)
= 0#D90A83203B41EF6E6E512B3E5FF54850

```

```

cmodule M[M0] h4_M : {}

```

```

sig
  val c[c0] : int name
end (valuable, valuable)
sig
  val c[c0] : int name
end
version h4_M
= struct
  let c[c0] = hash(int, "", hash(h0_N.f) %[int -> unit]) %[int]
end

where
  h4_M = hash(hmodule M : {}
    sig
      val c[c0] : int name
    end
    version myname
  = struct
    let c[c0] = hash(int, "", hash(h0_N.f) %[int -> unit]) %[int]
    end)
  = 0#2F7112C065BF44899C98205353679AD7

```

15.6 Compilation: type normalisation and marshalling within abstraction boundaries

The result of compilation for the example of marshalling within an abstraction boundary, §8.5, page 32, is below. Note here in the `cmodule struct` that the types at which marshalling and unmarshalling are done, in the `send` and `recv` fields, have both been normalised to `int` from the source-language `t`.

```

cmodule EvenCounter[M0] h0_EvenCounter : {}
sig
  type t[t0] : Type
  val start[start0] : t0
  val get[get0] : t0 -> int
  val up[up0] : t0 -> t0
  val send[send0] : t0 -> unit
  val recv[recv0] : unit -> t0
end (valuable, valuable)
sig
  type t[t0] : Eq(h0_EvenCounter.t)
  val start[start0] : h0_EvenCounter.t
  val get[get0] : h0_EvenCounter.t -> int
  val up[up0] : h0_EvenCounter.t -> h0_EvenCounter.t
  val send[send0] : h0_EvenCounter.t -> unit
  val recv[recv0] : unit -> h0_EvenCounter.t
end
version h0_EvenCounter
= struct
  type t[t0] = int
  let start[start0] = 0
  let get[get0] = function (x0 : int) -> x0
  let up[up0] = function (x0 : int) -> 2 + x0
  let send[send0] = function (x0 : int) -> IO[Lib_IO].send (marshal "StdLib" x0 : int)
  let recv[recv0] =
    function (ds0 : unit) -> match ds0 with () -> (unmarshal (IO[Lib_IO].receive ()) as int)

```

end

where

```
h0_EvenCounter = hash(hmodule EvenCounter : {}
  sig
    type t[t0] : Type
    val start[start0] : t0
    val get[get0] : t0 -> int
    val up[up0] : t0 -> t0
    val send[send0] : t0 -> unit
    val recv[recv0] : unit -> t0
  end
  version myname
  = struct
    type t[t0] = int
    let start[start0] = 0
    let get[get0] = function (x0 : int) -> x0
    let up[up0] = function (x0 : int) -> 2 + x0
    let send[send0] =
      function (x0 : int) -> h1_IO.send (marshal "StdLib" x0 : int)
    let recv[recv0] =
      function (ds0 : unit) ->
        match ds0 with () -> (unmarshal (h1_IO.receive ()) as int)
    end)
  end)
= 0#A896BA1BA88F408A0AEE9744742E2717
```

where

```
h1_IO = hash(hmodule IO : {}
  sig
    val print_int[print_int0] : int -> unit
    val print_string[print_string0] : string -> unit
    val print_newline[print_newline0] : unit -> unit
    val send[send0] : string -> unit
    val receive[receive0] : unit -> string
  end
  version myname
  = struct
    let print_int[print_int0] = function (x0 : int) -> h2_Pervasives.print_int x0
    let print_string[print_string0] =
      function (s0 : string) -> h2_Pervasives.print_string s0
    let print_newline[print_newline0] =
      function (ds0 : unit) -> match ds0 with () -> h2_Pervasives.print_newline ()
    let send[send0] = function (data0 : string) -> h3_Persist.write data0
    let receive[receive0] =
      function (ds0 : unit) -> match ds0 with () -> h3_Persist.read ()
    end)
  end)
= 0#2808905E9138A8AA18FF6FF8E169EDED
```

where

```
h2_Pervasives = hash(hmodule Pervasives : {}
  sig
    val string_of_int[string_of_int0] : int -> string
    val int_of_string[int_of_string0] : string -> int
    val print_string[print_string0] : string -> unit
    val print_int[print_int0] : int -> unit
    val print_endline[print_endline0] : string -> unit
    val print_newline[print_newline0] : unit -> unit
  end)
```



```

    end
    version myname
  = struct
    let string_of_int[string_of_int0] =
      function (ds0 : int) -> %"Apervasives_string_of_int" ds0
    let int_of_string[int_of_string0] =
      function (ds0 : string) -> %"Apervasives_int_of_string" ds0
    let print_string[print_string0] =
      function (ds0 : string) -> %"Apervasives_print_string" ds0
    let print_int[print_int0] =
      function (ds0 : int) -> %"Apervasives_print_int" ds0
    let print_endline[print_endline0] =
      function (ds0 : string) -> %"Apervasives_print_endline" ds0
    let print_newline[print_newline0] =
      function (ds0 : unit) -> %"Apervasives_print_newline" ds0
  end)
  = 0#4A5FE3EC8D80DFA70AD367461DD525AA
  h3_Persist = hash(hmodule Persist : {})
    sig
      val write[write0] : string -> unit
      val read[read0] : unit -> string
      val write2[write20] : string -> unit
      val read2[read20] : unit -> string
    end
    version myname
  = struct
    let write[write0] = function (ds0 : string) -> %"Persist_write" ds0
    let read[read0] = function (ds0 : unit) -> %"Persist_read" ds0
    let write2[write20] = function (ds0 : string) -> %"Persist_write2" ds0
    let read2[read20] = function (ds0 : unit) -> %"Persist_read2" ds0
  end)
  = 0#D90A83203B41EF6E6E512B3E5FF54850

```

```
EvenCounter[M0].send EvenCounter[M0].start
```

15.7 Compilation: imports

The result of compiling the `M` and `EvenCounter` import example, §8, page 28, is below. Note here that the `cmodule` `M` and the `cimport` `M` have quite different names, the hashes `h0_M` and `h1_M` respectively. It is the latter that appears in the hash `h2_EvenCounter` of the `EvenCounter` module, and that thus would appear in the runtime type names of any source-language `EvenCounter.t` types (there are no such occurrences in this example).

```

cmodule M[M0] h0_M : {}
  sig
    val f[f0] : int -> int
  end (valuable, valuable)
  sig
    val f[f0] : int -> int
  end
  version h0_M
= struct
  let f[f0] = function (x0 : int) -> x0 + 2
end

```

```

where
  h0_M = hash(hmodule M : {}
    sig
      val f[f0] : int -> int
    end
    version myname
  = struct
    let f[f0] = function (x0 : int) -> x0 + 2
    end)
  = 0#FBCF6A65CCD4F06635C5188503EA9B72

cimport M[M1] h1_M
: sig
  val f[f0] : int -> int
end (valuable, valuable)
sig
  val f[f0] : int -> int
end
version *
like struct      end
by Here_Already
= M[M0]

where
  h1_M = hash(himport M: sig   val f[f0] : int -> int   end   version *   like   struct      end)
  = 0#BD28AD1B690255427DBA10F9471C765B

mark "MK"
cmodule EvenCounter[M0] h2_EvenCounter : {}
  sig
    type t[t0] : Type
    val start[start0] : t0
    val get[get0] : t0 -> int
    val up[up0] : t0 -> t0
  end (valuable, valuable)
  sig
    type t[t0] : Eq(h2_EvenCounter.t)
    val start[start0] : h2_EvenCounter.t
    val get[get0] : h2_EvenCounter.t -> int
    val up[up0] : h2_EvenCounter.t -> h2_EvenCounter.t
  end
  version h2_EvenCounter
= struct
  type t[t0] = int
  let start[start0] = 0
  let get[get0] = function (x0 : int) -> x0
  let up[up0] = function (x0 : int) -> M[M1].f x0
end

where
  h2_EvenCounter = hash(hmodule EvenCounter : {}
    sig
      type t[t0] : Type
      val start[start0] : t0
      val get[get0] : t0 -> int
      val up[up0] : t0 -> t0
    end

```

```

        version myname
    = struct
        type t[t0] = int
        let start[start0] = 0
        let get[get0] = function (x0 : int) -> x0
        let up[up0] = function (x0 : int) -> h1_M.f x0
    end)
    = 0#A60A0BC55D9A1B0F753ED1FA69475D83

IO[Lib_IO].send
(marshal "MK"
 (function (ds0 : unit) ->
   match ds0 with () -> EvenCounter[M0].get (EvenCounter[M0].up EvenCounter[M0].start))
 : unit -> int)

```

15.8 Compilation: imports with abstract type fields

Here we show a fleshed-out version of the last two examples of §8.1, page 28. Consider the import below, which has a non-exact-name version and has a signature containing an abstract type field. It has a *likespec* like `M` specifying that the representation type for that type must be the same as that of the preceding module `M` (one could equivalently give the *likespec* explicitly, writing like `struct type t=int end`). The import is initially linked to `M`.

```

module M : sig   type t   val x:t end
  version 2.4.9
  = struct type t=int let x=17 end

import M : sig   type t   val x:t end
  version 2.4.7-
  like M
  = M

mark "MK"
(marshal "MK" M.x : M.t)

```

The result of compiling this code is below. Note that the *likespec* data appears in the import hash `h1_M`, which is used to form the type `h1_M.t` at which the final `marshal` is done. Type errors caused by rebinding the import to modules with different representation types are thus excluded.

```

cmodule M[M0] h0_M : {}
  sig
    type t[t0] : Type
    val x[x0] : t0
  end (valuable, valuable)
  sig
    type t[t0] : Eq(h0_M.t)
    val x[x0] : h0_M.t
  end
  version 2.4.9
= struct
  type t[t0] = int
  let x[x0] = 17
end

where
  h0_M = hash(hmodule M : {})
        sig

```

```

        type t[t0] : Type
        val x[x0] : t0
      end
      version 2.4.9
    = struct
      type t[t0] = int
      let x[x0] = 17
      end)
    = 0#D17E216FA11DBB5BDDDB0B020646900A3

cimport M[M1] h1_M
: sig
  type t[t0] : Type
  val x[x0] : t0
end (valuable, valuable)
sig
  type t[t0] : Eq(h1_M.t)
  val x[x0] : h1_M.t
end
version 2.4.7-
like struct type t[t0] = int end
by Here_Already
= M[M0]

where
  h1_M = hash(himport M
    : sig
      type t[t0] : Type
      val x[x0] : t0
    end
    version 2.4.7-
    like struct type t[t0] = int end)
  = 0#FE46E0350E1A6EAFB00547C6E836B6CB

mark "MK"
(marshal "MK" M[M1].x : h1_M.t)

```

15.9 Compilation: breaking abstractions

The result of compiling the `with!` example of §8.2, page 29, is below. Here `h1_EvenCounter` is the hash of the original module and `h0_EvenCounter` is the hash of the new version with a down operation. The type equation `{h1_EvenCounter.t=int}` is recorded in the `cmodule`.

```

cmodule EvenCounter[M0] h0_EvenCounter : {h1_EvenCounter.t=int}
  sig
    type t[t0] : Eq(h1_EvenCounter.t)
    val start[start0] : h1_EvenCounter.t
    val get[get0] : h1_EvenCounter.t -> int
    val up[up0] : h1_EvenCounter.t -> h1_EvenCounter.t
    val down[down0] : h1_EvenCounter.t -> h1_EvenCounter.t
  end (valuable, valuable)
  sig
    type t[t0] : Eq(h1_EvenCounter.t)
    val start[start0] : h1_EvenCounter.t
    val get[get0] : h1_EvenCounter.t -> int
    val up[up0] : h1_EvenCounter.t -> h1_EvenCounter.t
  end

```

```

    val down[down0] : h1_EvenCounter.t -> h1_EvenCounter.t
  end
  version h0_EvenCounter
= struct
  type t[t0] = int
  let start[start0] = 0
  let get[get0] = function (x0 : int) -> x0
  let up[up0] = function (x0 : int) -> 2 + x0
  let down[down0] = function (x0 : int) -> x0 - 2
end

where
  h0_EvenCounter = hash(hmodule EvenCounter : {h1_EvenCounter.t=int})
    sig
      type t[t0] : Eq(h1_EvenCounter.t)
      val start[start0] : h1_EvenCounter.t
      val get[get0] : h1_EvenCounter.t -> int
      val up[up0] : h1_EvenCounter.t -> h1_EvenCounter.t
      val down[down0] : h1_EvenCounter.t -> h1_EvenCounter.t
    end
    version myname
  = struct
    type t[t0] = int
    let start[start0] = 0
    let get[get0] = function (x0 : int) -> x0
    let up[up0] = function (x0 : int) -> 2 + x0
    let down[down0] = function (x0 : int) -> x0 - 2
  end)
  = 0#E5E0448DECB46AC6F6E22081B274831D

```

```

where
  h1_EvenCounter = hash(hmodule EvenCounter : {})
    sig
      type t[t0] : Type
      val start[start0] : t0
      val get[get0] : t0 -> int
      val up[up0] : t0 -> t0
    end
    version myname
  = struct
    type t[t0] = int
    let start[start0] = 0
    let get[get0] = function (x0 : int) -> x0
    let up[up0] = function (x0 : int) -> 2 + x0
  end)
  = 0#E09083A42C03366FA0698C81E0063682

```

15.10 Marshalled values

In these examples the `-hack_optimise` option of the implementation is used to suppress most vacuous coloured brackets, as described in §16.11.

The marshalled value of the first example of §3, page 13, is below. It contains simply a value and a type.

```
marshalled ({ }, { }, {}, {}, 5, int)
```

The marshalled value of the first example of §4.2, page 15, is below. Here the module M is shipped together with a function that refers to it.

```
marshalled (  
  { },  
  {cmodule M[M0] hO_M : {}  
    sig  
      val y[x] : int  
    end (valuable, valuable)  
    sig  
      val y[x] : int  
    end  
    version hO_M  
  = struct  
    let y[x] = 6  
  end  
  
  }, {},  
  {},  
  
  (function (x : unit) -> match x with () -> M[M0].y),  
  unit -> int)
```

The marshalled value of the second example of §4.2, page 15, is below. This includes an import for M1 and the module for M2, and a function that refers to both. The former is automatically generated for the module binding of M1 that is cut by the mark. It is constructed with an exact-name version constraint, here to the hash-generated name *hO_M1* of M1. The *likespec* of the import is also constructed based on the original module, though here that had no abstract types so the resulting *likespec* is empty.

```
marshalled (  
  { },  
  {cimport M1[M0] hO_M1  
    : sig  
      val y[x] : int  
    end (valuable, valuable)  
    sig  
      val y[x] : int  
    end  
    version name = hO_M1  
    like struct      end  
    by Here_Already  
    = unlinked  
  cmodule M2[M0] h1_M2 : {}  
    sig  
      val z[x] : int  
    end (valuable, valuable)  
    sig  
      val z[x] : int  
    end  
    version h1_M2  
  = struct  
    let z[x] = 3  
  end  
  
  },  
  (function (x : unit) -> match x with () -> M1[M0].y),  
  unit -> int)
```

```
}, {},  
{},
```

```
(function (x : unit) -> match x with () -> (M1[M0].y, M2[M0].z)),  
unit -> int * int)
```

The marshalled value of the third example of §4.2, page 16, is below. Here the marshalled import is essentially that supplied by the user above the mark (hence, that of the binding that is cut by the mark), not an automatically-generated default import.

```
marshalled (  
  { },  
  {cimport M1[M0] h0_M1  
   : sig  
     val y[x] : int  
     end (valuable, valuable)  
   sig  
     val y[x] : int  
     end  
     version name = h0_M1  
     like struct      end  
     by Here_Already  
     = unlinked  
   cimport M1[M1] h1_M1  
   : sig  
     val y[x] : int  
     end (valuable, valuable)  
   sig  
     val y[x] : int  
     end  
     version *  
     like struct      end  
     by Here_Already  
     = unlinked  
   cmodule M2[M0] h2_M2 : {}  
     sig  
       val z[x] : int  
       end (valuable, valuable)  
     sig  
       val z[x] : int  
       end  
     version h2_M2  
   = struct  
     let z[x] = 3  
     end  
  }, {}  
  {}  
)
```

```
(function (x : unit) -> match x with () -> (M1[M1].y, M2[M0].z)),  
unit -> int * int)
```

The marshalled value of the first example of §4.3, page 16, is below. Here one can see that the `M.y` under the `fun` in the source language has not been instantiated (and an import is shipped, binding that `M`) whereas the unguarded `M.y` has been instantiated by its value 6 before marshalling took place.

```

marshalled (
  { },
  {cimport M[M0] h0_M
   : sig
     val y[x] : int
     end (valuable, valuable)
   sig
     val y[x] : int
     end
   version name = h0_M
   like struct     end
   by Here_Already
   = unlinked
  cimport M[M1] h1_M
  : sig
     val y[x] : int
     end (valuable, valuable)
   sig
     val y[x] : int
     end
   version *
   like struct     end
   by Here_Already
   = unlinked

}, {},
{}),

(6, (function (x : unit) -> match x with () -> M[M1].y)),
int * (unit -> int))

```

The marshalled value of the example of §4.5, page 17, is below, with just an import being sent.

```

marshalled (
  { },
  {cimport M[M0] h0_M
   : sig
     val y[x] : int
     end (valuable, valuable)
   sig
     val y[x] : int
     end
   version name = h0_M
   like struct     end
   by Here_Already
   = unlinked
  cimport M[M1] h1_M
  : sig
     val y[x] : int
     end (valuable, valuable)
   sig
     val y[x] : int
     end
   version *
   like struct     end
   by Here_Already
   = unlinked

```



```
}, {},  
{},
```

```
(function (x : unit) -> match x with () -> M[M1].y),  
unit -> int)
```

The marshalled value of the example of §4.9, page 19, is below, with a store fragment mapping a single location to the value 5 and a store typing associating that location with type `int`. The expression part is just that location.

```
marshalled ({ }, { }, {<1> : int ref}, {<1> := 5}, <1>, int ref)
```

The marshalled value of the `thunkify` example of §9.11, page 37, is below. The body is a function which takes a `thunkkey` list containing a thread name and a mutex name and reconstructs the original thread and mutex state at those names.

```
marshalled (  
  { },  
  {cimport Pervasives[Lib_Pervasives] h0_Pervasives  
   : sig  
     val string_of_int[x] : int -> string  
     val int_of_string[x0] : string -> int  
     val print_string[x1] : string -> unit  
     val print_int[x2] : int -> unit  
     val print_endline[x3] : string -> unit  
     val print_newline[x4] : unit -> unit  
   end (valuable, valuable)  
  sig  
     val string_of_int[x] : int -> string  
     val int_of_string[x0] : string -> int  
     val print_string[x1] : string -> unit  
     val print_int[x2] : int -> unit  
     val print_endline[x3] : string -> unit  
     val print_newline[x4] : unit -> unit  
   end  
  version name = h0_Pervasives  
  like struct      end  
  by Here_Already  
  = unlinked  
  cimport Persist[Lib_Persist] h1_Persist  
  : sig  
     val write[x] : string -> unit  
     val read[x0] : unit -> string  
     val write2[x1] : string -> unit  
     val read2[x2] : unit -> string  
   end (valuable, valuable)  
  sig  
     val write[x] : string -> unit  
     val read[x0] : unit -> string  
     val write2[x1] : string -> unit  
     val read2[x2] : unit -> string  
   end  
  version name = h1_Persist  
  like struct      end  
  by Here_Already  
  = unlinked
```

```

cimport IO[Lib_IO] h2_IO
: sig
    val print_int[x] : int -> unit
    val print_string[x0] : string -> unit
    val print_newline[x1] : unit -> unit
    val send[x2] : string -> unit
    val receive[x3] : unit -> string
end (valuable, valuable)
sig
    val print_int[x] : int -> unit
    val print_string[x0] : string -> unit
    val print_newline[x1] : unit -> unit
    val send[x2] : string -> unit
    val receive[x3] : unit -> string
end
version name = h2_IO
like struct      end
by Here_Already
= unlinked

}, {},
{},

(function (x : thunkkey list) ->
    match x with Thread ((x2 : thread name), (x1 : thunkifymode))::Mutex (x0 : mutex name)::([] %[
        thunkkey]) ->
        unthunkify
        (Thunked_thread (x2,
            (function (x3 : unit) ->
                (let rec
                    x4 : int -> unit =
                    function
                        (x5 : int) ->
                            IO[Lib_IO].print_int x5; IO[Lib_IO].print_newline (); x4 (x5 + 1)
                    in
                        x4 ([4 ]int{} + 1))) ::
                    Thunked_mutex (x0, false) :: ([] %[thunklet]))),
                thunkkey list -> unit)
where
    h0_Pervasives = hash(hmodule Pervasives : {})
        sig
            val string_of_int[x] : int -> string
            val int_of_string[x0] : string -> int
            val print_string[x1] : string -> unit
            val print_int[x2] : int -> unit
            val print_endline[x3] : string -> unit
            val print_newline[x4] : unit -> unit
        end
        version myname
    = struct
        let string_of_int[x] = function (x : int) -> %"Apervasives_string_of_int" x
        let int_of_string[x0] =
            function (x0 : string) -> %"Apervasives_int_of_string" x0
        let print_string[x1] =
            function (x1 : string) -> %"Apervasives_print_string" x1
        let print_int[x2] = function (x2 : int) -> %"Apervasives_print_int" x2
        let print_endline[x3] =

```

```

        function (x3 : string) -> %"Apervasives_print_endline" x3
    let print_newline[x4] =
        function (x4 : unit) -> %"Apervasives_print_newline" x4
    end)
    = 0#4A5FE3EC8D80DFA70AD367461DD525AA
h1_Persist = hash(hmodule Persist : {})
    sig
        val write[x] : string -> unit
        val read[x0] : unit -> string
        val write2[x1] : string -> unit
        val read2[x2] : unit -> string
    end
    version myname
    = struct
        let write[x] = function (x : string) -> %"Persist_write" x
        let read[x0] = function (x0 : unit) -> %"Persist_read" x0
        let write2[x1] = function (x1 : string) -> %"Persist_write2" x1
        let read2[x2] = function (x2 : unit) -> %"Persist_read2" x2
        end)
    = 0#D90A83203B41EF6E6E512B3E5FF54850
h2_IO = hash(hmodule IO : {})
    sig
        val print_int[x] : int -> unit
        val print_string[x0] : string -> unit
        val print_newline[x1] : unit -> unit
        val send[x2] : string -> unit
        val receive[x3] : unit -> string
    end
    version myname
    = struct
        let print_int[x] = function (x : int) -> h0_Pervasives.print_int x
        let print_string[x0] = function (x0 : string) -> h0_Pervasives.print_string x0
        let print_newline[x1] =
            function (x1 : unit) -> match x1 with () -> h0_Pervasives.print_newline ()
        let send[x2] = function (x2 : string) -> h1_Persist.write x2
        let receive[x3] = function (x3 : unit) -> match x3 with () -> h1_Persist.read ()
        end)
    = 0#2808905E9138A8AA18FF6FF8E169EDED

```


Part III

Definition

16 Language Definition

16.1 Metavariables

<i>A</i>	set of module identifiers M_M and locations l
<i>BC</i>	bracket context
<i>C</i>	single-level evaluation context, or <i>definitions</i>
<i>CC</i>	evaluation context
<i>CVAL</i>	compile-time valuable context
<i>E</i>	type environment
<i>E_n</i>	type environment of global abstract names
<i>K</i>	kind
<i>L</i>	set of store locations
<i>M</i>	module identifier (external)
<i>MK</i>	mark (string literal \underline{s})
<i>M</i>	module identifier (internal)
<i>Mo</i>	module identifier option
<i>Ms</i>	sequence of module identifier
<i><u>N</u></i>	numeric hash
<i>P</i>	process
Φ	filesystem
<i>S</i>	set of module identifier
<i>SC</i>	structure evaluation context
<i>Sig</i>	signature
<i>Str</i>	structure
<i>T</i>	type
<i>TC</i>	thread top-level evaluation context
<i>TCC</i>	thread evaluation context
<i>TpubfromC</i>	type
<i>TrepfromC</i>	type
<i>URI</i>	Uniform Resource Identifier
<i>X</i>	module name or hash
<i>ahvc</i>	atomic hash version constraint
<i>ahvce</i>	atomic hash version constraint expression
<i>atomicresolvespec</i>	atomic resolve spec
<i>avc</i>	atomic version constraint
<i>avce</i>	atomic version constraint expression
<i>avn</i>	atomic version number

<i>avne</i>	atomic version number expression
<u><i>b</i></u>	boolean literal
<u><i>c</i></u>	character literal
<i>compilationunit</i>	compilation unit
<i>compilationunit</i>	compilation unit
<i>compileddefinition</i>	compiled definition
<i>compiledfilename</i>	filename of compiled file
<i>compiledunit</i>	compiled unit
<i>config</i>	runtime configuration
<i>definition</i>	module definition
<i>definitions</i>	module definitions
<i>dvc</i>	dotted version constraint
<i>dvce</i>	dotted version constraint expression
<i>e</i>	expression
<i>e_k</i>	expression
<i>ℓ</i>	transition label
<i>eo</i>	expression option
<i>eq</i>	type equation
<i>eqs</i>	type equation set
<i>h</i>	hash
<i>i</i>	index (from \mathbb{N})
<u><i>i</i></u>	integer literal
<i>j</i>	index (from \mathbb{N})
<i>k</i>	index (from \mathbb{N})
<i>l</i>	store location
<i>likespec</i>	likespec
<i>likestr</i>	structure (in a likespec)
<i>m</i>	index (from \mathbb{N})
<i>mode</i>	module or import mode
<i>mtch</i>	match
<i>mv</i>	marshalled value
<i>n</i>	index (from \mathbb{N})
<u><i>n</i></u>	natural number literal (from $\mathbb{N}_{2^{31}}$)
<i>n</i>	abstract name (from \mathcal{N})
n	abstract or hash name value
nn	abstract or hash name
<i>ns</i>	name list
<i>nset</i>	name set
<i>op</i>	operator
<i>p</i>	pattern
<i>resolvespec</i>	resolvespec
<i>ρ</i>	substitution of <i>T</i> 's for $M_M.t$'s
<i>s</i>	store
<i>σ</i>	substitution of <i>h.x</i> 's for $M_M.x$'s
<i>sig</i>	signature body
<i>sourcedefinition</i>	source language definition
<i>sourcefilename</i>	filename of source file

<i>sourcefilenames</i>	set of <i>sourcefilenames</i>
<i>s</i>	string literal
<i>str</i>	structure body
<i>strval</i>	structure body value
<i>t</i>	type identifier (internal)
<i>θ</i>	arbitrary syntactic entity
<i>thk</i>	thunk
<i>thks</i>	thunk list
<i>tk</i>	thunkkey
<i>tkls</i>	thunkkey list
<i>tmode</i>	thunkify mode
<i>t</i>	type identifier (external)
<i>u</i>	expression identifier (internal)
<i>v</i>	value
<i>valuability</i>	valuability
<i>valuabilities</i>	valuabilities
<i>vc</i>	version constraint
<i>vce</i>	version constraint expression
<i>vn</i>	version number
<i>vne</i>	version number expression
<i>weqs</i>	withspec type equation set
<i>withspec</i>	withspec
<i>x</i>	expression identifier (internal)
<i>xo</i>	expression identifier option
<i>x</i>	expression identifier (external)
<i>y</i>	expression identifier (internal)
<i>y</i>	expression identifier (external)
<i>z</i>	expression identifier (internal)
<i>z</i>	expression identifier (external)

16.2 Syntax

The definition involves several related languages:

1. The *concrete source* language is the language that programmers type, e.g. `function (x,y) -> x + y + M.z`. This is concrete — a set of character sequences.
2. The *sugared source internal* language is generated by parsing, scope resolution and type inference; for example `function (x : int, y : int) -> (+) x ((+) y MM.z)`. This is an abstract grammar, up to alpha equivalence. The *x*, *y* and *M* are internal identifiers, subject to alpha equivalence; the *z* and *M* are external identifiers, which are not. (In fact operators are eta-expanded to ensure they are fully applied.)
3. The *source internal* language is generated by desugaring, for example `function (u : int * int) -> match u with ((x : int), (y : int)) -> (+) x ((+) y MM.z)`.
4. The *compiled* language is generated by compilation, which here computes global type names for hashed abstract types, carries out *withspec* and *likespec* checks, etc. The operational semantics is defined over elements of the compiled language.

Note that the compiled language contains both compiled form and source internal form components. Specifically, a compiled program consists of compiled form definitions and/or source internal form **module fresh** definitions, and an optional compiled form expression.

The main definition is of the union of the grammars for the *source internal* and *compiled* languages. The differences are signposted with “*Source internal form:*” (or S) and “*Compiled form:*” (or C) respectively. The main type system is defined over this union.

Sugared forms (the *sugared source internal* additions to the *source internal* language) are signposted “*Sugared source internal language:*” (or G). Additional rules specify typing for the sugared forms.

Differences between the *sugared source internal* and the concrete syntax of the *concrete source* language are signposted “*Concrete source language:*”.

For any syntactic entity θ , we say `sugaredsourceinternalform(θ)`, `sourceinternalform(θ)` or `compiledform(θ)` to mean that entity is an element of the respective language.

Some syntactic requirements are not easily expressed in the BNF grammar itself. They are instead placed in the body of the text in paragraphs signposted with “*Syntactic requirement:*”.

Concrete source language: This definition does not fix character sets, comments, whitespace etc. The implementation generally follows OCaml.

Comment: The syntax generally follows OCaml for standard features. We have tried to resist any temptation to change or improve it, for three reasons: (1) to avoid time-consuming and unproductive syntactic debate; (2) to enable automated testing of the implementation of those standard features against OCaml’s behaviour; and (3) so that we and others can write Acute code without needing to learn new syntactic conventions. There are quite a number of things that should in principle be improved, however.

Identifiers

- x expression identifier (external)
- x* expression identifier (internal)
- t type identifier (external)
- t* type identifier (internal)
- M module identifier (external)
- M* module identifier (internal)

Concrete source language: Expression and type identifiers are not split into internal and external forms; they are uncapitalized. External module identifiers are capitalized; they can have an optional internal identifier, also capitalized.

We use x, x, t, t etc. both as metavariables (in most of this document) and as elements of their respective syntactic categories (in examples). Hence x_x (qua metavariables) ranges over x_x, x_y, y_x , etc (qua elements) – just because the two metavariables look similar does not mean that any concrete instance must be a pair related by the obvious isomorphism.

In ASCII when we need to write M_M, x_x , and t_t they are rendered, respectively, `MM[M]`, `xx[x]`, and `tt[t]`.

Kinds

K	<code>::=</code>	<code>TYPE</code>	kind of all types
		<code>EQ(<i>T</i>)</code>	kind of types equal to <i>T</i>

Types

$T ::=$ <ul style="list-style-type: none"> TC_0 $T \ TC_1$ $T_1 * .. * T_n \quad n \geq 2$ $T_1 + .. + T_n \quad n \geq 2$ n $T \rightarrow T'$ $M_M.t$ $h.t$ t $\forall t. T$ $\exists t. T$ 	$TC_0 ::=$ <ul style="list-style-type: none"> int bool string unit char void exn thread mutex cvar thunkifymode thunkkey thunklet unixerrorcode 	$TC_1 ::=$ <ul style="list-style-type: none"> list option ref name tie
	C	
		C

Here $M_M.t$ is a type field t from module M_M , and t (used within a structure or signature) is a type defined in a previous field.

Source internal form: $h.t$ is a global type name built from a module name; it is not permitted in source programs.

Compiled form: $M_M.t$ is not permitted in compiled form.

Type Environments

$E ::=$	<ul style="list-style-type: none"> empty $E, x : T$ $E, l : T \text{ ref}$ $E, M_M : Sig$ $E, t : K$ 	<ul style="list-style-type: none"> empty type environment

We write E, E' for the concatenation of two type environments, thereby asserting also that E and E' have disjoint domains. The domain $\text{dom}(E)$ of an E is a set of internal value identifiers, locations, module external/internal identifier pairs, and internal type identifiers.

Names

Take an infinite set \mathcal{N} of abstract names, ranged over by n . These are used to represent runtime and compile-time freshly-generated names.

We introduce a global type environment E_n associating abstract names with types, kind `TYPE`, or module/import data. Note that these can occur inside “closed” types, hashes etc.

$E_n ::=$	<ul style="list-style-type: none"> empty $E_n, n : \mathbf{nmodule}_{eqs} M : Sig_0 \ \mathbf{version} \ vne = Str$ $E_n, n : \mathbf{nimport} M : Sig_0 \ \mathbf{version} \ vc \ \mathbf{like} \ Str$ $E_n, n : \text{TYPE}$ $E_n, n : T \ \text{name}$

Comment: In the absence of first-class existentials freshly generated type names would not be required, as ML abstract types are a module-level feature.

Comment: We often write just E to stand for the pair E_n, E of a name environment and a type environment. In this case, $\text{namepart}(E)$ denotes the name environment component of this pair..

We let **nn** range over term names (hash- or fresh-generated) and **n** over their bracket-closure:

$$\frac{\begin{array}{l} \mathbf{nn} ::= \mathbf{hash}(h.x)_T \\ \mathbf{hash}(T', \underline{s})_T \\ \mathbf{hash}(T', \underline{s}, \mathbf{nn})_T \\ n_T \end{array}}{\mathbf{n} ::= \mathbf{nn} \quad [\mathbf{n}]_{eqs}^T}$$

In building these and other hashes we hash the abstract syntax up to alpha equivalence.

These are subgrammars of the e grammar; the e typing judgements apply.

Define the auxiliary typeof(**n**) to give the type subscript of the inner **nn**.

We suppose there is a fixed total order \leq over the **n**, taken (in the implementation) to depend on the hash / **n** only, ignoring the T subscripts).

Comment: Later we will also add name-indexed hashtables, which should respect the order.

We let h range over module names (hash- or fresh-generated).

$$\frac{\begin{array}{l} h ::= \mathbf{hash}(\mathbf{hmodule}_{eqs} M : Sig_0 \mathbf{version} vne = Str) \\ \mathbf{hash}(\mathbf{himport} M : Sig_0 \mathbf{version} vc \mathbf{like} Str) \\ n \end{array}}{\quad}$$

Comment: Note that only the external identifier M of a definition is included in the hash; the internal identifier is not. We will only ever deal with hashes in which eqs , Sig_0 and Str have been module-identifier-closed by substituting for $M'_{M'}.t$ and $M'_{M'}.x$. Furthermore, any internal module field references to type abbreviations will have been normalised away.

The version in a module hash is a version number expression, to permit it to include **myname**. This avoids the need for a recursive hash construction. In any context, **myname** may simply be interpreted as the hash in which it occurs. In contrast, the version in an import hash is not a version constraint expression but a version constraint. This is to force the evaluation of any M_M references, replacing them by the hash of the module named. Otherwise, the meaning of M_M would be (undesirably) context-dependent.

For the term part, we substitute $h.x$ for $M_M.x$ where h is the hash of whatever is bound to M_M . That gives a slightly more discriminating type equivalence than the ICFP calculus (which substituted code, not hash) for types that depend on the code containing that $h.x$, but it seems more intuitive, and is cheaper to implement. For the type part:

- where $M_M.t$ is abstract (of kind **TYPE** in the source definition) we substitute in $h.t$.
- where $M_M.t$ is concrete we must substitute in the type representation, otherwise we won't have enough type equalities later.

We assume a fixed function $\text{HASH}(\cdot)$ which takes a structured hash h to a numeric hash \underline{N} , where $\underline{N} \in \mathbb{H}$ for some set \mathbb{H} . Typically $\text{HASH}(\cdot)$ would be a well-known hash function such as MD5 or SHA1, and numeric hashes would just be long bit strings (128- or 160-bit, respectively).

With numeric hashes, runtime type safety for the language is only probabilistically guaranteed (though with rather high probability for reasonable usage); it depends on the assumption that $\text{HASH}(\cdot)$ is injective for the set of structured hashes in use.

The language is not intended to protect against the malicious forging of ill-formed hashes or marshalled values.

Implementation: In the implementation both **n** and all **hash**(...) forms will be represented by a long bitstring taken from \mathbb{H} . (So **hash**($h.x$) is represented by the hash of the pair of h and the external name x , not the pair of h and x .)

In the implementation, the representations of abstract names n will be generated randomly. More specifically: we do not want to require that the implementation generates each individual name randomly, as that might be too costly — it is acceptable to generate a random start point at the initialisation of each compilation and the initialisation of each runtime instance, and thereafter use some fast generation function for compile-time new and run-time new respectively. (Ideally the generation function would not be successor, to avoid triggering worst-case performance of naïve finite map implementations.) Nonetheless, a low-level attacker would often be able to tell whether two names originated from the same point, and that (for making real nonces etc) a more aggressively random `fresh` would be required.

Name representations could be generated lazily: as earlier discussed for FreshOCaml marshalling, we only really need an element of \mathbb{H} when a name is first marshalled; the implementation could keep a finite map associating internal-to-this-runtime names (represented just with pointers) and elements of \mathbb{H} that have been marshalled or have been unmarshalled from the outside. Whether we would gain very much by this is unclear, and we do not do it now. (However, it is important to make local channel use very cheap).

Implementation: In a production implementation, all occurrences of h would be implemented by occurrences of \underline{N} , with `HASH(.)` used where necessary to compute an \underline{N} from a `hash(.)` form of h .

In our current implementation we support both numeric hashes and structured hashes, the latter preserving all the structure above. A compiler option selects which are generated. This enables us (when using structured hashes) to typecheck the reachable intermediate states. Four points in the semantics describe a typecheck that can be performed if structured hashes are being used: in compilation when a compiled unit is imported; at runtime when a marshalled value is unmarshalled; at runtime during module field instantiation, when compiled definitions are taken from a URI; and at runtime after reduction steps (for the small-step evaluator, after every reduction step; for the big-step evaluator, only some of the intermediate points are reached). All these checks should always succeed, assuming that marshalled values and compiled files are not forged.

Our implementation takes a `HASH(.)` function that calculates the MD5 of a canonical pretty-print of structured hashes.

Hash equations

$$\frac{eq ::= h.t \approx T \quad M_M.t \approx T \quad S}{eqs ::= \emptyset \quad eq \quad eqs, eqs}$$

The `eqs` grammar is treated up to associativity, commutativity, idempotence, and identity.

The domain $\text{dom}(eqs)$ of an equation set is the set of types ($h.t$ or $M_M.t$ on the left-hand sides of equations in the set).

Comment: We believe that in fact, any equation set will consist either entirely of h -equations, or entirely of M_M -equations; the two will never be mixed. This is because M_M -equations appear in source form, and h -equations in compiled form. However, we do not (yet) model this in the abstract syntax, because we suspect carrying this through the type system would be painful. We should revisit this decision once the type system is stable, as there might be a clarity gain.

We occasionally use the metavariable X to stand for either h or M_M :

$$\frac{X ::= M_M}{h \quad C}$$

Compiled form: The M_M case of X is not permitted in compiled form.

Constructors The constructors are:

$()$: unit
 i : int
 b : bool
 c : char
 s : string
 $[]_T$: T list
NONE $_T$: T option
SOME : $T \rightarrow T$ option
TIECON : T name * $T \rightarrow T$ tie C
INJ $_i^{(T_1+\dots+T_n)}$: $T_i \rightarrow T_1 + .. + T_n$ $n \geq 2 \wedge i \in 1..n$
INTERRUPTING : thunkifymode
BLOCKING : thunkifymode
THREAD : thread name * thunkifymode \rightarrow thunkkey
MUTEX : mutex name \rightarrow thunkkey
CVAR : cvar name \rightarrow thunkkey
THUNKED_THREAD : thread name * (unit \rightarrow unit) \rightarrow thunklet C
THUNKED_MUTEX : mutex name * bool \rightarrow thunklet C
THUNKED_CVAR : cvar name \rightarrow thunklet C
RESOLVE_FAILURE : exn
MATCH_FAILURE : string * int * int \rightarrow exn
LIBRARY_ERROR : string \rightarrow exn
MARSHAL_FAILURE : exn
UNMARSHAL_FAILURE : string \rightarrow exn
FAILURE : string \rightarrow exn
INVALID_ARGUMENT : string \rightarrow exn
NOT_FOUND : exn
SYS_ERROR : string \rightarrow exn
END_OF_FILE : exn
DIVISION_BY_ZERO : exn
SYS_BLOCKED_IO : exn
NONEXISTENT_THREAD : exn
NONEXISTENT_MUTEX : exn
NONEXISTENT_CVAR : exn
MUTEX_EPERM : exn
EXISTENT_NAME : exn
THUNKIFY_EINTR : exn
THUNKIFY_SELF : exn
THUNKIFY_KEYLISTS_MISMATCH : exn
THUNKIFY_THREAD_IN_DEFINITION : exn
UNIXERROR : unixerrorcode * string * string \rightarrow exn

The unix error codes, all constructors of type unixerrorcode, are:

E2BIG
EACCES
EADDRINUSE
EADDRNOTAVAIL
EAFNOSUPPORT
EAGAIN
EWOULDBLOCK
EALREADY
EBADF

EBADMSG
EBUSY
ECANCELED
ECHILD
ECONNABORTED
ECONNREFUSED
ECONNRESET
EDEADLK
EDESTADDRREQ
EDOM
EDQUOT
EEXIST
EFAULT
EFBIG
EHOSTUNREACH
EIDRM
EILSEQ
EINPROGRESS
EINTR
EINVAL
EIO
EISCONN
EISDIR
ELOOP
EMFILE
EMLINK
EMSGSIZE
EMULTIHOP
ENAMETOOLONG
ENETDOWN
ENETRESET
ENETUNREACH
NFILE
ENOBUFS
ENODATA
ENODEV
ENOENT
ENOEXEC
ENOLCK
ENOLINK
ENOMEM
ENOMSG
ENOPROTOOPT
ENOSPC
ENOSR
ENOSTR
ENOSYS

ENOTCONN
 ENOTDIR
 ENOTEMPTY
 ENOTSOCK
 ENOTSUP
 ENOTTY
 ENXIO
 EOPNOTSUPP
 EOVERFLOW
 EPERM
 EPIPE
 EPROTO
 EPROTONOSUPPORT
 EPROTOTYPE
 ERANGE
 EROFS
 EPIPE
 ESRCH
 ESTALE
 ETIME
 ETIMEDOUT
 ETXTBSY
 EXDEV
 ESHUTDOWN
 EHOSTDOWN
 EUNKNOWN_UNIX_ERROR

Here i ranges over integer literals (the same as the underlying FreshOCaml ints), s ranges over strings of characters, and b ranges over `{true, false}`.

In addition to s , we let MK also range over string constants.

Note that constructors are all of arity 0 (C_0), arity 1 (C_1), or equal to `::` or `(-, ..., -)`. The typing and reduction rules treat the C_0 and C_1 cases uniformly and have special rules for the others.

Concrete source language: The type annotation subscripts are optional. If they are included (both here and in later forms), the linear ASCII rendering is e.g. `None [%[T]`.

The string `* int * int` in the v' for the `MATCH_FAILURE` case gives the position in the source file of the match code.

Many of the exception constructors are raised by embedded OCaml library functions, as follows:

- `INVALID_ARGUMENT` is raised by library functions to signal that the given arguments do not make sense.
- `FAILURE` is raised by library functions to signal that they are undefined on the given arguments.
- `NOT_FOUND` is raised by search functions when the desired object could not be found.
- `SYS_ERROR` is raised by the input/output functions to report an operating system error.
- `END_OF_FILE` is raised by input functions to signal that the end of file has been reached.
- `DIVISION_BY_ZERO` is raised by division and remainder operations when their second argument is null.
- `SYS_BLOCKED_IO` is a special case of `SYS_ERROR` raised when no I/O is possible on a non-blocking I/O channel.
- `UNIXERROR` carries the errors raised by the TCP libraries.
- `LIBRARY_ERROR` carries any unrecognised error raised by an E_{const} .

Comment: The Unix error codes above are the set of all those on our current Linux install, and the translation from integers to constructors is hard-wired into the Acute implementation (in `library.mlp`). This should be made more portable — at the least, that part of `library.mlp` should be automatically generated from the C header file.

Comment: Note that the polymorphic constructors exist as indexed families rather than using explicit polymorphism. This is a historical artifact.

Comment: We are not entirely consistent about the type annotations on constructors, operators, and expression forms. Acute was originally monomorphic, though with type inference for these annotations; the semantics was originally written to ensure that all values have unique types. That is no longer the case: the `raise` is *not* type-annotated (as to maintain that annotation during reduction would require notationally-heavy annotation of evaluation contexts and the other expression forms), so function values with a `raise` in the body may not be uniquely typable.

Standard Library

We suppose there is a fixed collection of special constants E_{const} , which is a finite partial map from internal value identifiers to types. Each is equipped with a natural-number arity, written x^n if x has arity n . The special constants are partitioned by a predicate $\text{os}(x^n)$ into the OS calls, which have labelled transitions in the semantics, and the internal built-in library calls, which have delta rules. The OS calls are further partitioned by a predicate $\text{fast}(x^n)$ specifying whether each is a fast or slow call.

Their internal identifiers are never shadowed, as specified below when we discuss binding.

The types of E_{const} s must be first order.

Comment: Before the addition of concurrency we permitted higher-order E_{const} s, e.g. to automatically embed the FreshOCaml `List.map` into Acute, but with concurrency that would be unduly complex.

Suppose further that there is a fixed list of library definitions $\text{definitions}_{\text{lib}}$, a finite list of module definitions. These have a special status in that their code can mention special constants from $\text{dom}(E_{\text{const}})$ whereas user-defined modules cannot (the running expression can also mention them, of course).

Note that the internal identifiers of $\text{definitions}_{\text{lib}}$ are fixed globally.

We generate names for these modules in the usual way when they are compiled (note there will be free internal identifiers inside, but that is not a problem).

Let E_{lib} be the partial map from module external/internal identifier pairs to signatures such that $E_{\text{const}} \vdash \text{definitions}_{\text{lib}} \triangleright E_{\text{lib}}$.

The upshot of this is that all types defined in a $\text{definitions}_{\text{lib}}$ module must have representation types that are expressible within the language, but the code can make use of E_{const} . We do not require that $\text{definitions}_{\text{lib}}$ terms are pure E_{const} . Programs can rebound to user-land replacements for $\text{definitions}_{\text{lib}}$ modules if needed, and can use them in *withspec* and *likespecs*.

Implementation: In the implementation $\text{definitions}_{\text{lib}}$ is composed of two parts. The first (`definitions_lib_auto.ac`) is automatically generated from a collection of OCaml interface files; each value component in these gives rise to an E_{const} . These interfaces are described in Section 21. Most are simple fragments of OCaml standard library interfaces, and are linked to those; some are linked to hand-written OCaml modules. Type embeddings and projections are dealt with automatically. The second part consists of various hand-written Acute modules. The two are combined into `definitions_lib.ac` as below.

```
includesource "definitions_lib_auto.ac"
(* includesource "io_template.ac"                (* simple IO for tcp          *) *)
includesource "io_persist.ac"                   (* simple IO for persistent store *)
```

E_{const} s of arity 0 are now supported by the automated generation tool but they give non-value expressions in the `definitions_lib.ac` structures rather than the actual values, so we use the `hash!` mode for these modules.

Comment: Note that the semantics has immutable strings, whereas OCaml has mutable strings. Our string library contains only the non-mutating part of the OCaml string library.

Operators Take *operators* op^n

ref _T	:	$T \rightarrow T$ ref	
(=) _T	:	$T \rightarrow T \rightarrow \text{bool}$	
(<), (<=), (>), (>=)	:	$\text{int} \rightarrow \text{int} \rightarrow \text{bool}$	
(+), (-), (*), (/), (mod)	:	$\text{int} \rightarrow \text{int} \rightarrow \text{int}$	
(land), (lor), (lxor)	:	$\text{int} \rightarrow \text{int} \rightarrow \text{int}$	
(lsl), (lsr), (asr)	:	$\text{int} \rightarrow \text{int} \rightarrow \text{int}$	
-	:	$\text{int} \rightarrow \text{int}$	
(@) _T	:	$T \text{ list} \rightarrow T \text{ list} \rightarrow T \text{ list}$	
(^)	:	$\text{string} \rightarrow \text{string} \rightarrow \text{string}$	
compare_name _T	:	$T \text{ name} \rightarrow T \text{ name} \rightarrow \text{int}$	
create_thread _T	:	$\text{thread name} \rightarrow (T \rightarrow \text{unit}) \rightarrow T \rightarrow \text{unit}$	
self	:	$\text{unit} \rightarrow \text{thread name}$	
kill	:	$\text{thread name} \rightarrow \text{unit}$	
create_mutex	:	$\text{mutex name} \rightarrow \text{unit}$	
lock	:	$\text{mutex name} \rightarrow \text{unit}$	
try_lock	:	$\text{mutex name} \rightarrow \text{bool}$	
unlock	:	$\text{mutex name} \rightarrow \text{unit}$	
create_cvar	:	$\text{cvar name} \rightarrow \text{unit}$	
wait	:	$\text{cvar name} \rightarrow \text{mutex name} \rightarrow \text{unit}$	
signal	:	$\text{cvar name} \rightarrow \text{unit}$	
broadcast	:	$\text{cvar name} \rightarrow \text{unit}$	
thunkify ¹	:	$\text{thunkkey list} \rightarrow (\text{thunkkey list} \rightarrow \text{unit})$	
unthunkify	:	$\text{thunklet list} \rightarrow \text{thunkkey list} \rightarrow \text{unit}$	C
exit _T	:	$\text{int} \rightarrow T$	

The superscript is the arity of the operator. Note in particular that **thunkify** has arity 1, not 2.

Concrete source language: The binary operators in brackets may be written infix, e.g. $e =_T e'$ for $(=_T) e e'$; we use Ocaml's precedence rules. If ref_T is not saturated, then it must be enclosed in parentheses in source forms. Same for **mod**, **land**, **lor**, **lxor**, **lsl**, **lsr**, **asr**. The type subscripts can be omitted, as above.

Comment: With locally-unique naming, there is no point in parameterising the **create_thread** function argument on its identity.

Comment: The type of **compare_name** follows the type of **compare** in OCaml.

The operators come in two families: the *type indexed*, consisting of those bearing a type subscript, and the unindexed. We write op^n for both.

Comment: The definition does not at present follow a consistent policy as to what should appear as an operator and what as an expression form (cf. the treatment of coloured arguments by the atomic evaluation contexts). Ultimately it should. The distinction between operators and E_{const} s comes from the implementation: the former are implemented within the Acute runtime; the latter by calling out to FreshOCaml.

Expressions

$e ::=$	C_0	C_0 a constructor of arity 0
	$C_1 e$	C_1 a constructor of arity 1
	$e_1 :: e_2$	Cons
	(e_1, \dots, e_n)	Tuple ($n \geq 2$)
	function $(x : T) \rightarrow e$	Function
	$op^n e_1 \dots e_n$	op an operator
	$x^n e_1 \dots e_n$	x^n an external constant
	x	Identifier

$M_{M.X}$	Module projection
$h.x$	* Module hash projection
if e_1 then e_2 else e_3	Conditional
while e_1 do e_2 done	Loop
$e_1 \&\& e_2$	Boolean short-circuit and
$e_1 e_2$	Boolean short-circuit or
$e_1 ; e_2$	Sequence
$e_1 e_2$	Application
$!_T e$	Deref
$e_1 :=_T e_2$	Assign
$e_1 :='_T e_2$	C Assign uncoloured
l	C Location
match e with $mtch$	Pattern match
let rec $x_1 : T = \mathbf{function} (x_2 : T') \rightarrow e_1$ in e_2	Recursive definition
raise e	Raise exception
try e with $mtch$	Handle exception(s)
marshal $e_1 e_2 : T$	Marshal
marshalz $\underline{s} e : T$	C Marshal (expression in uncoloured context)
unmarshal e as T	Unmarshal
fresh $_T$	run-time fresh name generation
cfresh $_T$	S compile-time fresh name generation
hash ($X.x$) $_{T'}$	create name from module value field
hash (T, e_1) $_{T'}$	create name from type and string
hash (T, e_1, e_2) $_{T'}$	create name from type, string, and name
n_T	C abstract name
swap e_1 and e_2 in e_3	polytypic swap
e_1 freshfor e_2	polytypic freshness test
support $_T e$	polytypic typed-name support
$M_M @ X$	S tie construction
name_of_tie e	tie inspection
val_of_tie e	tie inspection
$\Lambda t \rightarrow e$	type abstraction
$e T$	type application
$\{T, e\}$ as T'	existential package
let $\{t, x\} = e_1$ in e_2	unpackaging
namecase e_1 with $\{t, (x_1, x_2)\}$ when $x_1 = e \rightarrow e_2$ otherwise $\rightarrow e_3$	unpackaging and name equality
function $mtch$	G ($mtch \neq (x' : T' \rightarrow e)$)
fun $p_1..p_n \rightarrow e'$	G ($n \geq 1$)
let $p = e'$ in e''	G
let $x : T p_1..p_n = e'$ in e''	G ($n \geq 1$)
let rec $x : T = \mathbf{function} mtch$ in e	G ($mtch \neq (x' : T' \rightarrow e')$)
let rec $x : T p_1..p_n = e'$ in e''	G ($n \geq 1$)
$e_1 e_2$	G spawn e_1
op (op^n) $^n e_1 .. e_n$	C Primitive application of an operator
op (x^n) $^n e_1 .. e_n$	C Primitive application of an external constant

$[e]_{eqs}^T$	C	Coloured brackets
resolve ($M_M.x, M'_{M'}, resolvespec$)	C	<i>resolvespec</i> in progress
resolve_blocked ($M_M.x, M'_{M'}, resolvespec$)	C	<i>resolvespec</i> blocked waiting for data
RET _T	C	Await return from a 'fast' OS routine
SLOWRET _T	C	Await return from a 'slow' OS routine

We write x^n to denote an $x \in E_{\text{const}}$ that has arity n .

Sometimes we write $\text{op}(e)^n \dots$ and in this case the e ranges over op^n and x^n only.

Concrete source language: The type annotations in $!_T e$, $e_1 :=_T e_2$, **function** ($x : T$) $\rightarrow e$ and **let rec** $x_1 : T = \text{function}$ ($x_2 : T'$) $\rightarrow e_1$ **in** e_2 are all optional. In ASCII a type abstraction $\Lambda t \rightarrow e$ is written as **Function** $t \rightarrow e$ and a type application $e T$ is written $e \%[T]$.

Sugared source internal language: The $: T$ type annotation in **let** $x : T p_1..p_n$ and **let rec** $x : T p_1..p_n$ is prohibited (to be compatible with Ocaml). These type annotations are inserted by type inference and used in the desugaring process.

Sugared source internal language: The **hash**($M_M.x$)_{T'} and $M_M@x$ forms are only permitted within structures, not in the main expression. (This is not essential, but simplifies the hashify semantics.)

Compiled form: Module hash projections $h.x$ may only occur within other hashes (they are not executable).

Sugared source internal language: In source programs, $!_T$, $:=_T$, $||$, and $\&\&$ may all be written as prefix functions by wrapping them in parentheses, e.g. $(:=_T)$. In source programs, operators, external constants, $!_T$, $:=_T$, $||$, and $\&\&$ can be partially applied. The desugaring process is responsible for eta-expanding these. Type annotations in these sugared forms are likewise optional.

See Section 16.4 for details of the desugarings.

Comment: It is an invariant that constructible values v^{eqs} satisfy $\text{compiledform}(v^{eqs})$ and in addition contain none of **RET**_T, **SLOWRET**_T, **resolve**(...), or **resolve_blocked**(...). Values can contain the forms l , $[e]_{eqs}^T$, n_T , and also the (transient) $e_1 :='_T e_2$, **marshalz** $\underline{s} e : T$, $\text{op}(e)^n e_1 .. e_n$. (See the semantics for **thunkify**, which cannot create a thunk containing the first group.) Likewise, marshalled and stored values contain none of the first group.

Marshalled values

$mv ::= \text{marshalled}(E_n, E_s, s, \text{definitions}, e, T)$ Marshalled value

We suppose a fixed partial function `raw_unmarshal` from strings to marshalled values that includes all marshalled values in its range.

Syntactic requirement: The components θ of a marshalled value all satisfy $\text{compiledform}(\theta)$.

Comment: Here e is the core value being shipped, T its type, s a store, E_s a store typing, definitions is a sequence of module definitions, and E_n is a name environment.

The E_n and E_s would not be shipped in an production implementation, but are needed to state type preservation and for runtime typechecking of reachable states. They are shipped in our implementation only if literal hashes are not being used.

As with the other syntactic objects, marshalled values are taken up to alpha equivalence. Here: the name environment E_n binds in everything to the right and internally contains no cycles; the store environment E_s binds in everything to the right and may contain internal cycles; the store s and the definitions bind to the right and may mutually refer to each other; the s may contain internal cycles.

Implementation: The implementation of marshalled values should include a global type name for the Acute implementation representation type. As we are not bootstrapping, we should do this manually.

Matches

$mtch ::= p \rightarrow e$	
$p \rightarrow e mtch$	

Concrete source language: An initial bar may be added.

Patterns

$p ::= (_ : T)$	Wildcard
$(x : T)$	Identifier
C_0	C_0 a constructor of arity 0
$C_1 p$	C_1 a constructor of arity 1
$p_1 :: p_2$	Cons
(p_1, \dots, p_n)	Tuple ($n \geq 2$)
$(p : T)$	Typed pattern

Syntactic requirement: These are subject to the condition that all identifiers occurring in a pattern are distinct.

Concrete source language: the type annotations on wildcard and identifier patterns can be omitted.

Signatures

$sig ::= \text{empty}$	Empty signature body
$\text{val } x_x : T \text{ sig}$	Signature body extended with val spec
$\text{type } t_t : K \text{ sig}$	Signature body extended with type spec
$Sig ::= \text{sig sig end}$	Signature

Concrete source language: We write $t : \text{TYPE}$ as t , write $t : \text{EQ}(T)$ as $t = T$, and allow optional `;;` between each non-empty spec in a *sig*. We write a single identifier in place of x_x and t_t .

Structures

$str ::= \text{empty}$	Empty structure body
$\text{type } t_t = T \text{ str}$	Structure body extended with type component
$\text{let } x_x = e \text{ str}$	Structure body extended with expression component
$\text{let } x_x : T \text{ } p_1..p_n = e' \text{ G } (n \geq 1)$	
$Str ::= \text{struct str end}$	Structure

Concrete source language: We allow optional `;;` between each non-empty spec in a *str*. We write a single identifier in place of x_x and t_t . To match Ocaml the $: T$ is prohibited (but inserted by the type inference system).

Resolve specs

$atomicresolvespec ::=$	atomic resolve spec
STATIC_LINK	code should be statically linked
HERE_ALREADY	code should be here already, fail if not
URI	load module from file or web
$resolvespec ::=$	resolve spec (nonempty list of atomic ones)
$atomicresolvespec$	
$atomicresolvespec, resolvespec$	
$\text{URI} ::=$	a string literal of a URI...
	(a subgrammar of RFC2396's absoluteURI)

Implementation: The current implementation supports file, http, and ftp URIs.

Version languages We define version number and constraint *expressions* as follows.

<i>avne</i> ::=	\underline{n} \underline{N} h myname	Atomic version number expression natural number literal in $\mathbb{N}_{2^{31}}$ numeric name literal in \mathbb{H} C structured name literal the compiler will write the name of this module in as a literal
<i>vne</i> ::=	\underline{n} \underline{N} h myname	Version number expression atomic version dotted version
<i>ahvce</i> ::=	\underline{N} h M_M	Atomic hash version constraint expression numeric name literal in \mathbb{H} C structured name literal the compiler will write the hash of M_M in as a literal
<i>avce</i> ::=	\underline{n} \underline{N} h M_M	Atomic version constraint expression atomic name version constraint expression
<i>dvce</i> ::=	\underline{n} $\underline{n}-\underline{n}'$ $-\underline{n}$ $\underline{n}-$ $*$	Dotted version constraint atomic constraint closed interval left-open interval right-open interval anything
<i>vce</i> ::=	\underline{n} \underline{N} h M_M $\underline{n}-\underline{n}'$ $-\underline{n}$ $\underline{n}-$ $*$	Version constraint dotted version constraint dotted version constraint exact-name version constraint

Syntactic requirement: We define the version number and constraint *values* *avn*, *vn*, *avc*, *ahvc*, *dvce*, *vc* to be the relevant subgrammars with the **myname** and M_M clauses removed.

Source internal form: A user source program may not have an exact-name constraint of the form **name** = \underline{N} , or **name** = h , only **name** = M_M , as an in-scope module is required to provide the data to construct a *likestr*.

Comment: There is an important distinction between h and \underline{N} . In the semantics a structured name h can be supplied only by the compiler, and thus we may ensure and assume it is generated from a well-formed and well-typed module or import. A numeric name \underline{N} in a version expression may be supplied by the user as an arbitrary element of \mathbb{H} (e.g., 0#60139C0047463B6261112944981EBF92), and thus (for type-safety purposes) cannot be assumed to arise from a well-formed structured name (i.e. be either the $\text{HASH}(\cdot)$ of a well-formed structured hash or be an appropriate abstract name).

Comment: Note that the semantics of an exact-name version constraint **name** = *ahvce* is rather different from the other *vces* in that it is a constraint on the name, not the version, of the modules and imports that can be linked to an import with this constraint.

Comment: The basic part of the version grammar should be improved: the intervals are not very useful as given here.

Define an equivalence relation over *avc* (note that *avn* and *avc* coincide) as the least equivalence such that $h \cong \underline{N}'$ if $\text{HASH}(h) = \underline{N}'$. More explicitly:

$$\begin{aligned}
 \underline{n} \cong \underline{n}' &\iff \underline{n} = \underline{n}' \\
 \underline{N} \cong \underline{N}' &\iff \underline{N} = \underline{N}' \\
 h \cong h' &\iff h = h' \\
 h \cong \underline{N}' &\iff \text{HASH}(h) = \underline{N}' \\
 \underline{N} \cong h' &\iff \underline{N} = \text{HASH}(h')
 \end{aligned}$$

Define the set of vn denoted by each dvc as follows.

$$\begin{aligned}
 \llbracket N \rrbracket &= \{avn \mid avn \cong N\} \\
 \llbracket h \rrbracket &= \{avn \mid avn \cong h\} \\
 \llbracket n \rrbracket &= \{n\} \\
 \llbracket n_1 - n_2 \rrbracket &= \{n \mid n_1 \leq n \leq n_2\} \\
 \llbracket n_1^- \rrbracket &= \{n \mid n_1 \leq n\} \\
 \llbracket -n_2 \rrbracket &= \{n \mid n \leq n_2\} \\
 \llbracket * \rrbracket &= \{vn \mid \mathbf{true}\} \\
 \llbracket avc.dvc \rrbracket &= \{avn.vn \mid avn \cong avc \wedge vn \in \llbracket dvc \rrbracket\} \cup \{avn \mid avn \cong avc \wedge dvc = *\}
 \end{aligned}$$

We write $vn \in dvc$ for $vn \in \llbracket dvc \rrbracket$.

Say $vc \subseteq vc'$ if either (1) $vc = dvc$, $vc' = dvc'$ and $\llbracket dvc \rrbracket \subseteq \llbracket dvc' \rrbracket$, or (2) $vc = (\mathbf{name} = ahvc)$ and $vc' = (\mathbf{name} = ahvc')$ and $ahvc \cong ahvc'$.

Modes and Valuabilities

<i>mode</i> ::=	
hash	hash the structure of the module or import
cfresh	calculate a fresh name at compile time
fresh	calculate a fresh name at run time
hash!	hash the structure of the module or import, ignoring valuability
cfresh!	calculate a fresh name at compile time, ignoring valuability

<i>vub</i> ::=	valuable	is statically determined
	cvaluable	is statically determined after compile-time new
	nonvaluable	can only be calculated at run-time

We write $vubs$ for a pair of valuabilities. The first element of the pair refers to the status of the terms, the second to the status of the types.

Definitions Source definitions and compiled definitions (the latter ranged over by *definitions*) are as follows.

<i>sourcedefinition</i> ::=		
module $mode M_M : Sig$ version $vne = Str$ <i>withspec</i>		Module declaration
import $mode M_M : Sig$ version vce <i>likespec</i> by $resolvespec = Mo$		Module import
mark MK		Mark
module $M_M : Sig = M'_{M'}$		Module alias declaration

<i>definition</i> ::=		
cmodule $_{h;eqs;Sig_0}$ $vubs M_M : Sig_1$ version $vn = Str$		Module declaration
cimport $_{h;Sig_0}$ $vubs M_M : Sig_1$ version vc like Str by $resolvespec = Mo$		Module import
module fresh $M_M : Sig$ version $vne = Str$ <i>withspec</i>		... (initialisation-time fresh)
import fresh $M_M : Sig$ version vce <i>likespec</i> by $resolvespec = Mo$... (initialisation-time fresh)
mark MK		Mark

In the **module** the *eqs* are any equations arising from the **with !** clause; the *Sig₀* is the semicompiled signature, not name-selfified but otherwise normalised as far as possible, and *Sig₁* is the fully compiled signature.

<i>weqs</i>	::= \emptyset	user coercion spec
	$M_M.t \approx T, weqs$	
<i>withspec</i>	::= with ! <i>weqs</i>	
<i>likespec</i>	::= empty	like spec
	like M_M	
	like <i>Str</i>	
<i>Mo</i>	::= M_M	linked to M_M
	UNLINKED	unlinked

Syntactic requirement: Here *weqs* is up to associativity, commutativity and identity.

Concrete source language: The **version** *vne* in a **module** can be omitted, in which case it defaults to **version myname**. The *withspec* in a **module** is either empty (in which case it defaults to **with ! \emptyset**) or **with !weqs** in which case *weqs* is not empty. The **version** *vce* in an **import** can be omitted, in which case it defaults to **version ***. The **by** *resolvespec* in an **import** can be omitted, in which case it defaults to **by HERE_ALREADY**. The = *Mo* in an **import** can be omitted, in which case it defaults to = UNLINKED. If an **import** has an exact-name constraint **name = M_M** then the *likespec* must be empty.

Compilation Units

<i>compilationunit</i>	::= <i>eo</i>
	<i>sourcedefinition</i> ;; <i>compilationunit</i>
	includesource <i>sourcefilename</i> ;; <i>compilationunit</i>
	includecompiled <i>compiledfilename</i> ;; <i>compilationunit</i>
<i>compiledunit</i>	::= ($E_n, definitions$ <i>eo</i>)
<i>definitions</i>	::= empty
	<i>definition</i> ;; <i>definitions</i>
<i>eo</i>	::= empty
	<i>e</i> ;;

Concrete source language: We allow optional and repeated ;; (different rules apply for structures and signatures).

Compiled form: All ;; are omitted.

There must be at most one final *e*, which may be e.g. at the end of an include at the end of the top-level compilation unit.

In the current implementation, programs with no final expression are not executed; in particular, no module initialisation is performed for such programs. This restriction should be relaxed.

Filesystems

Say a *filesystem* Φ is a finite partial map from *sourcefilename* to *compilationunits* and from *compiledfilename* to *compiledunits*.

Conventionally, *sourcefilenames* are of the form *foo.ac* and *compiledfilenames* are of the form *foo.aco*.

Processes

$$\frac{P ::= 0 \quad P_1|P_2 \quad \mathbf{n} : \text{definitions } e \quad \mathbf{n} : \text{MX}(\underline{b}) \quad \mathbf{n} : \text{CV}}{}$$

We work up to the structural congruence on processes which is the least congruence for $|$ containing $P|0 \equiv P$, $P_1|P_2 \equiv P_2|P_1$, and $P_1|(P_2|P_3) \equiv (P_1|P_2)|P_3$.

Write $\text{dom}(P)$ for the set of names of entities in P , i.e. $\text{dom}(0) = \emptyset$, $\text{dom}(P_1|P_2) = \text{dom}(P_1) \cup \text{dom}(P_2)$, $\text{dom}(\mathbf{n} : \dots) = \{\mathbf{n}\}$.

Stores

Say a *store* s is a finite partial map from locations l to \emptyset -coloured values (values are defined on page 125).

Configurations (or States)

Take tuples $\langle E_s, s, \text{definitions}, P \rangle$ of some module *definitions*, a store typing E_s , a store s , and a process P . The store typing is not needed in an implementation. Note that (as we have module initialisation) the E_s, s scope in s, P and *definitions*, and the *definitions* scope in s and P .

16.2.1 Binding

Syntactic requirement: We work up to alpha equivalence throughout.

Syntactic requirement: The external constants $x^n \in \text{dom}(E_{\text{const}})$ may not appear in a binding position.

Syntactic requirement: For expression and type identifiers we have internal identifiers x and t as normal binders rather than the external/internal pair a binder (as in [BHS⁺03]). For module identifiers we have the M_M pairs be binders.

We write $\text{fv}(\dots)$ for the set of free identifiers x, t , and M_M in \dots .

Syntactic requirement: In expression **function** $(x : T) \rightarrow e$ the x binds in e . In expression **let rec** $x_1 : T = \text{function } (x_2 : T') \rightarrow e_1$ **in** e_2 the x_1 binds in e_1 and e_2 , and the x_2 binds in e_1 .

Syntactic requirement: In sugar expression **let** $p = e_1$ **in** e_2 the internal value identifiers of p bind in e_2 . In sugar expression **let rec** $x : T = \text{function } \text{mtch}$ **in** e the x binds in mtch and in e . In sugar expression **fun** $p_1..p_n \rightarrow e$ the internal value identifiers of $p_1..p_n$ bind in e . In sugar expression **let** $x p_1..p_n = e_1$ **in** e_2 the identifier x binds in e_2 , and the internal value identifiers of $p_1..p_n$ bind in e_1 . In sugar expression **let rec** $x p_1..p_n = e_1$ **in** e_2 the identifier x binds in e_1 and e_2 , and the internal value identifiers of $p_1..p_n$ bind in e_1 .

Syntactic requirement: In $\Lambda t \rightarrow e$ the t binds in the e . In **let** $\{t, x\} = e_1$ **in** e_2 the t and x bind in the e_2 . In **namecase** e_1 **with** $\{t, (x_1, x_2)\}$ **when** $x_1 = e \rightarrow e_2$ **otherwise** $\rightarrow e_3$, type identifier t , expression identifier x_2 , and the first occurrence of expression identifier x_1 all bind in e_2 ; the second occurrence of x_1 is bound by the first occurrence. Note that e_1, e , and e_3 all live in the outer scope (no extra bindings).

Syntactic requirement: In match $p \rightarrow e$ the internal expression identifiers of p bind in e

Syntactic requirement: In signatures, in **val** $x_x : T \text{ sig}$ the x binds in sig and in **type** $t_t : K \text{ sig}$ the t binds in sig .

Syntactic requirement: In structures, in **let** $x_x = e \text{ str}$ the x binds in str , in **let** $(x_x : T)p_1..p_n = e \text{ str}$ the internal value identifiers of $p_1..p_n$ bind in e and the x binds in str ; and in **type** $t_t = T \text{ str}$ the t binds in str .

Syntactic requirement: For any occurrence of a *sourcedefinition* or *definition*, the M_M binds in subsequent definitions. It also binds in any subsequent store typing E_s , store s and expression or process e or P , e.g. when the *definitions* appear in a configuration or marshalled body.

Comment: Note that **mark** MK does not involve any binding – marks are just strings, as marks must be shared across programs.

We've (arbitrarily) chosen not to have the store bind the locations of its domain, as would have to chose whether the E_s or the s bind, or agglomerate the two.

16.3 Typing

The typing judgements are listed in the contents pages. The typing rules are in Figures below, with particularly interesting rules flagged ★.

Most judgements are parameterised on a set *eqs* of type equations. These are kept as a subscript instead of as part of *E* so they be easily removed when one passes through brackets – unlike binders, they are not additive.

There is no $E \vdash \text{compilationunit} \text{ ok}$ as we need to substitute file contents in (recursively) before typechecking, not having introduced separate interfaces.

The source language type system must be considered together with the checks performed by compilation: several checks are not carried out in the type system because they involve the representation types of abstract types, and version data, from previous modules; these are not recorded in type environments and so are not accessible in the type system. Specifically: (i) formation of the equation $E \vdash M_M.t \approx T \text{ ok}$ (used especially for the *weqs* in the **module** rule), and (ii) link-checking of a loaded import in the **import** rule, are only weakly constrained by the source type system.

The compiled language type system checks these explicitly, and enforces additional facts, e.g. that in compiled form occurrences of $M_M.t$ have been hashified to the *h.t* form. Also, the *h.x* form appears only within hashes.

The dynamic semantics is only intended to make sense for configurations that typecheck in the compiled language type system.

16.3.1 Typing for Source Internal and Compiled Forms

$E_n \vdash \text{ok}$	
$\frac{n \notin \text{dom}(E_n) \quad E_n \vdash_{\emptyset} T : \text{TYPE}}{E_n, n : T \text{ name} \vdash \text{ok}}$	$\frac{n \notin \text{dom}(E_n)}{E_n, n : \text{TYPE} \vdash \text{ok}}$
$\frac{n \notin \text{dom}(E_n) \quad E_n \vdash \text{nmodule}_{eqs} M : Sig_0 \text{ version } vne = Str \text{ ok}}{E_n, n : \text{nmodule}_{eqs} M : Sig_0 \text{ version } vne = Str \vdash \text{ok}}$	
$\frac{}{\text{empty} \vdash \text{ok}}$	$\frac{n \notin \text{dom}(E_n) \quad E_n \vdash \text{nimport} M : Sig_0 \text{ version } vc \text{ like } Str \text{ ok}}{E_n, n : \text{nimport} M : Sig_0 \text{ version } vc \text{ like } Str \vdash \text{ok}}$
$E_n, E \vdash \text{ok}$	
$\frac{E_n \vdash \text{ok}}{E_n, \text{empty} \vdash \text{ok}}$	$\frac{x \notin \text{dom}(E) \quad E_n, E \vdash_{\emptyset} T : \text{TYPE}}{E_n, E, x : T \vdash \text{ok}}$
$\frac{l \notin \text{dom}(E) \quad E_n, E \vdash_{\emptyset} T : \text{TYPE}}{E_n, E, l : T \text{ ref} \vdash \text{ok}}$	
$\frac{M_M \notin \text{dom}(E) \quad E_n, E \vdash Sig \text{ ok}}{E_n, E, M_M : Sig \vdash \text{ok}}$	$\frac{t \notin \text{dom}(E) \quad E_n, E \vdash K \text{ ok}}{E_n, E, t : K \vdash \text{ok}}$
$E_n \vdash \text{nmodule}_{eqs} M : Sig_0 \text{ version } vne = Str \text{ ok} \quad E_n \vdash \text{nimport}_{eqs} M : Sig_0 \text{ version } vc \text{ like } Str \text{ ok}$	
$\frac{E_n, E_{\text{const}} \vdash_{eqs} Str : Sig_0 \quad \vdash Str \text{ flat} \quad \vdash Sig_0 \text{ flat}}{E_n \vdash \text{nmodule}_{eqs} M : Sig_0 \text{ version } vne = Str \text{ ok}}$	$\frac{E_n \vdash_{\emptyset} Str : \text{limitdom}(Sig_0) \quad E_n \vdash Sig_0 \text{ ok} \quad \vdash Str \text{ flat} \quad \vdash Sig \text{ flat}}{E_n \vdash \text{nimport}_{eqs} M : Sig_0 \text{ version } vc \text{ like } Str \text{ ok}}$
$E_n \vdash h \text{ ok}$	
$\frac{h = \text{hash}(\text{hmodule}_{eqs} M : Sig_0 \text{ version } vne = Str) \quad E_n, E_{\text{const}} \vdash_{eqs} Str : Sig_0 \quad \vdash Str \text{ flat} \quad \vdash Sig_0 \text{ flat}}{E_n \vdash h \text{ ok}}$	$\frac{h = n \quad (n : \text{nmodule}_{eqs} M : Sig_0 \text{ version } vne = Str \text{ ok}) \in E_n \quad E_n, E_{\text{const}} \vdash_{eqs} Str : Sig_0 \quad \vdash Str \text{ flat} \quad \vdash Sig_0 \text{ flat}}{E_n \vdash h \text{ ok}} \quad \star$
$\frac{h = \text{hash}(\text{himport} M : Sig_0 \text{ version } vc \text{ like } Str) \quad E_n \vdash_{\emptyset} Str : \text{limitdom}(Sig_0) \quad E_n \vdash Sig_0 \text{ ok} \quad \vdash Str \text{ flat} \quad \vdash Sig \text{ flat}}{E_n \vdash h \text{ ok}}$	$\frac{h = n \quad (n : \text{nimport} M : Sig_0 \text{ version } vc \text{ like } Str) \in E_n \quad E_n \vdash_{\emptyset} Str : \text{limitdom}(Sig_0) \quad E_n \vdash Sig_0 \text{ ok} \quad \vdash Str \text{ flat} \quad \vdash Sig \text{ flat}}{E_n \vdash h \text{ ok}} \quad \star$

Figure 1: Typing Rules – Type Environments, Hashes

$E \vdash K \text{ ok}$	
$\frac{E \vdash \text{ok}}{E \vdash \text{TYPE ok}}$	$\frac{E \vdash_{\emptyset} T : \text{TYPE}}{E \vdash \text{EQ}(T) \text{ ok}}$
$E \vdash_{eqs} K \approx K'$	
$\frac{E \vdash_{eqs} \text{ok}}{E \vdash_{eqs} \text{TYPE} \approx \text{TYPE}}$	$\frac{E \vdash_{eqs} T \approx T'}{E \vdash_{eqs} \text{EQ}(T) \approx \text{EQ}(T')}$
$E \vdash_{eqs} K <: K'$	
$\frac{E \vdash_{\emptyset} T : \text{TYPE} \quad E \vdash_{eqs} \text{ok}}{E \vdash_{eqs} \text{EQ}(T) <: \text{TYPE}}$	$\frac{E \vdash_{eqs} K \approx K'}{E \vdash_{eqs} K <: K'} \quad \text{trans is derivable}$

Figure 2: Typing Rules – Kinds

We define a metafunction <code>limitdom ()</code> that limits a signature to its abstract type fields as follows:	
<code>limitdom (type t_t : TYPE sig)</code>	<code>= type t_t : TYPE limitdom (sig)</code>
<code>limitdom (type t_t : EQ(T) sig)</code>	<code>= limitdom (sig)</code>
<code>limitdom (val x_x : T sig)</code>	<code>= limitdom (sig)</code>
<code>limitdom (empty)</code>	<code>= empty</code>
We define <code>t abstract in_{E_n} h</code> to hold if for some <code>t'</code> we have <code>(type t_{t'} : TYPE) ∈ Sig₀</code> where either <code>h = hash(hmodule_{eqs} M : Sig₀ version vne = Str)</code> , <code>h = n</code> and <code>(n : nmodule_{eqs} M : Sig₀ version vne = Str) ∈ E_n</code> , <code>h = hash(himport M : Sig₀ version vc like Str)</code> , or <code>h = n</code> and <code>(n : nimport M : Sig₀ version vc like Str) ∈ E_n</code> .	
<code>selffysig_X (type t_t : TYPE sig)</code>	<code>= (type t_t : EQ(X.t)) selffysig_X (sig)</code>
<code>selffysig_X (type t_t : EQ(T) sig)</code>	<code>= (type t_t : EQ(T)) selffysig_X (sig)</code>
<code>selffysig_X (val x_x : T sig)</code>	<code>= (val x_x : T) (selffysig_X (sig))</code>
<code>selffysig_X (empty)</code>	<code>= empty</code>
<code>selffysig_X (sig sig end)</code>	<code>= sig selffysig_X (sig) end</code>

Figure 3: Typing Rules – Auxiliaries




Figure 4: Typing Rules – Types, Type Equality

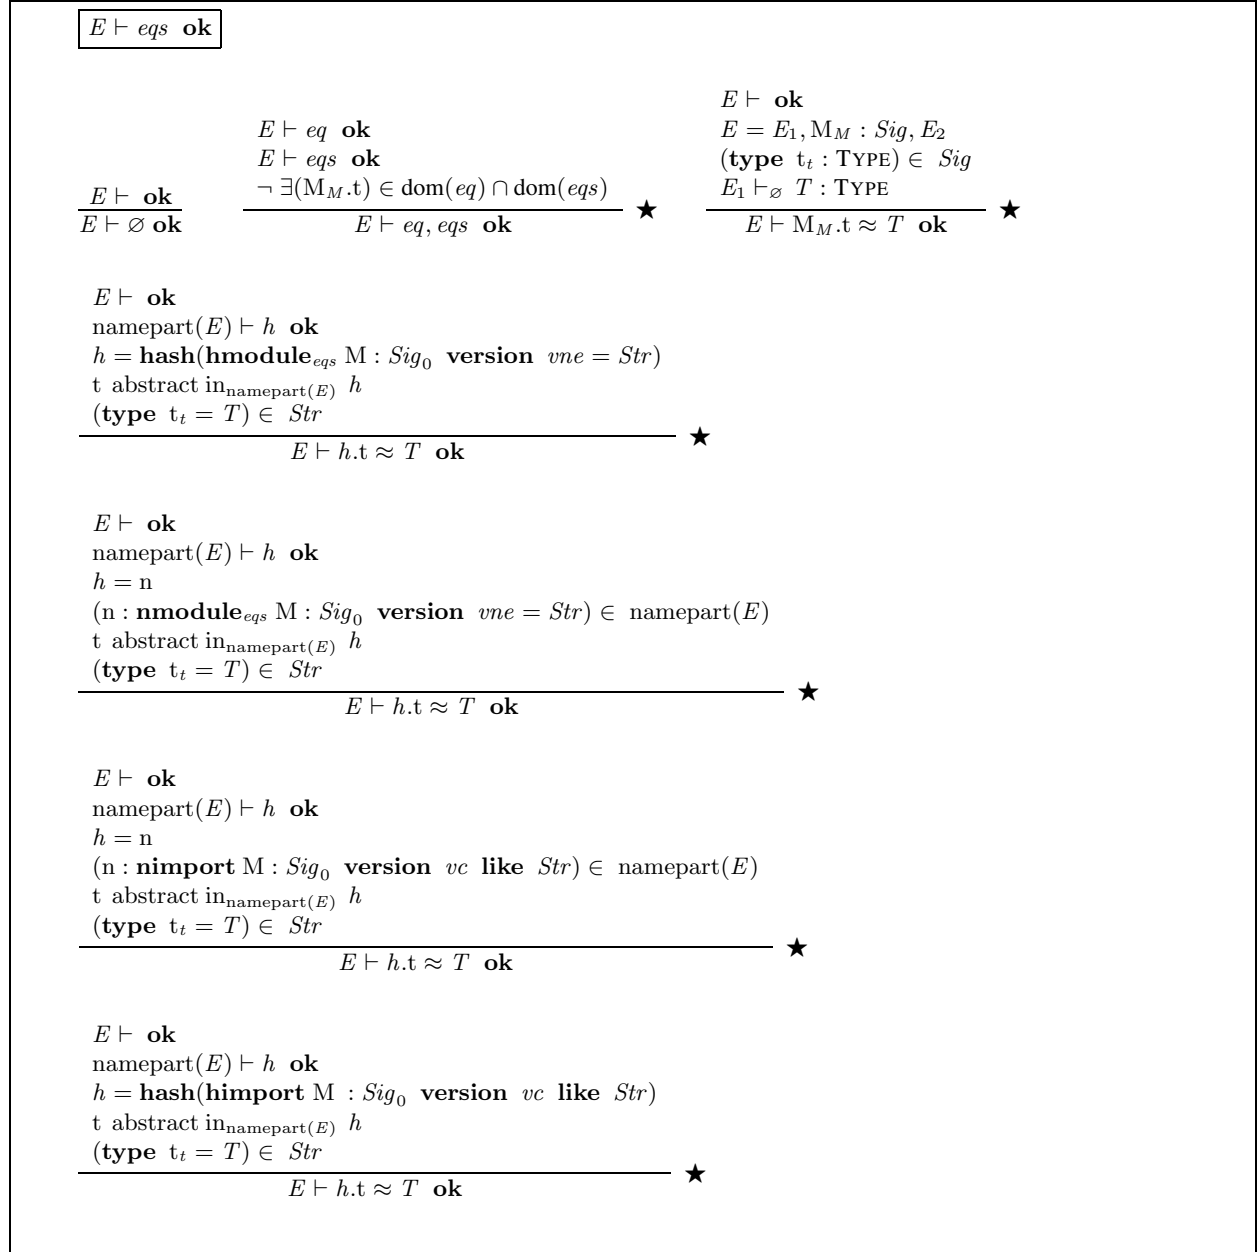


Figure 5: Typing Rules – Equation sets

$E \vdash sig \text{ ok}$		
$E \vdash \text{ok}$	$E, x : T \vdash sig \text{ ok}$ $x \notin \text{dom}(sig)$	$E, t : K \vdash sig \text{ ok}$ $t \notin \text{dom}(sig)$
$E \vdash \text{empty ok}$	$E \vdash \text{val } x_x : T \text{ sig ok}$	$E \vdash \text{type } t_t : K \text{ sig ok}$
$E \vdash_{eqs} sig <: sig'$		
$E \vdash_{eqs} \text{ok}$	$E \vdash_{eqs} T \approx T'$ $E, x : T \vdash_{eqs} sig <: sig'$ $x \notin \text{dom}(sig)$	$E \vdash_{eqs} K <: K'$ $E, t : K \vdash_{eqs} sig <: sig'$ $t \notin \text{dom}(sig)$
$E \vdash_{eqs} \text{empty} <: \text{empty}$	$E \vdash_{eqs} \text{val } x_x : T \text{ sig} <: \text{val } x_x : T' \text{ sig}'$	$E \vdash_{eqs} \text{type } t_t : K \text{ sig} <: \text{type } t_t : K' \text{ sig}'$
refl and trans are derivable		
$E \vdash_{eqs} sig \approx sig'$		
$E \vdash_{eqs} \text{ok}$	$E \vdash_{eqs} T \approx T'$ $E, x : T \vdash_{eqs} sig \approx sig'$ $x \notin \text{dom}(sig)$	$E \vdash_{eqs} K \approx K'$ $E, t : K \vdash_{eqs} sig \approx sig'$ $t \notin \text{dom}(sig)$
$E \vdash_{eqs} \text{empty} \approx \text{empty}$	$E \vdash_{eqs} \text{val } x_x : T \text{ sig} \approx \text{val } x_x : T' \text{ sig}'$	$E \vdash_{eqs} \text{type } t_t : K \text{ sig} \approx \text{type } t_t : K' \text{ sig}'$
refl and trans are derivable		
$E \vdash_{eqs} str : sig$		
$E \vdash_{eqs} \text{ok}$	$E, x : T \vdash_{eqs} str : sig$ $E \vdash_{eqs} v : T$ $x \notin \text{dom}(sig)$	$E, t : \text{EQ}(T) \vdash_{eqs} str : sig$ $E \vdash_{eqs} T : K$ $t \notin \text{dom}(sig)$
$E \vdash_{eqs} \text{empty} : \text{empty}$	$E \vdash_{eqs} \text{let } x_x = v \text{ str} : \text{val } x_x : T \text{ sig}$	$E \vdash_{eqs} \text{type } t_t = T \text{ str} : \text{type } t_t : K \text{ sig}$
$E \vdash_{eqs} T \approx T_1 \rightarrow \dots \rightarrow T_n \rightarrow T_0$ $E \vdash p_i : T_i \triangleright E_i \quad i = 1..n$ $E, x : T \vdash str : sig$ $E, E_1, \dots, E_n \vdash_{eqs} e : T_0$ $x \notin \text{dom}(sig)$ $n \geq 1$		
$E \vdash_{eqs} \text{let } x_x : T \text{ } p_1..p_n = e \text{ str} : \text{val } x_x : T \text{ sig}$		

Figure 6: Typing Rules – Signatures, Subsignaturing (part 1)




Figure 7: Typing Rules – Signatures, Subsignaturing (part 2)

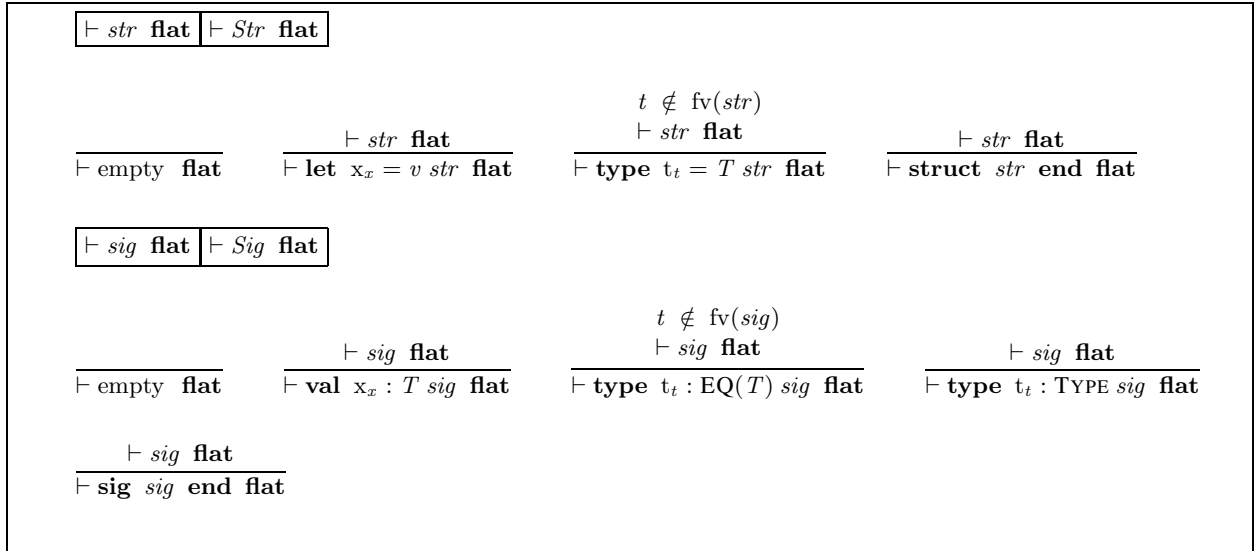


Figure 8: Typing Rules – flat predicates

$$E \vdash_{eqs} M_M : Sig$$

$$\frac{E_1, M_M : Sig, E_2 \vdash eqs \text{ ok}}{E_1, M_M : Sig, E_2 \vdash_{eqs} M_M : Sig}$$

$$\frac{E \vdash_{eqs} M_M : \mathbf{sig} \ sig_1 \ \mathbf{type} \ t_t : K \ sig_2 \ \mathbf{end}}{E \vdash_{eqs} M_M : \mathbf{sig} \ sig_1 \ \mathbf{type} \ t_t : EQ(M_M.t) \ sig_2 \ \mathbf{end}} \quad \frac{E \vdash_{eqs} M_M : Sig \quad E \vdash_{eqs} Sig <: Sig'}{E \vdash_{eqs} M_M : Sig'}$$

$$E \vdash_{eqs} h : Sig$$

$$\frac{h = \mathbf{hash}(\mathbf{hmodule}_{eqs'} M : Sig \ \mathbf{version} \ vne = Str) \quad \text{namepart}(E) \vdash h \ \mathbf{ok} \quad E \vdash eqs \ \mathbf{ok}}{E \vdash_{eqs} h : Sig} \star$$

$$\frac{h = n \quad (n : \mathbf{nmodule}_{eqs'} M : Sig \ \mathbf{version} \ vne = Str) \in \text{namepart}(E) \quad \text{namepart}(E) \vdash h \ \mathbf{ok} \quad E \vdash eqs \ \mathbf{ok}}{E \vdash_{eqs} h : Sig} \star$$

$$\frac{h = \mathbf{hash}(\mathbf{himport} M : Sig \ \mathbf{version} \ vc \ \mathbf{like} \ Str) \quad \text{namepart}(E) \vdash h \ \mathbf{ok} \quad E \vdash eqs \ \mathbf{ok}}{E \vdash_{eqs} h : Sig} \star$$

$$\frac{h = n \quad (n : \mathbf{nimport} M : Sig \ \mathbf{version} \ vc \ \mathbf{like} \ Str) \in \text{namepart}(E) \quad \text{namepart}(E) \vdash h \ \mathbf{ok} \quad E \vdash eqs \ \mathbf{ok}}{E \vdash_{eqs} h : Sig} \star$$

$$\frac{E \vdash_{eqs} h : \mathbf{sig} \ sig_1 \ \mathbf{type} \ t_t : K \ sig_2 \ \mathbf{end} \quad \text{t abstract in}_{\text{namepart}(E)} h}{E \vdash_{eqs} h : \mathbf{sig} \ sig_1 \ \mathbf{type} \ t_t : EQ(h.t) \ sig_2 \ \mathbf{end}} \star \quad \frac{E \vdash_{eqs} h : Sig \quad E \vdash_{eqs} Sig <: Sig'}{E \vdash_{eqs} h : Sig'}$$

Again h behaves much like M_M .
 For both this and M_M there is a stylistic choice as to how much selfification we do in one go; the rules deal with just a single field at a time. This judgement is wrt an E for uniformity (see the $h.x$ rule).

Figure 9: Typing Rules – Signatures of module identifiers and hashes

$E \vdash_{eqs} e : T$			
$E \vdash_{\emptyset} T : \text{TYPE}$	$E \vdash_{\emptyset} T \rightarrow T' : \text{TYPE}$	$E \vdash_{eqs} e_1 : T$	$E \vdash_{eqs} e_k : T_k \quad k \in 1..n$
$C_0 : T$	$C_1 : T \rightarrow T'$	$E \vdash_{eqs} e_2 : T \text{ list}$	$n \geq 2$
$E \vdash_{eqs} \text{ok}$	$E \vdash_{eqs} e : T$	$E \vdash_{eqs} e_1 :: e_2 : T \text{ list}$	$E \vdash_{eqs} (e_1, \dots, e_n) : T_1 * \dots * T_n$
$E \vdash_{eqs} C_0 : T$	$E \vdash_{eqs} C_1 e : T'$		
$x \notin \text{dom}(E_{\text{const}})$			
$E_1, x : T, E_2 \vdash_{eqs} \text{ok}$	$E \vdash_{eqs} M_M : \text{sig sig val } x_x : T \text{ sig}' \text{ end}$		
$E_1, x : T, E_2 \vdash_{eqs} x : T$	$E \vdash_{\emptyset} T : \text{TYPE}$	$E_1, l : T, E_2 \vdash_{eqs} \text{ok}$	
	$E \vdash_{eqs} M_M.x : T$	$E_1, l : T, E_2 \vdash_{eqs} l : T$	
$E \vdash_{eqs} e_1 : \text{bool}$			
$E \vdash_{eqs} e_2 : T$	$E \vdash_{eqs} e_1 : \text{bool}$	$E \vdash_{eqs} e_1 : \text{unit}$	
$E \vdash_{eqs} e_3 : T$	$E \vdash_{eqs} e_2 : \text{unit}$	$E \vdash_{eqs} e_2 : T$	
$E \vdash_{eqs} \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : T$	$E \vdash_{eqs} \text{while } e_1 \text{ do } e_2 \text{ done} : \text{unit}$	$E \vdash_{eqs} e_1 ; e_2 : T$	
$E \vdash_{eqs} e_1 : \text{bool}$			
$E \vdash_{eqs} e_2 : \text{bool}$			
$E \vdash_{eqs} e_1 \&\& e_2 : \text{bool}$	$E, x : T \vdash_{eqs} e : T'$	$E \vdash_{eqs} e_1 : T \rightarrow T'$	
$E \vdash_{eqs} e_1 \parallel e_2 : \text{bool}$	$E \vdash_{eqs} \text{function } (x : T) \rightarrow e : T \rightarrow T'$	$E \vdash_{eqs} e_2 : T$	
		$E \vdash_{eqs} e_1 e_2 : T'$	
$op^n : T_1 \rightarrow \dots \rightarrow T_n \rightarrow T$			
$E \vdash_{eqs} \text{ok}$			
$E \vdash_{eqs} e_j : T_j \quad j \in 1..n$			
$E \vdash_{eqs} op^n e_1 \dots e_n : T$ ★	$x^n \in \text{dom}(E_{\text{const}})$		
	$E_1, x^n : T', E_2 \vdash_{eqs} T' \approx T_1 \rightarrow \dots \rightarrow T_n \rightarrow T$		
	$E_1, x^n : T', E_2 \vdash_{eqs} e_j : T_j \quad j \in 1..n$		
	$E_1, x^n : T', E_2 \vdash_{eqs} x^n e_1 \dots e_n : T$ ★		
$E \vdash_{eqs} \text{ok}$			
$E \vdash_{\emptyset} e_0 : T_1 \rightarrow \dots \rightarrow T_n \rightarrow T$			
$E \vdash_{\emptyset} e_j : T_j \quad j \in 1..n$			
$E \vdash_{eqs} \text{op}(e_0)^n e_1 \dots e_n : T$ ★	$E \vdash_{\emptyset} T_1 + \dots + T_n : \text{TYPE}$		
	$E \vdash_{eqs} e : T_i$		
	$E \vdash_{eqs} \text{INJ}_i^{(T_1 + \dots + T_n)} e : T_1 + \dots + T_n$		
$E \vdash_{eqs} e : T$			
$E \vdash_{eqs} \text{mtch} : T \rightarrow T'$			
$E \vdash_{eqs} \text{match } e \text{ with } \text{mtch} : T'$	$E \vdash_{eqs} T_1 \approx T_2 \rightarrow T_3$		
	$E, x_1 : T_1, x_2 : T_2 \vdash_{eqs} e_3 : T_3$		
	$E, x_1 : T_1 \vdash_{eqs} e_4 : T_4$		
	$E \vdash_{eqs} \text{let rec } x_1 : T_1 = \text{function } (x_2 : T_2) \rightarrow e_3 \text{ in } e_4 : T_4$		
$E \vdash_{eqs} e : \text{exn}$			
$E \vdash_{eqs} \text{raise } e : T$	$E \vdash_{eqs} e : T$		$E \vdash_{\emptyset} T : \text{TYPE}$
	$E \vdash_{eqs} \text{mtch} : \text{exn} \rightarrow T$		$E \vdash_{eqs} \text{ok}$ ★
	$E \vdash_{eqs} \text{try } e \text{ with } \text{mtch} : T$		$E \vdash_{eqs} \text{RET}_T : T$
			$E \vdash_{eqs} \text{SLOWRET}_T : T$
$E \vdash_{eqs} e_1 : \text{string}$			
$E \vdash_{eqs} e_2 : T$			
$E \vdash_{eqs} \text{marshal } e_1 e_2 : T : \text{string}$ ★	$E \vdash_{eqs} e : \text{string}$		$E \vdash_{\emptyset} e : T$
	$E \vdash_{eqs} \text{unmarshal } e \text{ as } T : T$ ★		$E \vdash_{eqs} \text{marshalz } \underline{s} e : T : \text{string}$ ★
$E \vdash_{eqs} \text{ok}$			
$E \vdash_{\emptyset} T : \text{TYPE}$			
$E \vdash_{eqs'} e : T$			
$E \vdash_{eqs} [e]_{eqs'}^T : T$ ★	$E \vdash_{eqs} e : T$		
	$E \vdash_{eqs} T \approx T'$		
	$E \vdash_{eqs} e : T'$		

Figure 10: Typing Rules – Expressions (part 1)

$E \vdash_{eqs} e : T$ continued...		
$\frac{E \vdash_{eqs} e_1 : T \text{ ref} \quad E \vdash_{eqs} e_2 : T}{E \vdash_{eqs} e_1 :=_T e_2 : \text{unit}}$	$\frac{E \vdash_{eqs} e_1 : T \text{ ref} \quad E \vdash_{\emptyset} e_2 : T}{E \vdash_{eqs} e_1 :='_T e_2 : \text{unit}}$	$\frac{E \vdash_{eqs} e_1 : T \text{ ref}}{E \vdash_{eqs} !_T e_1 : T}$
$\frac{E \vdash_{eqs} h : \text{sig } sig \text{ val } x_x : T \text{ sig}' \text{ end} \quad E \vdash_{\emptyset} T : \text{TYPE}}{E \vdash_{eqs} h.x : T} \quad \star$	$\frac{E \vdash_{eqs} M'_{M'} : \text{Sig} \quad E \vdash_{eqs} M_M.x : T}{E \vdash_{eqs} \text{resolve}(M_M.x, M'_{M'}, \text{resolvespec}) : T} \quad \star$ $E \vdash_{eqs} \text{resolve_blocked}(M_M.x, M'_{M'}, \text{resolvespec}) : T$	
$\frac{E \vdash_{eqs} e_1 : T \text{ name} \quad E \vdash_{eqs} e_2 : T \text{ name} \quad E \vdash_{eqs} e_3 : T'}{E \vdash_{eqs} \text{swap } e_1 \text{ and } e_2 \text{ in } e_3 : T'}$	$\frac{E \vdash_{eqs} e_1 : T \text{ name} \quad E \vdash_{eqs} e_2 : T'}{E \vdash_{eqs} e_1 \text{ freshfor } e_2 : \text{bool}}$	$\frac{E \vdash_{\emptyset} T : \text{TYPE} \quad E \vdash_{eqs} e : T'}{E \vdash_{eqs} \text{support}_T e : T \text{ name list}}$
$\frac{E \vdash_{eqs} M_M.x : T}{E \vdash_{eqs} M_M@\text{x} : T \text{ tie}}$	$\frac{E \vdash_{eqs} e : T \text{ tie}}{E \vdash_{eqs} \text{name_of_tie } e : T \text{ name} \quad E \vdash_{eqs} \text{val_of_tie } e : T}$	$\frac{E \vdash_{eqs} e_1 : \text{unit} \quad E \vdash_{eqs} e_2 : T}{E \vdash_{eqs} e_1 e_2 : T}$
$\frac{E \vdash_{eqs} \text{ok} \quad E, t : \text{TYPE} \vdash_{eqs} e : T}{E \vdash_{eqs} \Lambda t \rightarrow e : \forall t. T}$	$\frac{E \vdash_{eqs} e : \forall t. T_1 \quad E \vdash_{\emptyset} T_2 : \text{TYPE}}{E \vdash_{eqs} e T_2 : \{T_2/t\} T_1}$	$\frac{E \vdash_{\emptyset} T_2 : \text{TYPE} \quad E \vdash_{eqs} e : \{T_2/t\} T_1}{E \vdash_{eqs} \{T_2, e\} \text{ as } \exists t. T_1 : \exists t. T_1}$
$\frac{E \vdash_{eqs} e_1 : \exists t. T \quad E, t : \text{TYPE}, x : T \vdash_{eqs} e_2 : T_2}{E \vdash \text{let } \{t, x\} = e_1 \text{ in } e_2 : T_2}$	$\frac{E \vdash_{eqs} e : T' \text{ name} \quad E \vdash_{eqs} e_1 : \exists t. t \text{ name} * T \quad E, t : \text{EQ}(T'), x_1 : T' \text{ name}, x_2 : T \vdash_{eqs} e_2 : T_2 \quad E \vdash_{eqs} e_3 : T_2}{E \vdash \text{namecase } e_1 \text{ with } \{t, (x_1, x_2)\} \text{ when } x_1 = e \rightarrow e_2 \text{ otherwise } \rightarrow e_3}$	
$\frac{E \vdash_{eqs} \text{ok} \quad E \vdash_{\emptyset} T : \text{TYPE}}{E \vdash_{eqs} \text{fresh}_T : T \text{ name}}$	$\frac{E \vdash_{eqs} \text{ok} \quad E \vdash_{\emptyset} T : \text{TYPE}}{E \vdash_{eqs} \text{cfresh}_T : T \text{ name}}$	$\frac{E \vdash_{eqs} X.x : T}{E \vdash_{eqs} \text{hash}(X.x)_T : T \text{ name}}$
$\frac{E \vdash_{\emptyset} T : \text{TYPE} \quad E \vdash_{eqs} e : \text{string}}{E \vdash_{eqs} \text{hash}(T, e)_T : T \text{ name}}$	$\frac{E \vdash_{\emptyset} T' : \text{TYPE} \quad E \vdash_{eqs} e_1 : \text{string} \quad E \vdash_{eqs} e_2 : T \text{ name}}{E \vdash_{eqs} \text{hash}(T', e_1, e_2)_{T'} : T' \text{ name}}$	$\frac{E_1, n : T \text{ name}, E_2 \vdash_{eqs} \text{ok}}{E_1, n : T \text{ name}, E_2 \vdash_{eqs} n_T : T \text{ name}}$

Figure 11: Typing Rules – Expressions (part 2)

$E \vdash_{eqs} e : T$ Sugared source forms	
$\frac{E \vdash_{eqs} mtch : T \rightarrow T'}{E \vdash_{eqs} \mathbf{function} \ mtch : T \rightarrow T'}$	
$\frac{E \vdash_{eqs} T_1 \approx T_2 \rightarrow T_3 \quad E, x : T_1 \vdash_{eqs} mtch : T_2 \rightarrow T_3 \quad E, x : T_1 \vdash_{eqs} e_4 : T_4}{E \vdash_{eqs} \mathbf{let} \ \mathbf{rec} \ x : T_1 = \mathbf{function} \ mtch \ \mathbf{in} \ e_4 : T_4}$	$\frac{E \vdash_{eqs} T \approx T_1 \rightarrow \dots \rightarrow T_n \rightarrow T' \quad E \vdash p_i : T_i \triangleright E_i \quad E, x : T, E_1, \dots, E_n \vdash_{eqs} e' : T' \quad E, x : T \vdash_{eqs} e'' : T''}{E \vdash_{eqs} \mathbf{let} \ \mathbf{rec} \ x : T \ p_1..p_n = e' \ \mathbf{in} \ e'' : T''}$
$\frac{E \vdash p : T_1 \triangleright E' \quad E \vdash_{eqs} e_1 : T_1 \quad E, E' \vdash_{eqs} e_2 : T_2}{E \vdash_{eqs} \mathbf{let} \ p = e_1 \ \mathbf{in} \ e_2 : T_2}$	$\frac{E \vdash_{eqs} T \approx T_1 \rightarrow \dots \rightarrow T_n \rightarrow T' \quad E \vdash p_i : T_i \triangleright E_i \quad i = 1..n \quad E, E_1, \dots, E_n \vdash_{eqs} e' : T' \quad E, x : T \vdash_{eqs} e'' : T''}{E \vdash_{eqs} \mathbf{let} \ x : T \ p_1..p_n = e' \ \mathbf{in} \ e'' : T''}$
$\frac{op^n : T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \quad E \vdash_{eqs} \mathbf{ok} \quad E \vdash_{eqs} e_j : T_j \quad j \in 1..k, k \in 0..n-1}{E \vdash_{eqs} op^n e_1 .. e_k : T_{k+1} \rightarrow \dots \rightarrow T_n \rightarrow T} \star$	$\frac{x^n \in \text{dom}(E_{\text{const}}) \quad E_1, x^n : T', E_2 \vdash_{eqs} T' \approx T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \quad E_1, x^n : T', E_2 \vdash_{eqs} e_j : T_j \quad j \in 1..k, k \in 0..n-1}{E_1, x^n : T', E_2 \vdash_{eqs} x^n e_1 .. e_k : T_{k+1} \rightarrow \dots \rightarrow T_n \rightarrow T} \star$

Figure 12: Typing Rules – Sugared Forms

$E \vdash p : T \triangleright E'$			
$\frac{E \vdash_{\emptyset} T : \text{TYPE}}{E \vdash (_ : T) : T \triangleright \text{empty}}$	$\frac{E \vdash_{\emptyset} T : \text{TYPE}}{E \vdash (x : T) : T \triangleright x : T}$	$\frac{C_0 : T \quad E \vdash \mathbf{ok}}{E \vdash C_0 : T \triangleright \text{empty}}$	$\frac{C_1 : T \rightarrow T' \quad E \vdash p : T \triangleright E'}{E \vdash C_1 p : T' \triangleright E'}$
$\frac{E \vdash p_1 : T \triangleright E_1 \quad E \vdash p_2 : T \ \text{list} \triangleright E_2}{E \vdash p_1 :: p_2 : T \ \text{list} \triangleright E_1, E_2}$	$\frac{E \vdash p_k : T_k \triangleright E_k \quad k \in 1..n \quad n \geq 2}{E \vdash (p_1, \dots, p_n) : T_1 * \dots * T_n \triangleright E_1, \dots, E_n}$	$\frac{E \vdash p : T \triangleright E'}{E \vdash (p : T) : T \triangleright E'}$	
$E \vdash_{eqs} mtch : T \rightarrow T'$			
$\frac{E \vdash p : T \triangleright E' \quad E, E' \vdash_{eqs} e : T'}{E \vdash_{eqs} p \rightarrow e : T \rightarrow T'}$	$\frac{E \vdash_{eqs} p \rightarrow e : T \rightarrow T' \quad E \vdash_{eqs} mtch : T \rightarrow T'}{E \vdash_{eqs} p \rightarrow e mtch : T \rightarrow T'}$		

Figure 13: Typing Rules – Patterns, Matches

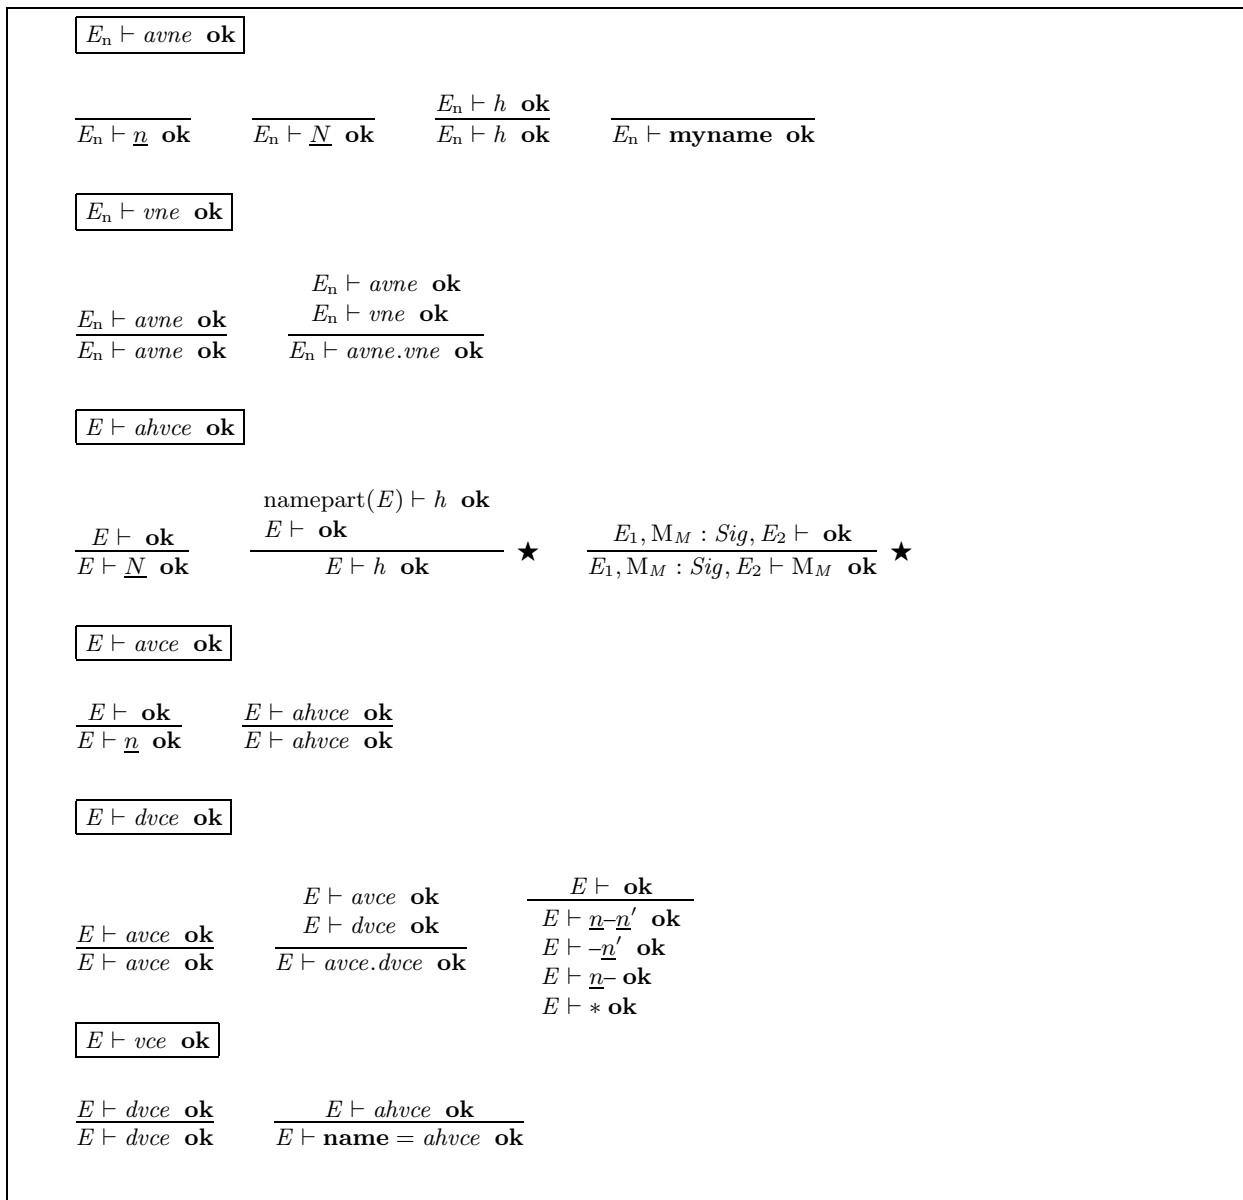


Figure 14: Typing Rules – Version number and constraint expressions

$E \vdash \text{likespec ok}$

$$\frac{E \vdash \text{ok}}{E \vdash \text{empty ok}}$$

$$\frac{E_1, M_M : \text{Sig}, E_2 \vdash \text{ok}}{E_1, M_M : \text{Sig}, E_2 \vdash \text{like } M_M \text{ ok}} \star$$

$$\frac{E \vdash_{\emptyset} \text{str} : \text{sig}}{E \vdash \text{like struct str end ok}} \star$$

$E \vdash = Mo : \text{Sig}$

$$\frac{E \vdash \text{Sig ok}}{E \vdash = \text{UNLINKED} : \text{Sig}}$$

$$\frac{E_1, M_M : \text{Sig}, E_2 \vdash \text{ok}}{E_1, M_M : \text{Sig}, E_2 \vdash = M_M : \text{Sig}}$$

$$\frac{E \vdash_{\emptyset} \text{Sig} <: \text{Sig}' \quad E \vdash = M_M : \text{Sig}}{E \vdash = M_M : \text{Sig}'}$$

Figure 15: Typing Rules – Definition auxiliaries

$E \vdash \text{sourcedefinition} \triangleright E'$

$$\frac{E \vdash \text{vne ok} \quad E \vdash_{\text{weqs}} \text{Str} : \text{Sig}}{E \vdash \text{module } M_M : \text{Sig} \text{ version } \text{vne} = \text{Str} \text{ with !weqs} \triangleright M_M : \text{Sig}} \star$$

$$\frac{E \vdash \text{vce ok} \quad E \vdash \text{likespec ok} \quad E \vdash = Mo : \text{Sig}}{E \vdash \text{import } M_M : \text{Sig} \text{ version } \text{vce likespec by resolvespec} = Mo \triangleright M_M : \text{Sig}} \star$$

We could additionally check that for all abstract type fields in *Sig* there is a corresponding type in an in-line *likestr*, or a type (which could reasonably be required to be abstract) in an $M'_{M'}$ *likestr*. At present this is left to compilation.

$$\frac{E \vdash \text{ok}}{E \vdash \text{mark MK} \triangleright \text{empty}} \star$$

$$\frac{E_1, M'_{M'} : \text{Sig}, E_2 \vdash \text{ok} \quad \text{Sig}' = \text{selffysig}_{M'_{M'}}(\text{Sig})}{E_1, M'_{M'} : \text{Sig}, E_2 \vdash \text{module } M_M : \text{Sig}' = M'_{M'} \triangleright M_M : \text{Sig}'} \star$$

Note that we have to selffy in the alias rule to avoid introducing new abstract types. We do not allow subsignaturing as we do not want to think about sealing here. We could allow sig equality.

(selffysig () is defined on page 93.)

Figure 16: Typing Rules – Source Definitions

16.3.2 Typing for Compiled and Executing Forms

$E \vdash \text{definition} \triangleright E'$

$\text{namepart}(E) \vdash h \text{ ok}$
 $\text{namepart}(E) \vdash \text{eqs} \text{ ok}$
 $\text{namepart}(E) \vdash \text{Sig}_0 \text{ ok}$
 $\text{eqs}_0 = \text{eqs_of_sign_str}(h, \text{Sig}_0, \text{Str})$
 $\text{Sig}_1 = \text{typeflattensig}(\text{selffysig}_h(\text{Sig}_0))$
 $E \vdash_{\text{eqs}_0, \text{eqs}} \text{Str} : \text{Sig}_1$
 $\vdash \text{Str} \text{ flat}$
 $\text{compiledform}(\text{eqs}, \text{Sig}_0, \text{Sig}_1, \text{Str})$

$E \vdash \text{cmodule}_{h; \text{eqs}; \text{Sig}_0} \text{ vubs } M_M : \text{Sig}_1 \text{ version } \text{vn} = \text{Str} \triangleright M_M : \text{Sig}_1 \quad \star$

The Sig_1 is now computable from the h and Sig_0 . We keep it in the data for the time being, however, as it has a clear conceptual role.

An alternative rule here (corresponding to that for user modules above) would have $E \vdash_{\text{eqs}_0} \text{Str} : \text{Sig}'_1$ and $E \vdash_{\text{eqs}_0, \text{eqs}} \text{Sig}'_1 \approx \text{Sig}_1$.

In typing compiled and hashed modules there is a stylistic choice as to how many of the properties that compilation establishes are captured in the typing rules. Here we choose to be rather tight, at the cost of some baroqueeness. Note that in the $E \vdash_{\text{eqs}} \text{Str} : \text{Sig}$ premise the E allows term components of earlier modules to be used (as is required), but also allows type components to be used. We prevent the latter with the $\text{compiledform}(\dots)$ premise, as they have been hashified by compilation.

$\text{namepart}(E) \vdash h \text{ ok}$
 $\text{Sig}_1 = \text{typeflattensig}(\text{selffysig}_h(\text{Sig}_0))$
 $\text{namepart}(E) \vdash \text{Sig}_0 \text{ ok}$
 $\text{namepart}(E) \vdash_{\emptyset} \text{Str} : \text{limitdom}(\text{Sig}_0)$
 $\vdash \text{Str} \text{ flat}$
 $\text{compiledform}(\text{Sig}_0, \text{Sig}_1, \text{Str})$
 $E \vdash = Mo : \text{Sig}_0$

$E \vdash \text{cimport}_{h; \text{Sig}_0} \text{ vubs } M_M : \text{Sig}_1 \text{ version } \text{vc} \text{ like } \text{Str} \text{ by } \text{resolvespec} = Mo \triangleright M_M : \text{Sig}_1 \quad \star$

Note that compilation has cut down the Str in likespec . The fact that this ensures it doesn't include any code (or extraneous types) legitimizes the empty. This rule does not explicitly check type coherence between the likespec and the Mo implementation, as the latter is not available in E , but note that Sig_1 is hashified.

$\frac{E \vdash \text{ok}}{E \vdash \text{mark MK} \triangleright \text{empty}}$

$\text{definition} = \text{module fresh } M_M : \text{Sig} \text{ version } \text{vne} = \text{Str} \text{ with !weqs}$
 no cfresh_T
 $E \vdash \text{vne} \text{ ok}$
 $E \vdash_{\text{weqs}} \text{Str} : \text{Sig}$

$E \vdash \text{definition} \triangleright M_M : \text{Sig} \quad \star$

$\text{definition} = \text{import fresh } M_M : \text{Sig} \text{ version } \text{vce} \text{ likespec} \text{ by } \text{resolvespec} = Mo$
 $E \vdash \text{vce} \text{ ok}$
 $E \vdash \text{likespec} \text{ ok}$
 $E \vdash = Mo : \text{Sig}$

$E \vdash \text{definition} \triangleright M_M : \text{Sig} \quad \star$

where $\text{eqs_of_sign_str}(h, \text{Sig}_0, \text{Str}) = \{h.t \approx T \mid \exists t. (\text{type } t_t : \text{TYPE} \in \text{Sig}_0 \wedge (\text{type } t_t = T) \in \text{Str})\}$

Figure 17: Typing Rules – Compiled Definitions

Define linkok ($E_n, definition', definition$) if

1. $definition$ is of the form **cimport** _{$h;Sig_0$} $vubs M_M : Sig_1$ **version** vc **like** Str **by** $resolvespec = Mo$
2. $definition'$ is either a **cmodule** or a **cimport** as below.

cmodule _{$h';eqs';Sig'_0$} $vubs' M'_{M'} : Sig'_1$ **version** $vn' = Str'$
cimport _{$h';Sig'_0$} $vubs' M'_{M'} : Sig'_1$ **version** vc' **like** Str' **by** $resolvespec' = Mo'$

3. the external names match: $M' = M$.
It is unclear whether we always want to require the above.
4. the interfaces match: $E_n \vdash_{\emptyset} Sig'_0 <: Sig_0$. In an implementation, we check only syntacticssubsig $Sig'_0 Sig_0$.
5. the versions match:
 - Case: the vc is not an exact-name constraint, i.e. $vc = dvc$ for some dvc . If $definition'$ is a **cmodule** check $vn' \in dvc$, otherwise if $definition'$ is a **cimport** check $vc' \subseteq vc$.
 - Case: the vc is an exact-name constraint, i.e. $name = ahvc$ for some $ahvc$. Check $h' \cong ahvc$.
6. the representation types match: $\forall(\mathbf{type} \ t_t = T) \in Str. \exists t', T'. (\mathbf{type} \ t_{t'} = T') \in Str' \wedge T = T'$.

Define linkok ($E_n, definitions$) if whenever

$definitions = definitions_1 ;; definition ;; definitions_2$
 $definition = \mathbf{cimport}_{h;Sig_0} vubs M_M : Sig_1$ **version** vc **like** Str **by** $resolvespec = M'_{M'}$

there exists a $definition'$ for $M'_{M'}$ in $definitions_1$ with linkok ($E_n, definition', definition$).

Define (on flattened signatures only) syntacticssubsig $Sig'_0 Sig_0$, a weak version of $E_n \vdash_{\emptyset} Sig'_0 <: Sig_0$:

$$\frac{T' = T \quad \text{syntacticssubsig } sig' sig}{\text{syntacticssubsig } (\mathbf{val} \ x_x : T' sig') (\mathbf{val} \ x_x : T sig)} \quad \frac{\text{syntacticssubsig } sig' sig}{\text{syntacticssubsig } (\mathbf{sig} \ sig' \ \mathbf{end}) (\mathbf{sig} \ sig \ \mathbf{end})}$$

$$\frac{\text{syntacticssubsig } sig' sig}{\text{syntacticssubsig } (\mathbf{type} \ t_t : \mathbf{TYPE} \ sig') (\mathbf{type} \ t_t : \mathbf{TYPE} \ sig)} \quad \frac{}{\text{syntacticssubsig } \mathbf{empty} \ \mathbf{empty}}$$

$$\frac{T' = T \quad \text{syntacticssubsig } sig' sig}{\text{syntacticssubsig } (\mathbf{type} \ t_t : \mathbf{EQ}(T') \ sig') (\mathbf{type} \ t_t : \mathbf{EQ}(T) \ sig)}$$

$$\frac{\text{syntacticssubsig } sig' (\{T'/t\}sig)}{\text{syntacticssubsig } (\mathbf{type} \ t_t : \mathbf{EQ}(T') \ sig') (\mathbf{type} \ t_t : \mathbf{TYPE} \ sig)}$$

Comment: The substitution is required in the case of a concrete type on the left and an abstract on the right, in order that inequalities such as the following are treated correctly: **sig type** $t = \mathbf{int} \ \mathbf{type} \ u = \mathbf{int} \ \mathbf{end} <: \mathbf{sig} \ \mathbf{type} \ t \ \mathbf{type} \ u = t \ \mathbf{end}$. There is no need to apply the substitution in the other concrete case, because compilation has already flattened Sig'_0 and Sig_0 . An implementation may avoid the type substitution by carrying a type environment. We don't check that the external value and type names are distinct, since compilation has already ensured that both signatures are well-formed.

Comment: Note that an implementation need not refer to E_n while doing the subsignature check. The E_n is required only to provide type equalities $n.t : \mathbf{EQ}(T)$ for concrete type fields of freshly-named modules. But hashify (step 3, page 119) has ensured that all concrete type fields in Sig_0 have already been substituted out. Thus only abstract type field references remain; and in this case it is sufficient to assume $n.t : \mathbf{TYPE}$ for all n, t . The same reasoning confirms that an implementation need not inspect the body of a **hash**(..)-form **hash** h .

Comment: Note that this does not permit a **cimport** to be linked to a **module** **fresh**. This is slightly unpleasant from the user's point of view, though **use** imports can usually be written with a **HERE_ALREADY** **resolvespec**. The restriction avoids the need to check subsignature or version of a **module** **fresh**, which (as they have no name) is problematic.

Figure 18: Link Checking

$$E ; E_s \vdash s \text{ store}$$

$\text{dom}(E_s)$ contains only locations
 $l \in \text{dom}(E_s) \iff l \in \text{dom}(s)$
 $\forall l : T \text{ ref} \in E_s.E, E_s \vdash_{\emptyset} s(l) : T \wedge \text{compiledform}(s(l))$

$$E ; E_s \vdash s \text{ store} \quad \star$$

$$E \vdash \text{definitions} \triangleright E'$$

$E \vdash \text{definition} \triangleright E'$
 $E, E' \vdash \text{definitions} \triangleright E''$
 $\text{dom}(E'), \text{dom}(E'') \text{ disjoint}$

$E \vdash \text{ok}$	$E \vdash \text{definition} \;; \text{definitions} \triangleright E', E''$
$E \vdash \text{empty} \triangleright \text{empty}$	

$$E \vdash \text{definitions eo ok}$$

$E \vdash \text{definitions} \triangleright E'$	$E \vdash \text{definitions} \triangleright E'$ $E, E' \vdash e : T$
$E \vdash \text{definitions ok}$	$E \vdash \text{definitions e ok}$

$$\vdash E_n ; \langle E_s, s, \text{definitions}, e \rangle : T$$

$E_n, E_{\text{const}}, E_s \vdash \text{definitions} \triangleright E$
 $\text{linkok}(E_n, \text{definitions})$
 $E_n, E_{\text{const}}, E ; E_s \vdash s \text{ store}$
 $E_n, E_{\text{const}}, E, E_s \vdash_{\emptyset} e : T$
 $\text{compiledform}(e)$

$$\vdash E_n ; \langle E_s, s, \text{definitions}, e \rangle : T \quad \star$$

$$\vdash E_n ; \langle E_s, s, \text{definitions}, P \rangle : T$$

$E_n, E_{\text{const}}, E_s \vdash \text{definitions} \triangleright E$
 $\text{linkok}(E_n, \text{definitions})$
 $E_n, E_{\text{const}}, E ; E_s \vdash s \text{ store}$
 $E_n, E_{\text{const}}, E, E_s \vdash P \text{ ok}$
 $\text{compiledform}(e)$

$$\vdash E_n ; \langle E_s, s, \text{definitions}, P \rangle : \text{unit} \quad \star$$

Figure 19: Typing Rules – Store, Configurations




Figure 20: Typing Rules – Marshalled Values

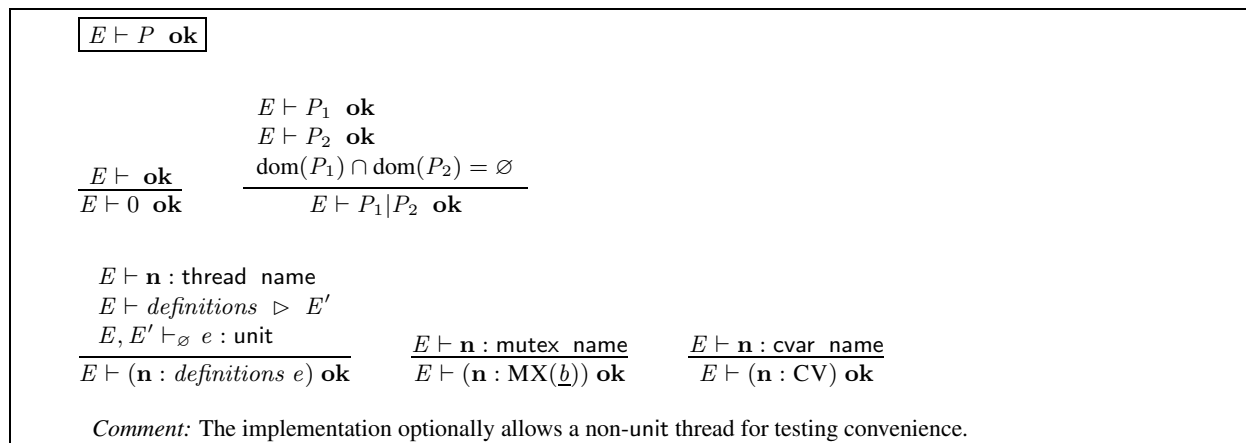


Figure 21: Typing Rules – Processes

16.4 Typed Desugaring

Desugaring is a function that takes a source expression and yields a source expression with all sugared forms eliminated. It is applied after type inference and typechecking and before hashification. Thus we assume that type annotations have been inserted that did not appear in user source programs.

Desugaring behaves as follows:

```

desugar( $e_1 ||| e_2$ )
  = create.thread freshthread (function () → desugar( $e_1$ )) () ; desugar( $e_2$ )
desugar(function  $mtch$ )
  = function ( $x : T$ ) → match  $x$  with desugarmtch( $mtch$ )
    where  $mtch \neq (x' : T' \rightarrow e)$ ,  $x$  fresh,  $T = \text{matchty}(mtch)$ 
desugar(fun  $p_1..p_n \rightarrow e'$ )
  = (function ( $x_1 : T_1$ ) → match  $x_1$  with  $p_1 \rightarrow ..$ 
    function ( $x_n : T_n$ ) → match  $x_n$  with  $p_n \rightarrow \text{desugar}(e')$ )
    where  $n \geq 1$ ,  $T_i = \text{patty}(p_i)$  for  $1 \leq i \leq n$ , and  $x_1..x_n$  fresh
desugar(let  $p = e_1$  in  $e_2$ )
  = match desugar( $e_1$ ) with  $p \rightarrow \text{desugar}(e_2)$ 
desugar(let  $x : T$   $p_1..p_n = e'$  in  $e''$ )
  = match (function ( $x_1 : T_1$ ) → match  $x_1$  with  $p_1 \rightarrow ..$ 
    function ( $x_n : T_n$ ) → match  $x_n$  with  $p_n \rightarrow \text{desugar}(e')$ )
    with ( $x : T$ ) → desugar( $e''$ )
    where  $n \geq 1$ ,  $T_i = \text{patty}(p_i)$  for  $1 \leq i \leq n$ , and  $x_1..x_n$  fresh
desugar(let rec  $x : T = \text{function } mtch$  in  $e$ )
  = let rec  $x : T = \text{function } (x' : T') \rightarrow (\text{match } x' \text{ with desugarmtch}(mtch))$  in desugar( $e$ )
    where  $mtch \neq (x'' : T'' \rightarrow e')$ ,  $x'$  fresh,  $T' = \text{matchty}(mtch)$ 
desugar(let rec  $x : T$   $p_1..p_n = e'$  in  $e''$ )
  = let rec  $x : T = (\text{function } (x_1 : T_1) \rightarrow \text{match } x_1 \text{ with } p_1 \rightarrow ..$ 
    function } (x_n : T_n) \rightarrow \text{match } x_n \text{ with } p_n \rightarrow \text{desugar}(e'))
    in desugar( $e''$ )
    where  $n \geq 1$ ,  $T_i = \text{patty}(p_i)$  for  $1 \leq i \leq n$ , and  $x_1..x_n$  fresh
desugarmtch( $p_1 \rightarrow e_1 | .. | p_n \rightarrow e_n$ )
  =  $p_1 \rightarrow \text{desugar}(e_1) | .. | p_n \rightarrow \text{desugar}(e_n)$ 

```

Desugaring of structure items:

```

desugar(let  $x_x : T$   $p_1..p_n = e'$ )
  = let  $x_x : T = \text{function } (x_1 : T_1) \rightarrow \text{match } x_1 \text{ with } p_1 \rightarrow ..$ 
    function } (x_n : T_n) \rightarrow \text{match } x_n \text{ with } p_n \rightarrow \text{desugar}(e'))
    where  $n \geq 1$ ,  $T_i = \text{patty}(p_i)$  for  $1 \leq i \leq n$ , and  $x_1..x_n$  fresh

```

Desugar also saturates operators, by performing eta-expansions.

```

desugar( $e e_1 .. e_m$ )
  = function ( $x_{m+1} : T_{m+1}$ ) → .. → function ( $x_n : T_n$ ) → ( $e e_1 .. e_m x_{m+1} .. x_n$ )
    where  $e : T_1 \rightarrow \dots \rightarrow T_n \rightarrow T_0$  and  $m < n$ 

```

for e an *op* or $e \in E_{\text{const}}$.

Likewise, we saturate partial applications of $(:=_T) : T \text{ ref} \rightarrow T \rightarrow \text{unit}$ and $(!_T) : T \text{ ref} \rightarrow T$ and $(\&\&) : \text{bool} \rightarrow \text{bool} \rightarrow \text{bool}$ and $(||) : \text{bool} \rightarrow \text{bool} \rightarrow \text{bool}$.

In all other cases, it merely descends recursively into the term without altering its structure. In particular,

$$\text{desugar}(\mathbf{function} (x : T) \rightarrow e) = \mathbf{function} (x : T) \rightarrow \text{desugar}(e)$$

Comment: We diverge from Ocaml in alternating between “function” and “match”; the expression ‘let f (Some x) (y:int) = x+1 in f None’ raises MATCH_FAILURE in Acute while Ocaml returns a thunk.

desugar(.) makes use of a function matchty (.) that determines the argument type of a match:

matchty ($p \rightarrow e mtch$)	=	matchty ($p \rightarrow e$)	
matchty ($p \rightarrow e$)	=	patty (p)	
patty ($_ : T$)	=	T	
patty ($x : T$)	=	T	
patty (C_0)	=	T	where $C_0 : T$
patty ($C_1 p$)	=	T'	where $C_1 : T \rightarrow T'$
patty ($p_1 :: p_2$)	=	patty (p_1) list	
patty (p_1, \dots, p_n)	=	patty (p_1) * .. * patty (p_n)	
patty ($p : T$)	=	T	

16.5 Errors

The possible outcomes of compilation and execution are as follows:

SUCCESS

DEADLOCK

LIBRARY

FAILURE

 COMPILE

 LEX

 PARSE

 TYPE

 INCLUDE

 CYCLE

 SYSTEM_ERROR

 NONFINAL_EXPRESSION

 HASHIFY

 WITHSPEC_EQUATION_FROM_IMPORT

 WITHSPEC_WRT_BAD_TYPE_FIELD

 WITHSPEC_TYPES_NOT_EQUAL

 LIKESPEC_MISSING_TYPE_FIELDS

 LINKOK_NOT

 BAD_SOURCEDEF_VALUABILITY

 NENV_MERGE_OF_COMPILEDUNIT

 TYPECHECK_OF_COMPILEDUNIT

 RUNTIME_MISMATCH

 RUN

 TYPECHECK_OF_CONFIGURATION

 TYPECHECK_ON_MARSHAL

 TYPECHECK_ON_UNMARSHAL

 TYPECHECK_ON_GET_URI

 INTERNAL

 NEVER_HAPPEN

 STUCK

 UNIMPLEMENTED

with the appropriate additional data (respectively a configuration, a set of definitions , or further information about the failure).

Of these, FAILURE.COMPILE.TYPECHECK_OF_COMPILEDUNIT, FAILURE.RUN.TYPECHECK_OF_CONFIGURATION, FAILURE.RUN.TYPECHECK_ON_UNMARSHAL, and FAILURE.RUN.TYPECHECK_ON_GET_URI, are signalled if a compile-time compiled module typecheck or a run-time typecheck fails. They should never happen (assuming that marshalled values and compiled files are not forged), and in a production implementation using numeric hashes these typechecks cannot be performed. They are therefore not currently mapped to internal Acute exceptions.

The FAILURE.INTERNAL.* errors should never happen.

The `FAILURE.COMPILE.RUNTIME_MISMATCH` error occurs when we try and parse something (a marshalled value, an included file, or an imported file) which was created with a runtime incompatible with the current one — for example one created using literal hashing when we are using structured hashing. If a `resolvespec` meets this error Acute fails immediately rather than proceeding to the next `resolvespec`.

Only some kinds of deadlock are detected by Acute.

Individual threads may exit cleanly, be killed, or raise an exception; none of these (in themselves) result in program termination, although they may cause a debugging message to be written to the console.

16.6 Valuability helper functions

Comment: These definitions should be disentangled.

We define `find_valuabilities(definitions, MM)` which looks up the valuabilities `vubs` of `MM`; `check_valuability_expr(definitions, e, vub)` which checks whether `e` can have valuability `vub`; and `derive_valuabilities(definitions, sourcedefinition)` which calculates the valuabilities `vubs` of the `sourcedefinition` for `sourcedefinition` not of the form `mark MK`.

As usual, we conflate the `definitions` with a finite map `C` from module names to definitions.

First, we define `find_valuabilities(definitions, MM)` as follows.

To do this, we perform case analysis on `C(MM)` to calculate `vubs'`:

- Case `cmoduleh;eqs;Sig0 vubs MM : Sig1 version vn = Str`: let `vubs' = vubs`.
- Case `cimporth;Sig0 vubs MM : Sig1 version vc like Str by resolvespec = Mo`: let `vubs' = vubs`.
- Case `module fresh MM : Sig version vne = Str withspec`: let `vubs' = (nonvaluable, nonvaluable)`.
- Case `import fresh MM : Sig version vce likespec by resolvespec = Mo`: let `vubs' = (nonvaluable, nonvaluable)`.

Now we define `check_valuability_expr(definitions, e, vub)`, which checks where `e` can have valuability `vub`.

Comment: The valuabilities are linearly ordered, with valuable implying `cvaluable` and `cvaluable` implying `nonvaluable`.

Say a *cval context* is a linear expression context of the grammar defined from the v^\emptyset grammar by taking all clauses with a sub- v^\emptyset metavariable, replacing one by a `_` and all others by an `e`, i.e. :

```

CVAL ::= _
        C1 CVAL
        CVAL :: e
        e :: CVAL
        (CVAL, ..., en) .. (e1, ..., CVAL)
        hash(T, CVAL)T'
        hash(CVAL, T, e)T'
        hash(e, T, CVAL)T'
        {T, CVAL} as T'

```

The `cval` contexts are used in the following to limit occurrences of `cfresh` to unguarded positions.

We perform case analysis on `vub`:

- Case `vub = valuable`: check that
 - `e` is a term of the grammar consisting of
 - * the clauses of the v^\emptyset grammar
 - * together with `MM.x`
 - * together with `MM@x`
 - * together with `x` (such as might occur if the expression is in a module structure and refers to an earlier field)

(This list allows `MM.x`, etc., to occur in unguarded positions, which is not allowed if one confines `e` to just the value grammar.)

- for all $M_M.x$ and $M_M@x$ occurring in e we have $\text{fst}(\text{find_valuabilities}(\text{definitions}, M_M)) = \text{valuable}$;
- for all $M_M.t$ occurring in e we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M_M)) = \text{valuable}$.
- there are no occurrences of **cfresh_T**.
- Case $vub = \text{cvaluable}$: check that
 - e is a term of the grammar consisting of the clauses of
 - * the v^\emptyset grammar
 - * together with **cfresh_T** where all occurrences of **cfresh_T** give a decomposition of e of the form CVAL.cfresh_T (i.e. all are in unguarded positions)
 - * together with $M_M.x$
 - * together with $M_M@x$
 - * together with x (i.e. an earlier field)
 - for all $M_M.x$ and $M_M@x$ occurring in e we have $\text{fst}(\text{find_valuabilities}(\text{definitions}, M_M)) \in \{\text{valuable}, \text{cvaluable}\}$;
 - for all $M_M.t$ occurring in e we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M_M)) \in \{\text{valuable}, \text{cvaluable}\}$.
- Case $vub = \text{nonvaluable}$: check that all occurrences of **cfresh_T** give a decomposition of e of the form CVAL.cfresh_T .

Comment: Note that we impose conditions on the valuability of *all* $M_M.x$ and $M_M@x$ occurring in e , not just the ones in unguarded positions, since we need to be sure that we can replace these by an appropriate $h.x$ in hashification during compilation.

Now we define $\text{derive_valuabilities}(\text{definitions}, \text{sourcedefinition})$ by case analysis on sourcedefinition :

- Case $\text{sourcedefinition} = (\text{module mode } M_M : \text{Sig } \text{version } vne = \text{Str withspec})$. We consider several subcases:
 - Case $\text{mode} = \text{hash}$: Check for all e on the rhs of Str we have $\text{check_valuability_expr}(\text{definitions}, e, \text{valuable})$. Check that for all $M'_{M'}.t$ occurring anywhere we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M'_{M'})) = \text{valuable}$. The result is $(\text{valuable}, \text{valuable})$.
 - Case $\text{mode} = \text{hash!}$: Check for all $M_M.x$ and $M_M@x$ occurring in an e on the rhs of Str we have $\text{fst}(\text{find_valuabilities}(\text{definitions}, M_M)) = \text{valuable}$. Check that for all $M'_{M'}.t$ occurring anywhere we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M'_{M'})) = \text{valuable}$. The result is $(\text{valuable}, \text{valuable})$.
 - Case $\text{mode} = \text{cfresh}$: Check for all e on the rhs of Str we have $\text{check_valuability_expr}(\text{definitions}, e, \text{cvaluable})$. Check that for all $M'_{M'}.t$ occurring anywhere we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M'_{M'})) \in \{\text{valuable}, \text{cvaluable}\}$.
The result is $(\text{cvaluable}, \text{cvaluable})$.
 - Case $\text{mode} = \text{cfresh!}$: Check for all e on the rhs of Str : (a) for all $M_M.x$ and $M_M@x$ occurring in e we have $\text{fst}(\text{find_valuabilities}(\text{definitions}, M_M)) \in \{\text{valuable}, \text{cvaluable}\}$; and (b) all occurrences of **cfresh_T** give a decomposition of e of the form CVAL.cfresh_T (i.e. all are in unguarded positions). Check that for all $M'_{M'}.t$ occurring anywhere we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M'_{M'})) \in \{\text{valuable}, \text{cvaluable}\}$. The result is $(\text{cvaluable}, \text{cvaluable})$.
 - Case $\text{mode} = \text{fresh}$: Check for all e on the rhs of Str we have $\text{check_valuability_expr}(\text{definitions}, e, \text{nonvaluable})$. The result is $(\text{nonvaluable}, \text{nonvaluable})$.

Comment: For the **hash!** case we do not ignore valuability checking altogether, as to build (at compile-time) a hash for this module requires names for any referenced modules and imports. There is an alternative

possibility, deferring the hash construction until run-time if necessary, but it seems that that would be confusing (too different from the `hash` semantics). We do similarly for the `cfresh!` case.

- Case `sourcedefinition = (import mode MM : Sig version vce likespec by resolvespec = Mo)`:

Comment: The valuabilities of $M''_{M''}$, where $Mo = M''_{M''}$, are unimportant.

- Case `mode = hash` and `mode = hash!`: We perform case analysis on the `likespec`:

- * Case `likespec = empty`: `true`.

- * Case `likespec = like Str`: Check that for all $M'_{M'}.t$ occurring anywhere in those fields of `Str` that are in the domain of `limitdom (Sig)` we have `snd(find_valuabilities(definitions, M'_{M'})) ∈ valuable`.

- * Case `likespec = like M'_{M'}`: check `snd(find_valuabilities(definitions, M'_{M'})) = valuable`

The result is (valuable, valuable).

- Case `mode = cfresh` and `mode = cfresh!`: We perform case analysis on the `likespec`:

- * Case `likespec = empty`: `true`.

- * Case `likespec = like Str`: Check that for all $M'_{M'}.t$ occurring anywhere in those fields of `Str` that are in the domain of `limitdom (Sig)` we have `snd(find_valuabilities(definitions, M'_{M'})) ∈ {valuable, cvaluable}`.

- * Case `likespec = like M'_{M'}`: check `snd(find_valuabilities(definitions, M'_{M'})) ∈ {valuable, cvaluable}`.

The result is (cvaluable, cvaluable)

- Case `mode = fresh`: The result is (nonvaluable, nonvaluable).

Comment: We used to regard expression projections from an import as always nonvaluable (as there is no unique value they are guaranteed to reduce to, in the presence of rebinding, except in the exact-name version constraint case). Now, we regard the expression and type valuabilities as the same, and so could return to a single valuability rather than a pair throughout.

- Case `sourcedefinition = (module MM : Sig = M'_{M'})`: Check $C(M'_{M'})$ is not of the form `cimport ;`. The result is `find_valuabilities(definitions, M'_{M'})` if neither element of this pair is equal to `nonvaluable`.

If any of these checks fail, we have the exception `COMPILE.HASHIFY.BAD_SOURCEDEF_VALUABILITY`.

16.7 Compilation

Formally, compilation is a relation from a name environment E_n , a `sourcefilename`, and a filesystem Φ to either a tuple of a source type environment E_0 , a compiled type environment E_1 , and a `compiledunit`, or an error.

Note that `compiledunit` includes a name environment E_n : this environment contains `cfresh` names created during compilation. This name environment has no implementation significance: its sole purpose is to allow included compiled units to be appropriately typechecked and the configuration produced by compilation to be typechecked. These two checks are both necessary for runtime typechecking, but not otherwise.

Note that compilation is not a function because the choice of name environment in the `compiledunit` is nondeterministic. This nondeterminism is common in many of the helper “functions” throughout, thus we take them all to be relations. For convenience, though, we write them as functions of their inputs, and use \rightsquigarrow rather than $=$ to relate the “input arguments” to the “results”.

Compilation has the form

$$\text{compile}_{\Phi}(\text{sourcefilename})E_n \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$$

defined to be

$$\text{compile}_{\Phi}^{\text{empty}} E_n E_{\text{const}} E_{\text{const}}(\text{includesource sourcefilename ;; empty}) \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$$

where the latter relation

$$\text{compile}_{\Phi}^{\text{definitions}} E_n E_0 E_1(\text{compilationunit}) \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$$

is defined inductively on the *compilationunit*. Here *sourcefilenames* is the filenames we've been through (used to detect cyclic includes), *definitions* is the accumulated compiled definitions, E_n is the accumulated name environment (all names created during compilation will be disjoint from $\text{dom}(E_n)$), E_0 is the accumulated source type environment (including E_{const} at the start), E_1 is the accumulated compiled type environment (including E_{const} at the start), and *compilationunit* is what we have left to do.

It uses auxiliaries $E_0(\text{definitions})$, $E_1(\text{definitions})$, $\text{derive_valuabilities}(\text{definitions}, \text{sourcedefinition})$, $\rho^{\text{definitions}}$, $\text{hashify_ties_and_hashes}(\text{definitions}, e)$, and $\text{hashify}_{\text{definitions}}(E_n, \text{sourcedefinition}, \text{vubs})$ defined in the rest of this section.

Consider the cases of *compilationunit*:

- Case *empty*.
 1. $(E_0, E_1, (E_n, (\text{definitions empty}))$.
- Case *e*.
 1. Check that for some T we have $E_0 \vdash_{\emptyset} e : T$ (otherwise `COMPILE.TYPE`).
 2. Calculate $E_n' = E_n$ and $e' = \text{desugar}(e)$.
 3. An expression is just like a field in a fresh module. Thus at initialisation time, when we have shunted all modules across, we should then apply ρ and rewrite $M'_{M'}@x$, etc, in e' . For now we don't do that, so we have to be very conservative:
 - Check that e' has no `cfresh` subexpressions.
 - Check that for all all $M'_{M'}@x$ and all $\text{hash}(M'_{M'}.x)_T$ subexpressions, $\text{fst}(\text{find_valuabilities}(\text{definitions}, M'_{M'})) \in \{\text{valuable}, \text{cvaluable}\}$.
 - Check that for all $M'_{M'}.t$, we have $\text{snd}(\text{find_valuabilities}(\text{definitions}, M'_{M'})) \in \{\text{valuable}, \text{cvaluable}\}$.
 4. $(E_0, E_1, (E_n', (\text{definitions} (\text{hashify_ties_and_hashes}(\text{definitions}, \rho^{\text{definitions}}(e'))))))$
- Case *(sourcedefinition ;; compilationunit)*.
 1. Check that for some E'_0 we have $E_0 \vdash \text{sourcedefinition} \triangleright E'_0$ (otherwise `COMPILE.TYPE`).
 2. Now we perform case analysis on the structure of *sourcedefinition*. Each case constructs E_n' and *definition'*.
 - Case *sourcedefinition = mark MK*: Let *definition'* = `mark MK` and $E_n' = E_n$.
 - Otherwise:
 - (a) Calculate *vubs* where $\text{vubs} = \text{derive_valuabilities}(\text{definitions}, \text{sourcedefinition})$ (otherwise `COMPILE.BAD_SOURCEDEF_VALUABILITY`).

(b) Now we do case analysis on *mode*:

- * Case $mode \in \{\mathbf{hash}, \mathbf{cfresh}, \mathbf{hash!}, \mathbf{cfresh!}\}$: Calculate $(E_n', definition')$ where $hashify_{definitions}(E_n, desugar(sourcedefinition), vubs) \rightsquigarrow (E_n', definition')$ (otherwise any of the `COMPILE.HASHIFY.*`).
 - * Case $mode = \mathbf{fresh}$: Calculate $E_n' = E_n$ and $definition' = sourcedefinition$.
3. Calculate $E_1' = E_1(definition')$.
 4. Calculate a result of $compile_{\Phi}^{(definitions ;; definition') E_n' (E_0, E_0') (E_1, E_1')}(compilationunit)$ (otherwise any of the compilation errors).
- Case $(\mathbf{includesource} sourcefilename ;; compilationunit)$.
 1. Check $sourcefilename \notin sourcefilenames$ (otherwise `COMPILE.INCLUDE.CYCLE`).
 2. Look up $compilationunit' = \Phi(sourcefilename)$ (otherwise `COMPILE.INCLUDE.SYSTEM_ERROR`).
 3. Calculate $(E_0', E_1', (E_n', definitions' eo'))$ where $compile_{\Phi}^{definitions E_n E_0 E_1}(sourcefilenames, sourcefilename)(compilationunit') \rightsquigarrow (E_0', E_1', (E_n', definitions' eo'))$ (otherwise any of the compilation errors).
 4. Check $eo' = \text{empty} \vee compilationunit = \text{empty}$ (otherwise `COMPILE.NONFINAL_EXPRESSION`).
 5. Calculate a result of $compile_{\Phi}^{definitions' E_n' E_0' E_1'}(compilationunit ;; eo')$ (otherwise any of the compilation errors).
 - Case $(\mathbf{includecompiled} compiledfilename ;; compilationunit)$
 1. Look up $(E_n', (definitions' eo')) = \Phi(compiledfilename)$ (otherwise `COMPILE.INCLUDE.SYSTEM_ERROR`).
 2. Check $eo' = \text{empty} \vee compilationunit = \text{empty}$ (otherwise `COMPILE.NONFINAL_EXPRESSION`).
 3. If using structured hashes,
 - Check $E_n' \vdash definitions' eo'$ **ok**.
 - Check that eo' has no **cfresh** subexpressions.
 - Check that eo' has no ties $M'_{M'}@x$ and no unhashed hashes $\mathbf{hash}(M'_{M'}x)_T$.
 (otherwise `COMPILE.TYPECHECK_OF_COMPILEDUNIT`).
 4. Let $E_n'' = \text{merge_nenvs}(E_n, E_n')$, raising `COMPILE.NENV_MERGE_OF_COMPILEDUNIT` in case of an error. The function `merge_nenvs` merges two name environments, checking that their ranges are identical where their domains intersect.
 5. Let $E_0' = E_0(definitions')$.
 6. Let $E_1' = E_1(definitions')$.
 7. Calculate a result of $compile_{\Phi}^{definitions ;; definitions' E_n'' (E_0, E_0') (E_1, E_1')}(compilationunit ;; eo')$ (otherwise any of the compilation errors).

Note that compilation as defined here operates on abstract syntax. The implementation operates on bytestrings, and so can result also in `FAILURE.COMPILE.LEX` and `FAILURE.COMPILE.PARSE`. We do not specify where these can arise in any more detail.

The definition is given rather algorithmically – it might be nice to rephrase it in a way that makes explicit use of type normalization.

Now we describe the helper functions and relations.

$E_0(\text{definitions})$ and $E_1(\text{definitions})$ extract the bindings for, respectively, the source and compiled signature of a definition

$$\begin{array}{ll}
 E_i(\mathbf{cmodule}_{h;eqs;Sig_0} \text{ vubs } M_M : Sig_1 \text{ version } vn = Str \;; \text{ definitions}) & = M_M : Sig_i, E_i(\text{definitions}) \\
 E_i(\mathbf{cimport}_{h;Sig_0} \text{ vubs } M_M : Sig_1 \text{ version } vc \text{ like } Str \text{ by } \text{resolvespec} = Mo \;; \text{ definitions}) & = M_M : Sig_i, E_i(\text{definitions}) \\
 E_i(\mathbf{module fresh } M_M : Sig \text{ version } vne = Str \text{ withspec} \;; \text{ definitions}) & = M_M : Sig, E_i(\text{definitions}) \\
 E_i(\mathbf{import fresh } M_M : Sig \text{ version } vce \text{ likespec } \text{ by } \text{resolvespec} = Mo \;; \text{ definitions}) & = M_M : Sig, E_i(\text{definitions}) \\
 E_i(\mathbf{mark MK} \;; \text{ definitions}) & = E_i(\text{definitions}) \\
 E_i(\text{empty}) & = \text{empty}
 \end{array}$$

In what follows we let C range over *definitions*.

The substitution $\rho^{\text{definitions}}$, or ρ^C , is defined by

$$\cup \left\{ \begin{array}{l} M'_{M'}.t \mapsto T \mid \left. \begin{array}{l} C(M'_{M'}) = \mathbf{cmodule}_{h';eqs';Sig'_0} \text{ vubs}' M'_{M'} : Sig'_1 \text{ version } vn' = Str' \\ \wedge (\mathbf{type } t_t : EQ(T)) \in Sig'_1 \text{ end} \end{array} \right\} \\ M'_{M'}.t \mapsto T \mid \left. \begin{array}{l} C(M'_{M'}) = \mathbf{cimport}_{h';Sig'_0} \text{ vubs}' M'_{M'} : Sig'_1 \text{ version } vc' \text{ like } Str' \text{ by } \text{resolvespec}' = Mo' \\ \wedge (\mathbf{type } t_t : EQ(T)) \in Sig'_1 \text{ end} \end{array} \right\} \end{array} \right.$$

Comment: ρ^C does have any affect on $M'_{M'}.t$ if $M'_{M'}$ is a **module fresh** or **import fresh** definition in C . These cases will never arrive (thanks to valuability checking) when ρ^C is used in hashification.

The relation `evalcfresh` nondeterministically transforms a pair of a name environment and an expression. This expression is a value modulo the presence of **cfresh**s in `cval` contexts.

$$\text{evalcfresh}(E_n, e) \rightsquigarrow (E'_n, e')$$

It replaces all **cfresh** used in a `cval` context (see §??) by a fresh name and is extended in a standard way to structures:

$$\begin{array}{lll}
 \text{evalcfresh}(E_n, \mathbf{cfresh}_T) & \rightsquigarrow ((E_n, n : T \text{ name}), n) & \text{for } n \notin \text{dom}(E_n) \\
 \text{evalcfresh}(E_n, C_1 e) & \rightsquigarrow (E'_n, C_1 e') & \text{for } \text{evalcfresh}(E_n, e) \rightsquigarrow (E'_n, e') \\
 \text{evalcfresh}(E_n, \mathbf{hash}(T, e)) & \rightsquigarrow (E'_n, \mathbf{hash}(T, e')) & \text{ditto} \\
 \text{evalcfresh}(E_n, \{T, e\} \text{ as } T') & \rightsquigarrow (E'_n, \{T, e'\} \text{ as } T') & \text{ditto} \\
 \text{evalcfresh}(E_n, e_1 :: e_2) & \rightsquigarrow (E'_{n_2}, e'_1 :: e'_2) & \text{for } \text{evalcfresh}(E_n, e_1) \rightsquigarrow (E'_{n_1}, e'_1) \\
 & & \text{evalcfresh}(E'_{n_1}, e_2) \rightsquigarrow (E'_{n_2}, e'_2) \\
 \text{evalcfresh}(E_n, \mathbf{hash}(T, e_1, e_2)_{T'}) & \rightsquigarrow (E'_{n_2}, \mathbf{hash}(T, e'_1, e'_2)_{T'}) & \text{ditto} \\
 \text{evalcfresh}(E_n, (e_1, \dots, e_k)) & \rightsquigarrow (E'_{n_k}, (e'_1, \dots, e'_k)) & \text{for } E_{n_0} = E_n \\
 & & \text{evalcfresh}(E_{n_0}, e_1) \rightsquigarrow (E'_{n_1}, e'_1) \\
 & & \dots \\
 & & \text{evalcfresh}(E'_{n_{k-1}}, e_k) \rightsquigarrow (E'_{n_k}, e'_k) \\
 \text{evalcfresh}(E_n, \mathbf{struct } str \text{ end}) & \rightsquigarrow (E'_n, \mathbf{struct } str' \text{ end}) & \text{for } \text{evalcfresh}(E_n, str) \rightsquigarrow (E'_n, str') \\
 \text{evalcfresh}(E_n, (\mathbf{let } x_x = e) str) & \rightsquigarrow (E''_n, (\mathbf{let } x_x = e') str') & \text{for } \text{evalcfresh}(E_n, e) \rightsquigarrow (E'_n, e') \\
 & & \text{evalcfresh}(E'_n, str) \rightsquigarrow (E''_n, str') \\
 \text{evalcfresh}(E_n, (\mathbf{type } t_t = T) str) & \rightsquigarrow (E'_n, (\mathbf{type } t_t = T) str') & \text{for } \text{evalcfresh}(E_n, str) \rightsquigarrow (E'_n, str')
 \end{array}$$

Now we define a helper function `hashify_ties_and_hashes(C, e)` which compiles all ties and hashes as follows: in e each $M'_{M'}@x$ is replaced by `TIECON(hash($\sigma^C(M'_{M'}.x)$) $_{T'}$, $M'_{M'}.x$)` and each `hash($M'_{M'}.x$) $_T$` is replaced by `hash($\sigma^C(M'_{M'}.x)$) $_T$` , where σ is defined below and T is the type of $M'_{M'}.x$ in the relevant signature in C . We extend the domain to structures, `hashify_ties_and_hashes(C, str)`, by applying the expression-level version pointwise to the fields of str .

The relation $\text{hashify}_C(E_n, \text{sourcedefinition}, \text{vubs})$ is defined by case analysis as follows. Here C is a list of *definitions*, which we also regard as a partial function mapping each M_M to a *definition*. It returns a new E_n' (containing E_n and any new names) and a *definition*, or fails with one of FAILURE.COMPILE.HASHIFY.*.

Note that it is used

- during compilation for **hash** and **cfresh** modules and imports;
- during run time for initialisation of **fresh** modules and imports.

We consider the following cases:

- Case $\text{sourcedefinition} = \text{module } \text{mode } M_M : \text{Sig } \text{version } \text{vne} = \text{Str withspec}$, where $\text{Sig} = \text{sig } \text{sig } \text{end}$ and $\text{Str} = \text{struct } \text{str } \text{end}$.

(Convention: things subscripted by n are roughly the result of the n -th step.)

1. Remove other-module type dependencies with ρ^C . Let $(\text{str}_1 : \text{sig}_1) = \rho^C(\text{str} : \text{sig})$.

This replaces references to $M'_{M'}.t$ by the manifest type from the compiled module or import signature – and in compiled modules and imports, all types in the Sig_1 have been made manifest.

2. Let $\text{str}'_1 = \text{hashify_ties_and_hashes}(C, \text{str}_1)$.
3. Then, remove same-module type dependencies, i.e. references t to previous type fields, as far as possible, with $\text{typeflattenstruct}()$ and $\text{typeflattensig}()$.

$$\begin{aligned} \text{typeflattenstruct}(\text{type } t_t = T \text{ str}) &= (\text{type } t_t = T) \text{ typeflattenstruct}(\{T/t\}\text{str}) \\ \text{typeflattenstruct}(\text{let } x_x = v^{eqs} \text{ str}) &= (\text{let } x_x = v^{eqs}) \text{ typeflattenstruct}(\text{str}) \\ \text{typeflattenstruct}(\text{empty}) &= \text{empty} \end{aligned}$$

$$\begin{aligned} \text{typeflattensig}(\text{type } t_t : \text{TYPE } \text{sig}) &= (\text{type } t_t : \text{TYPE}) \text{ typeflattensig}(\text{sig}) \\ \text{typeflattensig}(\text{type } t_t : \text{EQ}(T) \text{ sig}) &= (\text{type } t_t : \text{EQ}(T)) \text{ typeflattensig}(\{T/t\}\text{sig}) \\ \text{typeflattensig}(\text{val } x_x : T \text{ sig}) &= (\text{val } x_x : T) (\text{typeflattensig}(\text{sig})) \\ \text{typeflattensig}(\text{empty}) &= \text{empty} \end{aligned}$$

Let $\text{str}_2 : \text{sig}_2 = \text{typeflattenstruct}(\text{str}'_1) : \text{typeflattensig}(\text{sig}_1)$.

This $\text{typeflattensig}()$ leaves internal references to abstract type fields – that is forced, as we cannot yet calculate the h required to build their replacements.

This type normalisation (both the ρ and type flattening here, and selfification below) amounts to treating modules up to type equality when making **hmodule** s and **nmodule** s, instead of using exactly the abstract syntax. Working up to type equality seems intuitively preferable, though it makes compilation seem even more algorithmic.

4. Generate **cfresh** names if needed. We proceed by case analysis on mode and establish str_3 and E_{n3} :

- Case $\text{mode} \in \{\text{hash}, \text{fresh}, \text{hash!}\}$:

Let $\text{str}_3 = \text{str}_2$ and $E_{n3} = E_n$.

Comment: Note that $\text{hashify}(\ , \ , \)$ is called at run time, not compile time, for **module fresh**. For such modules, the **cfresh** subexpressions are thus eliminated by compilation and not here.

- Case $\text{mode} = \text{cfresh}, \text{cfresh!}$:

Calculate E_{n3}, str_3 by applying evalcfresh to eliminate all **cfresh** expressions: $\text{evalcfresh}(E_n, \text{str}_2) \rightsquigarrow (E_{n3}, \text{str}_3)$.

5. Check the *withspec* (see-through semantics).

Suppose $withspec = \mathbf{with} !weqs$.

Let $eqs = \{M'_{M'.t} \approx \rho^C T \mid M'_{M'.t} \approx T \in weqs\}$.

For all $M'_{M'.t} \approx T \in weqs$, if $C(M'_{M'})$ is an import then fail (`WITHSPEC_EQUATION_FROM_IMPORT`). Otherwise, suppose

$$C(M'_{M'}) = \mathbf{cmodule}_{h';eqs';Sig'_0} vubs' M'_{M'} : Sig'_1 \mathbf{version} vn' = Str' \wedge (\mathbf{type} \ t_t : \mathbf{TYPE}) \in Sig'_0 \wedge (\mathbf{type} \ t_t = TrefromC) \in Str'$$

Comment: $M'_{M'.t}$ exists and is abstract, since the source definition is well-typed. It's unclear whether abstractness is forced but there does not seem to be any reason to permit seeing through a non-abstract type.

Check $TrefromC = \rho^C T$ (or fail with `WITHSPEC_TYPES_NOT_EQUAL`).

Comment: $TrefromC$ was already closed by the ρ from an earlier stage, so applying (the current) ρ to the T in $weqs$ means that syntactic type equality is the appropriate check.

Comment: There are two possible semantics for see-through $\mathbf{with} !$. Currently we permit eqs to be used anywhere in typing the structure part; alternatively one could allow it to be used only in a final subsignature step. Unclear which is preferable in practice.

We need to construct a closed set of equations for use inside the `nmodule` or `hmodule`.

Let $eqs' = \{\rho^C(M'_{M'.t}) \approx \rho^C T \mid (M'_{M'.t} \approx T) \in weqs\}$.

6. Remove other-module term dependencies with $\sigma^C =$

$$\cup \left\{ \begin{array}{l} M'_{M'.x} \mapsto h'.x \\ M'_{M'.x} \mapsto h'.x \end{array} \left| \begin{array}{l} C(M'_{M'}) = \mathbf{cmodule}_{h';eqs';Sig'_0} vubs' M'_{M'} : Sig'_1 \mathbf{version} vn' = Str' \\ \wedge (\mathbf{val} \ x_x : T) \in Sig'_0 \\ C(M'_{M'}) = \mathbf{cimport}_{h';Sig'_0} vubs' M'_{M'} : Sig'_1 \mathbf{version} vc' \mathbf{like} Str' \mathbf{by} resolvespec' = Mo' \\ \wedge (\mathbf{val} \ x_x : T) \in Sig'_0 \end{array} \right. \right\}$$

Let $str_5 = \sigma^C str_3$.

7. We do case analysis on the mode, calculating h and E_n' .

– Case $mode = \mathbf{hash}$:

Let $E_n' = E_{n_3}$ and $h = \mathbf{hash}(\mathbf{hmodule} \ eqs' \ M : \mathbf{sig} \ sig_2 \ \mathbf{end} \ \mathbf{version} \ vne = \mathbf{struct} \ str_5 \ \mathbf{end})$.

Comment: A possible alternative semantics would be to substitute the eqs out in the body of the hash, making hash equality slightly coarser. It is unclear whether that would be preferable or not.

– Case $mode = \mathbf{cfresh}, \mathbf{fresh}$:

Let n be a fresh name not in the domain of E_{n_3} (to serve as the new name of this module). We extend E_{n_3} to E_n' accordingly:

$$E_n' = E_{n_3, n} : \mathbf{nmodule}_{eqs'} M : \mathbf{sig} \ sig_2 \ \mathbf{end} \ \mathbf{version} \ vne = \mathbf{struct} \ str_5 \ \mathbf{end}$$

Let $h = n$.

8. Selfify with respect to that h , to remove same-module type references. (`selffysig ()` is defined on page ??.)

Let $sig_7 = \text{typeflattensig}(\text{selffysig}_h(sig_2))$.

Comment: Note that **hmodule** and **nmodule** generation happens before abstract-type selfification (and evaluation of vne). That is forced, as we need the h to selfify abstract type components (and to do that evaluation). However, manifest type components do get substituted out before.

Comment: Stylistic choice: you have to flatten the sig again, either in the definition of $\text{selffysig}_h()$ (as we used to), or with $\text{typeflattensig}(\dots)$ again. We do the latter, so that $\text{selffysig}_h()$ can be used in the module alias typing rule.

Comment: In previous versions, and in [Sew01], selfify not only replaced **TYPE** by $\text{EQ}(h)$ in the signature, but also replaced t by $h.t$ in the structure (letting h range over hashes and new names). In [Sew01] that was because functors took type fields and term fields from their argument *struct*, not their argument sig, and so to not replace t by $h.t$ would have been broken. An unfortunate consequence of doing that $\{h.t/t\}$ in the struct, however, is that you need to keep the t elsewhere (in [Sew01], formally in the global type environment) for representation type checking of **with !**. Now we realise that that was not really forced. As signatures have always been fully EQified before you apply a functor (one case) or construct a ρ^C to use later (the other case), we can have both functors and ρ pull out type fields from sigs instead of structs.

Comment: With respect to marshalling (or fresh name generation etc.) inside an abstraction boundary (cf. §8.5), however, doing $\{h.t/t\}$ in the structure might well be preferable. That would require changes to the construction of *eqs*, for which one might want to do the substitution in expression fields but not in the definitions of abstract types.

9. Evaluate the version number expression.

Let $vn = \{h/\text{myname}\}(vne)$.

10. Finally, put that all together, writing in the h and the equations.

Let $\text{definition}' = \text{cmodule}_{h;eqs';\text{sig } sig_2} \text{end vubs } M_M : \text{sig } sig_7 \text{ end version } vn = \text{struct } str_2 \text{ end}$ and let E_n' be as built above.

- Case *sourcedefinition* = **import** *mode* $M_M : Sig$ **version** *vc* *likespec* **by** *resolvespec* = Mo and $Sig = sig$ **sig** *sig* **end**.

Comment: Note we have an *h* subscript on **cimport**, *s* too, for convenience during compilation.

1. Normal case: the *vc* is not an exact-name constraint, ie $vc = dvce$ for some *dvce*.

(a) Calculate a *likestr'* without internal type dependencies or other-module type dependencies. There are three cases: either *likespec* was empty, or an in-line structure, or a module identifier (in the last case, we allow import-bound identifiers, as there seems no reason to exclude them).

– Case *likespec* = empty. If this import is linked to $M'_{M'}$, then the empty *likespec* defaults to **like** $M'_{M'}$ (see below). Otherwise, take *likestr* = empty.

– Case *likespec* = **like struct** *str* **end**.

Use the auxiliary `typeflattenstruct()` to substitute out occurrences of internal type field names *t* within *str*.

Let *likestr* = `typeflattenstruct(ρ^C str)`.

– Case *likespec* = **like** $M'_{M'}$. Either

$$\begin{aligned} C(M'_{M'}) &= \mathbf{cmodule}_{h';eqs';Sig'_0} vubs' M'_{M'} : Sig'_1 \mathbf{version} vn' = Str' \\ Str' &= \mathbf{struct} likestr \mathbf{end} \end{aligned}$$

or

$$\begin{aligned} C(M'_{M'}) &= \mathbf{cimport}_{h';Sig'_0} vubs' M'_{M'} : Sig'_1 \mathbf{version} vc' \mathbf{like} Str' \mathbf{by} resolvespec' = Mo' \\ Str' &= \mathbf{struct} likestr \mathbf{end} \end{aligned}$$

Comment: These are the only two cases we need consider: if we are being called at compile time then *mode* $\in \{\text{valuable}, \text{cvaluable}\}$; valuability then ensures that $M'_{M'}$ cannot be a fresh module or import. If we are being called at run time, then all the previous definitions have been hashified already.

Comment: In the second case it might be more intuitive to insist that the *likestr* has exactly the needed fields, rather than (as here) permit it to have more.

Now calculate a *likestr'* by cutting down the *likestr* to the abstract type part of *Sig*. To do that we define the auxiliary function `filter str sig` which calculates the subsequence of *str* with the external type fields of *sig*. It is assumed that *sig* contains no value fields. It is a partial function, failing if there are not enough type fields in *str*, and only constructs a sensible struct if the struct argument is type-flattened.

$$\begin{aligned} \text{filter}(\mathbf{type} \ t_t = T \ str)(\mathbf{type} \ t_{t'} : K \ sig) &= (\mathbf{type} \ t_t = T) (\text{filter} \ str \ sig) \\ \text{filter}(\mathbf{type} \ t_t = T \ str)(\mathbf{type} \ t'_{t'} : K \ sig) &= (\text{filter} \ str(\mathbf{type} \ t'_{t'} : K \ sig)) \text{ if } t \neq t' \\ \text{filter}(\mathbf{let} \ x_x = v^{eqs} \ str) \ sig &= \text{filter} \ str \ sig \\ \text{filter}(\mathbf{type} \ t_t = T \ str) \ \text{empty} &= \text{empty} \\ \text{filter} \ \text{empty} \ \text{empty} &= \text{empty} \\ \text{filter} \ \text{empty} \ sig &\text{ undefined if } sig \text{ is non empty} \end{aligned}$$

Let *likestr'* = `filter likestr(limitdom(sig))` (or fail with `LIKESPEC_MISSING_TYPE_FIELDS` if this is undefined).

Comment: This semantics permits the *likestr* to contain more fields than are required (or will appear in the constructed *likestr'* of this import when compiled). Inelegant?

Comment: Because we cut *likestr* down to a *likestr'*, a structure containing only type fields, we have no need to worry about *likestr'* containing uses of `cfresh` or `fresh`.

- (b) Let $sig_0 = \text{typeflattensig}(\rho^C sig)$. Let $Sig_0 = \mathbf{sig} \ sig_0 \ \mathbf{end}$.
- (c) Let vc be the result of evaluating vce with respect to C , replacing any $M'_{M'}$ by the hash associated with the module or import in C , i.e. $vc = \rho^C vce$.
- (d) ** Now we do case analysis on $mode$ to construct h and E_n' :
- Case $mode \in \{\mathbf{hash}, \mathbf{hash!}\}$: Let $h = \mathbf{hash}(\mathbf{himport} \ M \ : \ Sig_0 \ \mathbf{version} \ vc \ \mathbf{like} \ \mathbf{struct} \ likestr' \ \mathbf{end})$ and let $E_n' = E_n$.
 - Case $mode \in \{\mathbf{cfresh}, \mathbf{fresh}, \mathbf{cfresh!}\}$: Let n be a fresh name not in the domain of E_n . Let $h = n$ and let

$$E_n' = E_n, n : \mathbf{nimport} \ M : Sig_0 \ \mathbf{version} \ vc \ \mathbf{like} \ \mathbf{struct} \ likestr' \ \mathbf{end}$$

Comment: We choose not to include *resolvespecs* in *hmodules* or *nmodules* of imports. This is debatable – the argument against including them is that it is useful to be able to change location without breaking code (local code mirror, changing web site to avoid MSBlast.exe, etc.).

- (e) Selfify the sig. Let $Sig_1 = \mathbf{sig} \ \text{typeflattensig}(\text{selfifysig}_h(sig_0)) \ \mathbf{end}$.
- (f) Calculate $definition' = \mathbf{cimport}_{h;Sig_0} \ vubs \ M_M : Sig_1 \ \mathbf{version} \ vc \ \mathbf{like} \ \mathbf{struct} \ likestr' \ \mathbf{end} \ \mathbf{by} \ \text{resolvespec} = Mo$.
- (g) If $Mo = M''_{M''}$ then check $\text{linkok}(E_n', definition'', definition')$ where $C(M''_{M''}) = definition''$ (or fail with LINKOK_NOT). Otherwise if $Mo = \text{UNLINKED}$ check true.
- (h) The result is $(E_n', definition')$.

2. exact-name case: if the vce is an exact-name constraint $\mathbf{name} = M'_{M'}$, then we must have $likespec =$ empty (this is enforced by a syntactic requirement).

The name of this import will be exactly the name of $M'_{M'}$.

Construct

$$\text{sourcedefinition}_1 = \mathbf{import} \ mode \ M_M \ : \ Sig \ \mathbf{version} \ \mathbf{name} = M'_{M'} \ \mathbf{like} \ M'_{M'} \ \mathbf{by} \ \text{resolvespec} = Mo$$

and use the normal-case algorithm as above except that we take the h to be the one associated with the module or import $M'_{M'}$ in C in the step marked **.

Comment: You could hash this import instead of using h in $definition'$. This gives a slightly coarser type equality, which might sometimes be handy, but when you come make up exact-name imports the choice is forced: for type preservation those have to have exactly the hash of the module they are made up from.

- Case $\text{sourcedefinition} = \mathbf{module} \ M_M : Sig = M'_{M'}$.

Note that by typing $M'_{M'}$ cannot be a **module fresh** or **import fresh**.

- Case $C(M'_{M'}) = \mathbf{cmodule}_{h';eqs';Sig'_0} \ vubs' \ M'_{M'} : Sig'_1 \ \mathbf{version} \ vn' = Str'$
Take $definition' = \mathbf{cmodule}_{h';eqs';Sig'_0} \ vubs' \ M_M : Sig'_1 \ \mathbf{version} \ vn' = Str'$ (identical except for the M_M).
- Case $C(M'_{M'}) = \mathbf{cimport}_{h';Sig'_0} \ vubs' \ M''_{M''} : Sig'_1 \ \mathbf{version} \ vc' \ \mathbf{like} \ Str' \ \mathbf{by} \ \text{resolvespec}' = Mo'$
Take $definition' = \mathbf{cimport}_{h';Sig'_0} \ vubs' \ M_M : Sig'_1 \ \mathbf{version} \ vc' \ \mathbf{like} \ Str' \ \mathbf{by} \ \text{resolvespec}' = Mo'$ (identical except for the M_M).

Comment: There are two options here: either copy the *definition* from $M'_{M'}$ – but that is semantically odd when one does any rebinding – or just keep the alias in the resulting code – but then both C and runtime lookup need to go through aliases transparently.

However, as aliases are present just to get module names into scope for **with !** and version annotations, to avoid formalising a filesystem containing modules, for now it is not worth doing anything more elaborate than the above, even though it is strange to copy modules and imports across marks.

- Case *sourcedefinition* = **mark** MK.

Take $definition' = sourcedefinition$.

16.8 Operational semantics

16.8.1 The judgements

We define a labelled transition system over configurations with judgements as follows.

• $E_n ; \langle E_s, s, definitions, P \rangle \xrightarrow{n:\ell} E_n' ; \langle E'_s, s', definitions', P' \rangle$	Process reduction.
• $E_n ; \langle E_s, s, definitions, P \rangle \rightarrow \mathbf{TERM}$	Program termination.
• $E_n ; \langle E_s, s, definitions, e \rangle \xrightarrow{\ell}_{eqs} E_n' ; \langle E'_s, s', definitions', e' \rangle$	Expression reduction.
• $e \xrightarrow{\ell}_{eqs} e'$	
• $E_n ; e \xrightarrow{\ell}_{eqs} E_n' ; e'$	
• $P \xrightarrow{\ell} P'$	
where	
$\ell ::=$	
empty	internal reduction step
$x^n v_1^\emptyset \dots v_n^\emptyset$ for $x^n \in \text{dom}(E_{\text{const}}) \wedge \text{os}(x^n)$	invocation of OS call
$\text{Ok}(v^\emptyset)$	return from OS call
$\text{Ex}(v^\emptyset)$	return from OS call
$\text{GetURI}(URI)$	request for code at URI
$\text{DeliverURI}(definitions)$	resulting code
CannotFindURI	nothing found at URI
We write $\xrightarrow{\text{empty}}$ simply as \rightarrow .	

In addition the runtime implementation might fail with the RUN or INTERNAL errors, though it should not.

16.8.2 Values

The set of values is indexed by a colour, to control the administrative bracket-pushing reductions, as follows. Note that colour is a spectral phenomenon — two entities have the same colour iff they are indexed by the same set of equations

$$\begin{aligned}
 v^{eqs} ::= & C_0 \\
 & C_1 v^{eqs} \\
 & v^{eqs} :: v^{eqs} \\
 & (v_1^{eqs}, \dots, v_n^{eqs}) \\
 & \mathbf{function} (x : T) \rightarrow e \\
 & l \\
 & \mathbf{nn} \\
 & \Lambda t \rightarrow e \\
 & \{T, v^{eqs}\} \mathbf{as} T' \\
 & [v^{eqs'}]_{eqs'}^T \mathbf{ref} \\
 & [v^{eqs'}]_{eqs'}^T \mathbf{name} \\
 & [v^{eqs'}]_{eqs'}^{h.t} \text{ where } h.t \in \text{dom}(eqs') \text{ and } h.t \notin \text{dom}(eqs)
 \end{aligned}$$

Comment: The different families of values collapse into a single one when there are no coloured brackets, such is the case with user source programs.

This just says that within values, brackets may only appear at types with no visible structure or at a T ref or T name type. This is achieved by the bracket-pushing reductions below. For discussion of these and of the possible design choices for the T ref and T name cases, see §16.8.4 (page 130).

16.8.3 Reduction contexts and closure rules

We use redex-time instantiation for module identifiers, but as we have marks only between the (second-class) modules, there is no need to be anything other than conventional call-by-value λ_c *within* the running expression. Evaluation contexts are therefore conventional, except that we must track colours.

Single-level evaluation contexts and colour-changing evaluation contexts

$C_{eqs} ::=$	C_1 - $_ :: e$ $v^{eqs} :: _$ $(e_1, \dots, e_{m-1}, _ , v_{m+1}^{eqs}, \dots, v_n^{eqs}) \quad n \geq 2, 1 \leq m \leq n$ if $_$ then e_1 else e_2 $_ \&\& e$ $_ e$ $_ ; e$ $_ e$ $v^{eqs} _$ $e \ e_1 \dots e_{m-1} _ v_{m+1}^{eqs} \dots v_n^{eqs} \quad 1 \leq m \leq n, e = op^n$ or $e = x^n$ $!_T _$ $_ :=_T e$ $v^{eqs} :=_T _$ match $_$ with $mtch$ raise $_$ try $_$ with $mtch$ marshal $MK _ : T$ marshal $_ e_2 : T$ unmarshal $_$ as T swap $_$ and e_2 in e_3 swap v_1^{eqs} and $_$ in e_3 swap v_1^{eqs} and v_2^{eqs} in $_$ $_$ freshfor e_2 v_1^{eqs} freshfor $_$ support $_T$ - name_of_tie $_$ val_of_tie $_$ $_ T$ let $\{t, x\} = _$ in e_2 namecase $_$ with $\{t, (x_1, x_2)\}$ when $x_1 = e \rightarrow e_2$ otherwise $\rightarrow e_3$ namecase v^{eqs} with $\{t, (x_1, x_2)\}$ when $x_1 = _ \rightarrow e_2$ otherwise $\rightarrow e_3$
$C_{eqs_1}^{eqs_2} ::=$	C_{eqs_1} $[-]_{eqs_2}^T$ marshalz $MK _ : T$ $l :='_T _$ $eqs_1 = eqs_2$ $eqs_2 = \emptyset$ $eqs_2 = \emptyset$

$\mathbf{op}(e_0)^n$	$e_1 \dots e_{i-1} \cdot v_{i+1}^\emptyset \dots v_n^\emptyset$	$1 \leq i \leq n, eqs_2 = \emptyset$
$\mathbf{hash}(T, _)$	T'	$eqs_2 = \emptyset$
$\mathbf{hash}(T, v_1^\emptyset, T)$	T'	$eqs_2 = \emptyset$
$\mathbf{hash}(T, _, e_2)$	T'	$eqs_2 = \emptyset$

It is sometimes convenient to refer to bracket contexts (sequences of nested brackets).

Bracket contexts

BC	$::=$	$_$
		$BC.[_]_{eqs}^T$

We follow the evaluation order of `ocamlpt` (not `ocamlc`), except that `ocamlpt` treats saturated applications of operators (such as `(+)`) specially, whereas we treat all functions and operators uniformly. Specifically, we evaluate applications from left to right in all cases, whereas `ocamlpt` evaluates a saturated operator (either `e1 + e2` or `(+) e1 e2`, but not `((+) e1) e2`) from right-to-left, and `ocamlc` evaluates all applications from right-to-left. Note that in both, tuples are evaluated right-to-left.

Evaluation contexts and Colour changing evaluation contexts

CC_{eqs}	$::=$	$_$	$CC_{eqs_1}^{eqs_1}$	$::=$	$_$
		$CC_{eqs} \cdot C_{eqs}$			$CC_{eqs_1}^{eqs_1} \cdot C_{eqs_2}^{eqs_2}$

Structure evaluation contexts and thread evaluation contexts

TC_{eqs}	$::=$	$_$
		$(\mathbf{cmodule}_{h;eqs;Sig_0} \text{ vubs } M_M : Sig_1 \mathbf{version} \text{ vn} = \mathbf{struct} \text{ } SC_{eqs} \mathbf{end}) \text{ definitions } e$
TCC_{eqs}	$::=$	$_$
		$TC_{eqs_2} \cdot CC_{eqs}^{eqs_2}$
SC_{eqs}	$::=$	$\mathbf{let} \text{ } x_x = _ \text{ str}$
		$\mathbf{let} \text{ } x_x = v^{eqs} SC_{eqs} \text{ where } x \notin \text{fv } SC_{eqs}$
		$\mathbf{type} \text{ } t_t = T \text{ } SC_{eqs}$
$strval_{eqs}$	$::=$	$\mathbf{let} \text{ } x_x = v^{eqs} \text{ strval}_{eqs}$
		$\mathbf{type} \text{ } t_t = T \text{ strval}_{eqs}$

Module and definition values Say a *module value* is a *definition* of the form `$\mathbf{cmodule}_{h;eqs;Sig_0} \text{ vubs } M_M : Sig_1 \mathbf{version} \text{ vn} = \mathbf{struct} \text{ } strval_{eqs} \mathbf{end}$` where there are no internal expression field dependencies in `$strval_{eqs}$` .

Say a *definition value* is a `$\mathbf{cmodule}$` value, a `$\mathbf{cimport}$` , or a `\mathbf{mark}` `MK` .

These induce the following reductions.

$$\begin{array}{c}
\text{if } \mathit{definition} \text{ is a definition value} \\
\hline
\frac{E_n ; \langle E_s, s, \mathit{definitions}_0, P | \mathbf{n} : (\mathit{definition} \ \mathit{definitions} \ e) \rangle \xrightarrow{\mathbf{n}:\text{empty}} E_n ; \langle E_s, s, \mathit{definitions}_0 \ \mathit{definition}, P | \mathbf{n} : (\mathit{definitions} \ e) \rangle}{\text{if } \mathit{definition} \text{ is of the form } \mathbf{module} \ \mathbf{fresh} \ \mathbf{or} \ \mathbf{import} \ \mathbf{fresh} \\ \text{and hashify}_{\mathit{definitions}_0}(E_n, \mathit{definition}, (\text{nonvaluable}, \text{nonvaluable})) \rightsquigarrow (E_n', \mathit{definition}')} \\
\hline
\frac{E_n ; \langle E_s, s, \mathit{definitions}_0, P | \mathbf{n} : (\mathit{definition} \ \mathit{definitions} \ e) \rangle \xrightarrow{\mathbf{n}:\text{empty}} E_n' ; \langle E_s, s, \mathit{definitions}_0, P | \mathbf{n} : (\mathit{definition}' \ \mathit{definitions} \ e) \rangle}{e \xrightarrow{\ell}_{eqs} e'} \\
\hline
\frac{E_n ; \langle E_s, s, \mathit{definitions}, e \rangle \xrightarrow{\ell}_{eqs} E_n ; \langle E_s, s, \mathit{definitions}, e' \rangle}{E_n ; e \xrightarrow{\ell}_{eqs} E_n', e'} \\
\hline
\frac{E_n ; \langle E_s, s, \mathit{definitions}, e \rangle \xrightarrow{\ell}_{eqs} E_n' ; \langle E'_s, s', \mathit{definitions}', e' \rangle}{E_n ; \langle E_s, s, \mathit{definitions}, C_{eqs_2}^{eqs_1}.e \rangle \xrightarrow{\ell}_{eqs_1} E_n' ; \langle E'_s, s', \mathit{definitions}', C_{eqs_2}^{eqs_1}.e' \rangle} \\
\hline
\frac{E_n ; \langle E_s, s, \mathit{definitions}, e \rangle \xrightarrow{\ell}_{eqs} E_n' ; \langle E'_s, s', \mathit{definitions}', e' \rangle}{E_n ; \langle E_s, s, \mathit{definitions}, P | \mathbf{n} : TC_{eqs}.e \rangle \xrightarrow{\mathbf{n}:\ell} E_n' ; \langle E'_s, s', \mathit{definitions}', P | \mathbf{n} : TC_{eqs}.e' \rangle} \\
\hline
\frac{P \xrightarrow{\mathbf{n}:\ell} P'}{E_n ; \langle E_s, s, \mathit{definitions}, P \rangle \xrightarrow{\mathbf{n}:\ell} E_n ; \langle E_s, s, \mathit{definitions}, P' \rangle}
\end{array}$$

If the hashification of fresh definition *definition* in rule 2 above fails, the error is reported at toplevel as FAILURE.COMPILE.HASHIFY and the program terminates. This is unpleasant, but unavoidable in the absence of exception handlers around threads.

16.8.4 Simple expression forms

Eliminating internal field dependencies When performing module initialisation we evaluate each field in order, and for each replace all later uses of it by its value. This ensures we do not need to consider marshalling or placing in the store a thunk containing a free x . For simplicity, we do this systematically to all **modules** even when it is not forced (which could be detected by looking at their valuabilities).

This strategy has some impact on rebinding: if one has a module $M_{M\dots} = \mathbf{struct} \ \mathbf{let} \ x_x = 3 \ \mathbf{let} \ y_y = \mathbf{function} \ () \rightarrow x \ \mathbf{end}$ before initialisation then the x will be eliminated in the y_y field. Thus later externally instantiating $M_{M.y}$ gives $\mathbf{function} \ () \rightarrow 3$ rather than $\mathbf{function} \ () \rightarrow M_{M.x}$.

The strategy also has implication for **swap** — as there's no need to follow x uses in a value.

$$\frac{}{\text{n:empty} \rightarrow E_n ; \langle E_s, s, \text{definitions}, P | \mathbf{n} : (\text{definition definitions}' e) \rangle}$$

where

$$\begin{aligned} \text{definition} &= (\mathbf{cmodule}_{h;eqs;Sig_0} \text{ vubs } M_M : Sig_1 \mathbf{version} \text{ vn} = \mathbf{struct} \text{ str} \mathbf{end}) \\ \text{definition}' &= (\mathbf{cmodule}_{h;eqs;Sig_0} \text{ vubs } M_M : Sig_1 \mathbf{version} \text{ vn} = \mathbf{struct} \text{ str}' \mathbf{end}) \\ \text{str} &= \text{strval}_{eqs} \mathbf{let} \ x_x = v^{eqs} \text{ str}_0 \\ \text{str}' &= \text{strval}_{eqs} \mathbf{let} \ x_x = v^{eqs} \{v^{eqs}/x\} \text{ str}_0 \\ x &\in \text{fv}(\text{str}_0) \\ \text{dom}(\text{strval}_{eqs}) &\text{ does not intersect the free expression identifiers of } \text{str} \end{aligned}$$

Note the $x \in \text{fv}(\text{str}_0)$ condition to prevent divergence.

Comment: This rule is terminating because of the $x \in \text{fv}(\text{str})$ condition. This condition, in combination with the free identifier side condition on SC_{eqs} contexts used in module field initialisation forces the two reduction rules to be disjoint. It might be more tasteful to work with a **cmodule** that is split into the post- and pre-evaluation parts, but it would be notationally heavy.

Matching Define a partial function $\text{matchsub}_{eqs}(-, -)$ taking a value of colour eqs and a pattern (in which all identifiers are distinct) and giving a set of substitutions, adding suitable brackets:

$$\begin{aligned} \text{matchsub}_{eqs}(v^{eqs}, (- : T)) &= \emptyset \\ \text{matchsub}_{eqs}(v^{eqs}, (x : T)) &= \{[v^{eqs}]_{eqs}^T/x\} \\ \text{matchsub}_{eqs}(v^{eqs}, (p : T)) &= \text{matchsub}_{eqs}(v^{eqs}, p) \\ \text{matchsub}_{eqs}(C_0, C_0) &= \emptyset \\ \text{matchsub}_{eqs}(C_1 v^{eqs}, C_1 p) &= \text{matchsub}_{eqs}(v^{eqs}, p) \\ \text{matchsub}_{eqs}(v_1^{eqs} :: v_2^{eqs}, p_1 :: p_2) &= \text{matchsub}_{eqs}(v_1^{eqs}, p_1) \cup \text{matchsub}_{eqs}(v_2^{eqs}, p_2) \\ \text{matchsub}_{eqs}((v_1^{eqs}, \dots, v_n^{eqs}), (p_1, \dots, p_n)) &= \text{matchsub}_{eqs}(v_1^{eqs}, p_1) \cup \dots \cup \text{matchsub}_{eqs}(v_n^{eqs}, p_n) \quad n \geq 2 \\ \text{matchsub}_{eqs}(v^{eqs}, p) &= \text{undefined otherwise} \end{aligned}$$

Reduction Axioms

$$\begin{aligned} \mathbf{if} \ \mathbf{true} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 &\rightarrow_{eqs} e_1 \\ \mathbf{if} \ \mathbf{false} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 &\rightarrow_{eqs} e_2 \\ \mathbf{false} \ \&\& \ e &\rightarrow_{eqs} \mathbf{false} \\ \mathbf{true} \ \&\& \ e &\rightarrow_{eqs} e \\ \mathbf{false} \ || \ e &\rightarrow_{eqs} e \\ \mathbf{true} \ || \ e &\rightarrow_{eqs} \mathbf{true} \\ () ; e &\rightarrow_{eqs} e \\ \mathbf{while} \ e_1 \ \mathbf{do} \ e_2 \ \mathbf{done} &\rightarrow_{eqs} \mathbf{if} \ e_1 \ \mathbf{then} \ (e_2 ; \mathbf{while} \ e_1 \ \mathbf{do} \ e_2 \ \mathbf{done}) \ \mathbf{else} \ () \\ (\mathbf{function} \ (x : T) \rightarrow e) \ v^{eqs} &\rightarrow_{eqs} \{[v^{eqs}]_{eqs}^T/x\}e \\ \mathbf{match} \ v^{eqs} \ \mathbf{with} \ p_1 \rightarrow e_1 | \dots | p_n \rightarrow e_n &\rightarrow_{eqs} \text{matchsub}_{eqs}(v^{eqs}, p_k) e_k \quad (a) \\ \mathbf{match} \ v^{eqs} \ \mathbf{with} \ p_1 \rightarrow e_1 | \dots | p_n \rightarrow e_n &\rightarrow_{eqs} \mathbf{raise} \ \text{MATCH_FAILURE} \ v' \quad (b) \\ \mathbf{let} \ \mathbf{rec} \ x_1 : T = \mathbf{function} \ (x_2 : T') \rightarrow e_1 \ \mathbf{in} \ e_2 &\rightarrow_{eqs} \{[[\mathbf{let} \ \mathbf{rec} \ x_1 : T = \mathbf{function} \ (x_2 : T') \rightarrow e_1 \ \mathbf{in} \ x_1]_{eqs}^T/x_1]\mathbf{function} \ (x_2 : T') \rightarrow e_1]_{eqs}^T/x_1\}e_2 \\ C_{eqs}.\mathbf{raise} \ v^{eqs} &\rightarrow_{eqs} \mathbf{raise} \ v^{eqs} \quad (c) \\ [\mathbf{raise} \ v^{eqs'}]_{eqs'}^T &\rightarrow_{eqs} \mathbf{raise} \ [v^{eqs'}]_{eqs'}^{\text{exn}} \\ \mathbf{try} \ \mathbf{raise} \ v^{eqs} \ \mathbf{with} \ p_1 \rightarrow e_1 | \dots | p_n \rightarrow e_n &\rightarrow_{eqs} \text{matchsub}_{eqs}(v^{eqs}, p_k) e_k \quad (a) \\ \mathbf{try} \ v^{eqs} \ \mathbf{with} \ p_1 \rightarrow e_1 | \dots | p_n \rightarrow e_n &\rightarrow_{eqs} v^{eqs} \\ \mathbf{marshal} \ \text{MK} \ v^{eqs} : T &\rightarrow_{eqs} \mathbf{marshalz} \ \text{MK} \ [v^{eqs}]_{eqs}^T : T \end{aligned}$$

(a) $\text{matchsub}_{eqs}(v^{eqs}, p_k)$ is defined and there is no $k' < k$ with $\text{matchsub}_{eqs}(v^{eqs}, p_{k'})$ defined

- (b) (i) not exists $k \in 1..n$ such that $\text{matchsub}_{eqs}(v^{eqs}, p_k)$ is defined, and
(ii) v'^{eqs} is an arbitrary value such that $\vdash v'^{eqs} : \text{string} * \text{int} * \text{int}$
(c) there does not exist $p_1 \rightarrow e_1 | \dots | p_n \rightarrow e_n$ and k st $C_{eqs} = \text{try} - \text{with } p_1 \rightarrow e_1 | \dots | p_n \rightarrow e_n$ and $\text{matchsub}_{eqs}(v^{eqs}, p_k)$ defined

Comment: Note that in several places the semantics involves not-quite-value substitutions: substitutions of a value surrounded by an extra pair of brackets. Bracket reduction is effect-free and terminating, so this is not a problem – it would be notationally awkward to reduce before substituting.

Bracket-pushing (administrative) reductions

Brackets are used to represent abstraction boundaries, and their location is carefully controlled during evaluation. Brackets are purely annotations, however, and the semantics obtained by erasing them corresponds precisely to the given (coloured) semantics. An implementation will typically use the erased semantics; the coloured semantics and the correspondence property serve to give confidence that the implementation respects abstraction boundaries.

Frequently desired reductions will involve subterms on both sides of an abstraction boundary - for example, applying a module function $M_M.z$ to a value outside that module will yield a term $[\text{function } x : T \rightarrow e]_{eqs}^{T' \rightarrow T''} v$. Rather than give reduction rules for each permutation of brackets, we give reduction rules only for the bracket-free case, and add administrative reductions to move brackets out of the way. Specifically, erased-values (that is, terms that correspond to values in the erased semantics) may not yet be values in the coloured semantics; to make them so, we push the brackets inwards by application of the bracket-pushing rules. In the example above, we may push the brackets inwards through the lambda, to obtain $(\text{function } x : T' \rightarrow \{[x]_{T'}^{eqs}/x\}e)_{\emptyset}^{T''} v$, and now the ordinary function application rule yields $\{[v]_{T'}^{eqs}/x\}e_{\emptyset}^{T''}$. These administrative reductions apply only to erased-values.

The bracket-pushing rules are as follows:

pushing through constructors:

$$\begin{array}{ll}
 [[]_{eqs'}^T \text{ list} & \rightarrow_{eqs} []_T \\
 [\text{NONE}_{T'}]_{eqs'}^T \text{ option} & \rightarrow_{eqs} \text{NONE}_T \\
 [\text{INJ}_i^{(T'_1 + \dots + T'_n)} v_{eqs'}]_{eqs'}^{T_1 + \dots + T_n} & \rightarrow_{eqs} \text{INJ}_i^{(T_1 + \dots + T_n)} [v_{eqs'}]_{eqs'}^{T_i} \\
 [v_1^{eqs'} \dots v_n^{eqs'}]_{eqs'}^T \text{ list} & \rightarrow_{eqs} [v_1^{eqs'}]_{eqs'}^T \dots [v_n^{eqs'}]_{eqs'}^T \text{ list} \\
 [(v_1^{eqs'}, \dots, v_n^{eqs'})]_{eqs'}^{T_1 * \dots * T_n} & \rightarrow_{eqs} ((v_1^{eqs'}]_{eqs'}^{T_1}, \dots, [v_n^{eqs'}]_{eqs'}^{T_n}) \quad n \geq 2 \\
 [C^n v_1^{eqs'} \dots v_n^{eqs'}]_{eqs'}^{T_0} & \rightarrow_{eqs} C^n [v_1^{eqs'}]_{eqs'}^{T_1} \dots [v_n^{eqs'}]_{eqs'}^{T_n} \\
 & C^n : T_1 \rightarrow \dots \rightarrow T_n \rightarrow T_0 \text{ any other constructor}
 \end{array}$$

pushing through lambda:

$$[\text{function } (x : T) \rightarrow e]_{eqs'}^{T' \rightarrow T''} \rightarrow_{eqs} \text{function } (x : T') \rightarrow \{[x]_{eqs'}^T/x\}e_{eqs'}^{T''}$$

pushing through type-lambda and pack:

$$\begin{array}{ll}
 [\Lambda t \rightarrow e]_{eqs'}^{\forall t, T} & \rightarrow_{eqs} \Lambda t \rightarrow [e]_{eqs'}^T \\
 \{T, v^{eqs'}\} \text{ as } T' \exists t, T'' & \rightarrow_{eqs} \{T, [v^{eqs'}]_{eqs'}^{\{T/t\}T''}\} \text{ as } \exists t, T''
 \end{array}$$

bracket type revelation:

$$[v^{eqs'}]_{eqs'}^{h.t} \rightarrow_{eqs} [v^{eqs'}]_{eqs'}^T \quad (h.t \approx T) \in eqs \wedge h.t \in \text{dom}(eqs')$$

bracket elimination:

$$[[v^{eqs''}]_{eqs''}^{h.t}]_{eqs'}^{h.t} \rightarrow_{eqs} [v^{eqs''}]_{eqs''}^{h.t} \quad h.t \notin \text{dom}(eqs')$$

It is straightforward to show that all these rules are type-preserving.

Comment: Note that brackets are handled specially in the case of names and store locations; there are no bracket-pushing rules for these forms.

Comment: Note that type revelation does not introduce non-termination, because the equation formation rules ensure that for any equation $X.t \approx T$, T is well-formed in the environment prior to the definition of X .

These rules ensure that any erased-value can be reduced to a (coloured) value, by pushing brackets inward as far as they can go, and eliminating double brackets.

Note that the rule for pushing brackets through a **function** depends on the fact that functions bind a single argument identifier, functions with more complex patterns being treated as syntactic sugar for a single-argument function with a **match** as the body. Without this, bracket pushing would have to be much more elaborate.

Store- and name-related bracket-pushing

Bracket handling for locations, dereferencing, assignment, and names is subtle. Notice, for example, that a module may return a location to its caller at an abstract type, and allow the caller to store abstract values in it, and then internally pull them out at the concrete one. Worse, a module may create a ref cell, and return its location twice, once at an abstract type and once at a concrete type. There seems no good reason to prohibit this arbitrary aliasing of pointers, where each alias may have different type transparency depending on the locally available *eqs*. In this respect we differ from Zdancewic et al. [GMZ00, §4.2]. For names the issue is simply that a name records its type as that at which it was created, but use at multiple types, according to context, must be permitted.

Since `ref` is treated as a vanilla operator (brackets are not involved), we do not discuss its semantics here.

In the value grammar we allow names and locations to be wrapped in brackets in order to express the variety of type transparency that aliases of the name or location may have. Thus, if we have a bracketted (!) or (:=), we *pull* the brackets outside, changing the type annotations accordingly. The goal is to peel away the brackets surrounding a location so as to expose the location itself to dereference or assignment:

$$\begin{array}{l} !_T [v^{eqs'}]_{eqs'}^{T'} \text{ref} \quad \rightarrow_{eqs} \quad [!_{T'} v^{eqs'}]_{eqs'}^{T'} \\ [v^{eqs'}]_{eqs'}^{T'} \text{ref} :=_T v^{eqs} \quad \rightarrow_{eqs} \quad [v^{eqs'} :=_{T'} [v^{eqs}]_{eqs'}^{T'}]_{eqs'}^{\text{unit}} \end{array}$$

Comment: When bracket pulling through $!_T$ it is not immediately obvious why the bracket on the RHS is at T' and not T . It is correct (even though the type of the whole expression must be T) because we may deduce from the LHS that $E \vdash_{eqs} T \approx T'$, and it is necessary because we cannot deduce $E \vdash_{eqs'} T \approx T'$, which would be needed in order to type the alternative.

Values in the store are always black (v^\emptyset). When we get a raw location, $!_T$ can dereference it:

$$E_n ; \langle E_s, (s, l \mapsto v^\emptyset), \text{definitions}, !_T l \rangle \rightarrow_{eqs} E_n ; \langle E_s, (s, l \mapsto v^\emptyset), \text{definitions}, v^\emptyset \rangle$$

Comment: Note that the correctness of this rule relies on the fact that typing is monotonic with respect to the *eqs* set. By hypothesis, $E_n, E_s \vdash_{\emptyset} v^\emptyset : T_0$ where $E_s(l) = T_0$ and $E_n, E_s \vdash_{eqs} \text{ok}$ and $E_n, E_s \vdash_{eqs} T_0 \approx T$. This implies $E_n, E_s \vdash_{eqs} v^\emptyset : T_0$ since having more equalities can't hurt, hence $E_n, E_s \vdash_{eqs} v^\emptyset : T$ as desired.

When we get a raw location, $:=_T$ prepares the value to be put in the store; when that value becomes a value in \emptyset (see the discussion of operator reduction below), we can install it in the store:

$$\begin{array}{l} l :=_T v^{eqs} \quad \rightarrow_{eqs} \quad l :='_T [v^{eqs}]_{eqs}^T \\ E_n ; \langle E_s, (s, l \mapsto v'^\emptyset), \text{definitions}, l :='_T v^\emptyset \rangle \rightarrow_{eqs} E_n ; \langle E_s, (s, l \mapsto v^\emptyset), \text{definitions}, () \rangle \end{array}$$

For names there is no other argument to which the brackets must be transferred; instead, we define all operators which operate on names to ignore brackets surrounding those names.

Comment: In the current semantics, we treat $[v^{eqs'}]_{eqs'}^{T'} \text{ref}$ as a value in *eqs* for arbitrary instantiations of the metavariables (and similarly for `name`). This is not desirable because it fails to distinguish between those T' `ref` brackets that are really necessary and those that are not. We have some ideas of how to proceed, but for the present leave the removal of this technical infelicity to future work.

Operators and Special Constants Before evaluating the application of a primitive operator or of a primitive constant, we make the arguments be \emptyset -coloured values. Once this is done, we perform the actual evaluation by a delta-rule.

$$\begin{array}{l} e^n v_1^{eqs} \dots v_n^{eqs} \quad \rightarrow_{eqs} \quad \text{op}(e^n)^n [v_1^{eqs}]_{eqs}^{T_1} \dots [v_n^{eqs}]_{eqs}^{T_n} \\ \text{where } e^n : T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \text{ is } x^n \text{ in } E_{\text{const}} \text{ or } \text{op}^n \end{array}$$

Note the rule includes the case $x^0 : T$.

For simple operators the delta rules are as follows:

$$\begin{array}{ll} \mathbf{op}(=_T)^2 v^\varnothing v'^\varnothing \rightarrow_{eqs} \mathbf{true} & \text{erase_brackets}(v^\varnothing) = \text{erase_brackets}(v'^\varnothing) \\ \mathbf{op}(=_T)^2 v^\varnothing v'^\varnothing \rightarrow_{eqs} \mathbf{false} & \text{otherwise} \end{array}$$

with similar rules for the other arithmetic and logical operators (noting that equality raises `INVALID_ARGUMENT` if used on a function or existential package and division can raise `DIVISION_BY_ZERO`).

The delta rule for the reference operator is:

$$E_n ; \langle E_s, s, \text{definitions}, (\mathbf{op}(\mathbf{ref}_T)^1 v^\varnothing) \rangle \rightarrow_{eqs} E_n ; \langle (E_s, l : T \text{ ref}), (s, l \mapsto v^\varnothing), \text{definitions}, l \rangle \quad l \notin \text{dom}(s)$$

Note that the rule for `ref` introduces nondeterminism. That could be avoided by working up to cyclic bindings – seems slightly simpler without, but there is little in it.

Special constants from E_{const} are of two classes. Some have are internal to the language; for these we should have further delta rules, but do not write them here. The others — the x^n such that $\text{os}(x^n)$, which are all of function type — are calls to OS routines. For these we have labelled transitions for invocations and returns:

$$\begin{array}{ll} E_n ; \mathbf{op}(x^n)^n v_1^\varnothing \dots v_n^\varnothing & \xrightarrow{x^n v_1^\varnothing \dots v_n^\varnothing}_{eqs} E_n ; \mathbf{RET}_T \quad (x^n : T_1 \rightarrow \dots T_n \rightarrow T) \in E_{\text{const}} \wedge \text{os}(x^n) \wedge \text{fast}(x^n) \\ E_n ; \mathbf{RET}_T & \xrightarrow{\text{Ok}(v^\varnothing)}_{eqs} E_n ; v^\varnothing \quad \text{if } E_n, E_{\text{const}} \vdash_\varnothing v^\varnothing : T \\ E_n ; \mathbf{RET}_T & \xrightarrow{\text{Ex}(v^\varnothing)}_{eqs} E_n ; \mathbf{raise}(v^\varnothing) \quad \text{if } E_n, E_{\text{const}} \vdash_\varnothing v^\varnothing : \text{exn} \end{array}$$

and

$$\begin{array}{ll} E_n ; \mathbf{op}(x^n)^n v_1^\varnothing \dots v_n^\varnothing & \xrightarrow{x^n v_1^\varnothing \dots v_n^\varnothing}_{eqs} E_n ; \mathbf{SLOWRET}_T \quad (x^n : T_1 \rightarrow \dots T_n \rightarrow T) \in E_{\text{const}} \wedge \text{os}(x^n) \wedge \neg \text{fast}(x^n) \\ E_n ; \mathbf{SLOWRET}_T & \xrightarrow{\text{Ok}(v^\varnothing)}_{eqs} E_n ; v^\varnothing \quad \text{if } E_n, E_{\text{const}} \vdash_\varnothing v^\varnothing : T \\ E_n ; \mathbf{SLOWRET}_T & \xrightarrow{\text{Ex}(v^\varnothing)}_{eqs} E_n ; \mathbf{raise}(v^\varnothing) \quad \text{if } E_n, E_{\text{const}} \vdash_\varnothing v^\varnothing : \text{exn} \end{array}$$

Comment: The semantics allows the OS return values to be typed with respect to E_n, E_{const} , though for the extant OS call types this makes no difference.

Termination For program termination we have the axiom below.

$$E_n ; \langle E_s, s, \text{definitions}, P \rangle \rightarrow \mathbf{TERM}$$

where either (a) there are no threads $\mathbf{n} : \text{definitions } e$ in P , or (b) there is at least one thread $\mathbf{n} : \text{definitions } e$ in P but for all such we have \mathbf{n} internally blocked in P .

Case (b) is a useful and sound but very coarse approximation to deadlock detection. In this case the implementation prints a warning.

In addition the programmer can force termination with `exit` as below.

$$E_n ; \langle E_s, s, \text{definitions}, P | \mathbf{n} : TCC_{eqs} . \mathbf{op}(\mathbf{exit}) v \rangle \rightarrow \mathbf{TERM}$$

Comment: At present we do not distinguish between successful and unsuccessful termination — cf. the thread exception semantics, which specifies that threads that reduce to a value or to a raised exception silently exit.

16.8.5 Marshalling and unmarshalling

Marshalling

Here we define the reduction step for $E_n ; \langle E_s, s, definitions, \mathbf{marshalz} \text{ MK } v^\emptyset : T \rangle$, where

$$\begin{aligned} definitions &= definitions_1 ;; \mathbf{mark} \text{ MK} ;; definitions_2 \\ \mathbf{mark} \text{ MK} &\notin definitions_2 \end{aligned}$$

that constructs a marshalled value.

If there is no $\mathbf{mark} \text{ MK}$ in $definitions$, we fail with $E_n ; \langle E_s, s, definitions, \mathbf{raise} \text{ MARSHAL_FAILURE} \rangle$.

In outline, what we do is prune $definitions_2$, omitting any modules that are not needed and on the way calculating which modules from $definitions_1$ we refer to. We then go through $definitions_1$ making up an import for each of those (this does not unload imports in $definitions_2$ that point within $definitions_1$, instead generating an additional import at the boundary).

Note that this does not involve any $definitions$ of the executing thread.

Write $\text{fmv}(\dots)$ for the set of free module external/internal identifier pairs in a gadget (note hashes are all fmv -closed). We make explicit some interesting cases of fmv (first on terms, then on Mos):

$$\begin{aligned} \text{fmv}(M_M.x) &= \{M_M\} \\ \text{fmv}(h) &= \emptyset \\ \\ \text{fmv}(M_M) &= \{M_M\} \\ \text{fmv}(\text{UNLINKED}) &= \emptyset \end{aligned}$$

Write $\text{locs}(\dots)$ for the set of locations occurring in a gadget.

Now, given the configuration above, with its $definitions$ and s , define a reachability relation \rightsquigarrow over the union of the set of M_M defined by the $definitions$ and the l in the domain of the store s .

- For M_M defined in $definitions_2$ by $definition = (\mathbf{cmodule}_{h;Sig_0;vubs} M_M : Sig_1 \mathbf{version} \text{ vn} = Str)$:

$$\begin{aligned} M_M &\rightsquigarrow M'_{M'} \quad \mathbf{if} \quad M'_{M'} \in \text{fmv}(Str) \\ M_M &\rightsquigarrow l \quad \mathbf{if} \quad l \in \text{locs}(definition) \end{aligned}$$

- For M_M defined in $definitions_2$ by $definition = (\mathbf{cimport}_{h;Sig_0} vubs M_M : Sig_1 \mathbf{version} \text{ vc like } Str \mathbf{by} \text{ resolvespec} = Mo)$:

$$\begin{aligned} M_M &\rightsquigarrow M'_{M'} \quad \mathbf{if} \quad Mo = M'_{M'} \\ M_M &\rightsquigarrow l \quad \mathbf{if} \quad l \in \text{locs}(definition) \end{aligned}$$

- For l ,

$$\begin{aligned} l &\rightsquigarrow M'_{M'} \quad \mathbf{if} \quad M'_{M'} \in \text{fmv}(s(l)) \\ l &\rightsquigarrow l' \quad \mathbf{if} \quad l' \in \text{locs}(s(l)) \end{aligned}$$

(Note there are no clauses for M_M defined in $definitions_1$. Note that in the import case the free module identifiers of the Str will always be empty, as it is a struct that consists exclusively of hashified types, and, similarly, in both cases the free module identifiers of the signatures will be empty.) Let A be the smallest set containing $fmv(v^\emptyset) \cup locs(v^\emptyset)$ and closed under \rightsquigarrow .

Let S_1 , S_2 , and L be the partition of A into its module identifiers defined by $definitions_1$, those defined by $definitions_2$, and the locations.

Let $definitions'_2$ be the subsequence of $definitions_2$ containing the definitions of modules in S_2 together with all **mark s**.

Now `makeimports definitions S` constructs imports for the needed modules based on their compiled $definitions_1$.

```
makeimports(definitions ;; cmoduleh;eqs;Sig0 vubs  $M_M : Sig_1$  version vn = Str)S =
  if  $M_M \in S$  then
    (makeimports(definitions)(S - { $M_M$ })) ;;
    cimporth;Sig0  $M_M : Sig_1$  version name = h like filter Str(limitdom(Sig0))
    by HERE_ALREADY = UNLINKED
  )
  else
    makeimports definitions S
```

```
makeimports(definitions ;; cimporth;Sig0  $M_M : Sig_1$  version vc like Str by resolvespec = Mo)S =
  if  $M_M \in S$  then
    (makeimports(definitions)(S - { $M_M$ })) ;;
    cimporth;Sig0  $M_M : Sig_1$  version vc like Str by resolvespec = UNLINKED
  )
  else
    makeimports definitions S
```

```
makeimports(definitions ;; mark MK)S =
  makeimports(definitions)S
```

```
makeimports(empty)S = empty
```

Comment: You might think that in the module-initialisation world, reachability would need to go via earlier fields of this module (occurrences of x under a lambda, say) and via top-level $definitions$ ($M_M.x$, say). However, see module field instantiation, internal field case: we have chosen an x -substitution semantics, and so the former case does not arise (x has been substituted away by the time we reach it).

Let $definitions' = \text{makeimports}(definitions_1)S_1 ;; definitions'_2$

Below we write $X \downarrow L$ for X restricted to L . Let $E_{s'} = E_s \downarrow L$. Let $s' = s \downarrow L$.

Let E_n' be the smallest subsequence of E_n including all the abstract names of E'_s , s' , $definitions'$, v^\emptyset , T and all **nmodules** in E_n' .

The E_n' can be omitted in a production implementation.

Note that marshalling preserves all the original marks we pass through in $definitions_2$, putting them in $definitions'$ and thus in the marshalled value. It does not include the MK we are marshalling with respect to.

Finally, then, we have:

$$\begin{aligned} & E_n ; \langle E_s, s, definitions, \text{marshalz MK } v^\emptyset : T \rangle \\ \rightarrow_{eqs} & E_n ; \langle E_s, s, definitions, \underline{s} \rangle \end{aligned}$$

where

$$\text{raw_unmarshal}(\underline{s}) = \mathbf{marshalled}(E_n', E_{s'}, s', \text{definitions}', v'^{\emptyset}, T)$$

If marshal-time typechecking is specified, additionally check $\vdash \mathbf{marshalled}(E_n', E_{s'}, s', \text{definitions}', v'^{\emptyset}, T)$ **ok**. Fail with `RUN.TYPECHECK_ON_MARSHAL` otherwise.

Unmarshalling

Choosing here to do linking as late as possible, so not doing any linking at unmarshal-time.

$$\begin{aligned} & E_n ; \langle E_s, s, \text{definitions}, \mathbf{unmarshal} \ \underline{s} \ \text{as} \ T \rangle \\ \rightarrow_{eqs} & E_n'' ; \langle (E_s, \sigma \ E_{s'}), (s, (\sigma \ s') \cdot \sigma^{-1}), (\text{definitions} \ ;\ ; \ \text{definitions}'), (\sigma \ (v'^{\emptyset})) \rangle \end{aligned}$$

where

$$\begin{aligned} \text{raw_unmarshal}(\underline{s}) &= \mathbf{marshalled}(E_n', E_{s'}, s', \text{definitions}', v'^{\emptyset}, T') \\ \text{the module binders of } &\text{definitions}' \text{ are distinct from those of } \text{definitions} \\ \sigma \text{ is a location injection with domain } &\text{dom}(s') \text{ and with } \text{ran}(\sigma) \text{ disjoint from } \text{dom}(s) \\ T &= T' \\ E_n'' &= \text{merge_nenvs}(E_n, E_n') \end{aligned}$$

(writing $\sigma \ X$ for the result of applying σ as a substitution to X , and so $\sigma \ s'$ for the result of doing that pointwise to the range of s').

As usual, the calculation of E_n'' is superfluous if we are not doing run-time type checking.

Note that marshalled modules are always fully evaluated so at unmarshal-time they can be added to the per-runtime definitions not to the thread definitions.

If marshal-time typechecking is specified, additionally check $\vdash \mathbf{marshalled}(E_n', E_{s'}, s', \text{definitions}', v'^{\emptyset}, T')$ **ok**, and check that the `merge_nenvs`(E_n, E_n') above succeeds. Fail with `RUN.TYPECHECK_ON_UNMARSHAL` otherwise.

$$\begin{aligned} & E_n ; \langle E_s, s, \text{definitions}, \mathbf{unmarshal} \ \underline{s} \ \text{as} \ T \rangle \\ \rightarrow_{eqs} & E_n ; \langle E_s, s, \text{definitions}, \mathbf{raise} \ \text{UNMARSHAL_FAILURE} \ \underline{s}' \rangle \end{aligned}$$

where `raw_unmarshal`(\underline{s}) undefined, or `raw_unmarshal`(\underline{s}) = $\mathbf{marshalled}(E_n', E_{s'}, s', \text{definitions}', v'^{\emptyset}, T')$ and $\neg T = T'$. \underline{s}' is a string describing the cause of the unmarshal failure.

Comment: Note that unmarshalling will cause existing marks to be shadowed by the marks contained in `definitions'`. This is sometimes desirable, but not always – really, this is a defect of the linear mark/module structure.

Comment: Note that marshalling permits one to see through abstraction boundaries in limited fashion, by equality testing (or even more detailed examination) of the marshalled strings for abstract types.

16.8.6 Module field instantiation

Module field instantiation – module case, via import sequence

$$E_n ; \langle E_s, s, definitions, M_M.x \rangle \rightarrow_{eqs} E_n ; \langle E_s, s, definitions, [v']_{eqs'}^T \rangle$$

where

$definitions = definitions_0 ;; definition ;; definitions_1 ;; definition_1 ;; definitions_2 ;; \dots ;; definition_n ;; definitions_n$
 $definition = \mathbf{cmodule}_{h;eqs_0;Sig_0} vubs M_{0M_0} : Sig_1 = Str$
 $\forall i \in 1..n. \quad definition_i = \mathbf{cimport}_{h_i;Sig_{0_i}} M_{iM_i} : Sig_{1_i} \mathbf{version} vc_i \mathbf{like} Str_i \mathbf{by} resolvespec_i = M_{i-1}M_{i-1}$
 $M_M = M_{nM_n}, \quad definitions_n \text{ doesn't_define } M_M$
 $(\mathbf{val} x_x : T) \in Sig_{1_n}$
 $(\mathbf{let} x_x = v^{eqs_0}) \in Str$
 $v' = v$
 $eqs' = eqs_0, eqs_of_sign_str(h, Sig_0, Str), eqs_of_sign_str(h_n, Sig_{0_n}, Str_n)$

Comment: Note that we include `eqs_of_sign_str` from the `cmodule` and from the ultimate `cimport` (if any), not from any intermediate imports.

Comment: There are two choices here, dependent on the module initialisation semantics. Before, module values could have an expression value field containing free expression identifiers of earlier fields. Then the v' here had to be mutated, taking v with each y free in v replaced by $M_{nM_n}.y$ (if $n > 0$) or by $M_M.y$ (if $n = 0$). Now module initialisation substitutes out fields as it goes, so this is no longer needed.

Note that in the earlier semantics that term selfification of the value v depends on that fact that the signature check in `linkok` does not allow width subsignaturing. What the behaviour should be if one allowed width subsignaturing is unclear.

Module field instantiation – unlinked import case; start looking

$$E_n ; \langle E_s, s, definitions, M_M.x \rangle \rightarrow_{eqs} E_n ; \langle E_s, s, definitions, \mathbf{resolve}(M_M.x, M_{0M_0}, resolvespec) \rangle$$

where

$definitions = definitions_0 ;; definition_0 ;; definitions_1 ;; definition_1 ;; definitions_2 ;; \dots ;; definition_n ;; definitions_n$
 $definition = \mathbf{cimport}_{h_0;Sig_0} M_{0M_0} : Sig_1 \mathbf{version} vc_0 \mathbf{like} Str_0 \mathbf{by} resolvespec_0 = \mathbf{UNLINKED}$
 $\forall i \in 1..n. \quad definition_i = \mathbf{cimport}_{h_i;Sig_{0_i}} M_{iM_i} : Sig_{1_i} \mathbf{version} vc_i \mathbf{like} Str_i \mathbf{by} resolvespec_i = M_{i-1}M_{i-1}$
 $M_M = M_{nM_n}$
 $definitions_n \text{ doesn't_define } M_M$
 $(\mathbf{val} x_x : T) \in Sig_{1_n}$

Module field instantiation – resolve URI

$$\frac{\text{GetURI}(URI)}{\rightarrow_{eqs}} E_n ; \langle E_s, s, definitions, \mathbf{resolve}(M_M.x, M'_{M'}, (URI, resolvespec)) \rangle$$

Module field instantiation – resolve case, HERE_ALREADY

$$E_n ; \langle E_s, s, definitions, \mathbf{resolve}(M_M.x, M'_{M'}, (\mathbf{HERE_ALREADY}, resolvespec_0)) \rangle \rightarrow_{eqs} E_n ; \langle E_s, s, definitions_9, e \rangle$$

where

$definitions = definitions_1 ;; definition ;; definitions_2$
 $definitions_2 \text{ doesn't_define } M'_{M'}$
 $definition = \mathbf{cimport}_{h;Sig_0} M'_{M'} : Sig_1 \mathbf{version} vc \mathbf{like} Str \mathbf{by} resolvespec = \mathbf{UNLINKED}$

Let Ms be the sequence of the $M''_{M''}$ defined by a $definition'$ in $definitions_1$ satisfying $linkok(E_n, definition', definition \oplus (= M'_{M'}))$, where $definition \oplus (= M'_{M'})$ is as $definition$ but with $= M'_{M'}$ replacing $= UNLINKED$.

If Ms is nonempty then, taking $M''_{M''}$ to be its last element,

$$\begin{aligned} definitions_9 &= definitions_1 ;; definition_9 ;; definitions_2 \\ definition_9 &= \mathbf{cimport}_{h;Sig_0} M'_{M'} : Sig_1 \mathbf{version} \mathit{vc} \mathbf{like} \mathit{Str} \mathbf{by} \mathit{resolvespec} = M''_{M''} \\ e &= M_{M.x} \end{aligned}$$

otherwise if Ms is empty take

$$definitions_9 = definitions \text{ and } e = \mathbf{resolve}(M_{M.x}, M'_{M'}, \mathit{resolvespec}_0).$$

Module field instantiation – resolve case, STATIC_LINK

$$\begin{aligned} &E_n ; \langle E_s, s, definitions, \mathbf{resolve}(M_{M.x}, M'_{M'}, (\mathbf{STATIC_LINK}, \mathit{resolvespec}_0)) \rangle \\ \rightarrow_{eqs} &E_n ; \langle E_s, s, definitions, \mathbf{raise} \mathbf{RESOLVE_FAILURE} \rangle \end{aligned}$$

The intention for imports with $\mathbf{STATIC_LINK}$ $\mathit{resolvespecs}$ is that they should have been statically linked, so if we reach one at runtime it is an error. We do not yet define a separate static linking phase, however, so they are not yet very useful.

Module field instantiation – resolveblocked case, got some $definitions'$

$$\frac{\text{DeliverURI}(E_n', definitions')}{\rightarrow_{eqs}} \begin{aligned} &E_n ; \langle E_s, s, definitions, \mathbf{resolve_blocked}(M_{M.x}, M'_{M'}, \mathit{resolvespec}_0) \rangle \\ &E_{n9} ; \langle E_s, s, definitions_9, e \rangle \end{aligned}$$

where

$$\begin{aligned} definitions &= definitions_1 ;; definition ;; definitions_2 \\ definitions_2 \text{ doesn't_define } M'_{M'} & \\ definition &= \mathbf{cimport}_{h;Sig_0} M'_{M'} : Sig_1 \mathbf{version} \mathit{vc} \mathbf{like} \mathit{Str} \mathbf{by} \mathit{resolvespec} = UNLINKED \end{aligned}$$

Let $E_{n9} = \text{merge_nenvs}(E_n, E_n')$. This is superfluous if we are not doing run-time type checking.

Note that we disallow module field instantiation with non-value definitions. This ensures the new $definitions'$ can be inserted into the existing per-runtime $definitions$ before the $\mathbf{cimport}$ which must be linked to them, without breaking the invariant that the per-runtime $definitions$ are always fully evaluated. To relax this would need additional mechanism to block instantiation from a linked but not-yet-evaluated module.

If

- $definitions'$ consists only of definition values.
- $E = E_1(definitions')$
- $\text{dom}(definitions') \cap \text{dom}(definitions) = \emptyset$ (achievable by alpha equivalence)
- Letting Ms be the sequence of the $M''_{M''}$ defined by a $definition'$ in $definitions'$ satisfying $linkok(E_{n9}, definition', definition \oplus (= M'_{M'}))$, we have Ms nonempty with a last element $M''_{M''}$.

then

$$\begin{aligned} definitions_9 &= definitions_1 ;; definitions' ;; definition_9 ;; definitions_2 \\ definition_9 &= \mathbf{cimport}_{h;Sig_0} M'_{M'} : Sig_1 \mathbf{version} \mathit{vc} \mathbf{like} \mathit{Str} \mathbf{by} \mathit{resolvespec} = M''_{M''} \\ e &= M_{M.x} \end{aligned}$$

else

$definitions_0 = definitions$ and $e = \mathbf{resolve}(M_M.x, resolvespec_0)$.

In the implementation, if a byte string that does not lex or parse as a $definitions'$ is returned for this URI it is treated as a CannotFindURI transition.

If doing run-time type checking, check additionally $E_{const} \vdash definitions' \triangleright E$ and $\mathit{linkok}(E_n', definitions')$ and that the merge_nenvs succeeds (or fail with TYPECHECK_ON_GET_URI).

Module field instantiation – resolveblocked case, didn't get any $definitions'$

$$\frac{\text{CannotFindURI}}{\rightarrow_{eqs}} \quad E_n ; \langle E_s, s, definitions, \mathbf{resolve_blocked}(M_M.x, M'_{M'}, resolvespec) \rangle$$

$$E_n ; \langle E_s, s, definitions, \mathbf{resolve}(M_M.x, M'_{M'}, resolvespec) \rangle$$

Module field instantiation – run out of $resolvespec$

$$E_n ; \langle E_s, s, definitions, \mathbf{resolve}(M_M.x, M'_{M'}, \mathbf{empty}) \rangle \rightarrow_{eqs} E_n ; \langle E_s, s, definitions, \mathbf{raise} \text{ RESOLVE_FAILURE} \rangle$$

16.8.7 Name operations

We write $\mathit{fn}(v)$ for the set of names \mathbf{n} occurring in v and $\mathit{fns}(v)$ for the set of simple names \mathbf{nn} occurring in v . The primitive swapping function $\mathit{SSwap}_{eqs}(BC_1.\mathbf{nn}_1, BC_2.\mathbf{nn}_2)$ in v yields the result of replacing, in v , all occurrences of \mathbf{nn}_i with $\mathit{revbc}_{eqs}(BC_i).BC_{2-i}.\mathbf{nn}_{2-i}$. These are defined homomorphically through the abstract syntax, save that they do *not* propagate through $\mathit{hash}(\dots)$. (Note that $\mathit{marshalled}(\dots)$ does not occur (hereditarily) in the abstract syntax for expressions). The auxiliary function revbc reverses the sequence of brackets in a bracket context, and is defined as follows:

$$\mathit{revbc}_{eqs}(-) = -$$

$$\mathit{revbc}_{eqs}([_]^{T'} \cdot BC) = \mathit{revbc}_{eqs'}(BC) \cdot [_]^{T'}$$

Given a configuration with $definitions$ and s , define a reachability relation \rightsquigarrow over the domain of the store s .

- $l \rightsquigarrow l'$ if $l' \in \mathit{locs}(s(l))$

Let the reachable locations from a value v^{eqs} be the smallest set A containing $\text{locs}(v^{eqs})$ and closed under \rightsquigarrow .

$$\frac{\mathbf{n} \notin \text{dom}(E_n)}{E_n ; \langle E_s, s, \text{definitions}, \text{fresh}_T \rangle \rightarrow_{eqs} E_n, \mathbf{n} : T \text{ name} ; \langle E_s, s, \text{definitions}, \mathbf{n}_T \rangle}$$

L is the set of locations reachable from v^{eqs}
 σ is a location injection with domain L and with $\text{ran}(\sigma)$ disjoint from $\text{dom}(s)$
 $E_{s'} = \sigma(E_s \downarrow L)$
 $s' = \lambda l \in \text{ran}(\sigma). \text{SSwap}_{eqs}(BC.\mathbf{nn}, BC'.\mathbf{nn}') \text{ in } \sigma s(\sigma^{-1}(l))$
 $v_2^{eqs} = \text{SSwap}_{eqs}(BC.\mathbf{nn}, BC'.\mathbf{nn}') \text{ in } \sigma^{-1} v^{eqs}$

$$E_n ; \langle E_s, s, \text{definitions}, \text{swap}(BC.\mathbf{nn}) \text{ and } (BC'.\mathbf{nn}') \text{ in } v^{eqs} \rangle \rightarrow_{eqs} E_n ; \langle (E_s, E_{s'}), (s, s'), \text{definitions}, v_2^{eqs} \rangle$$

$$\frac{\begin{array}{l} L \text{ is the set of locations reachable from } v_2^{eqs} \\ \underline{b} = (\text{erase_brackets}(v_1^{eqs}) \in \text{fns}(v_2^{eqs}) \cup \text{fns}(\text{ran}(s \downarrow L))) \end{array}}{E_n ; \langle E_s, s, \text{definitions}, v_1^{eqs} \text{ freshfor } v_2^{eqs} \rangle \rightarrow_{eqs} E_n ; \langle E_s, s, \text{definitions}, \underline{b} \rangle}$$

$$\frac{\begin{array}{l} L \text{ is the set of locations reachable from } v^{eqs} \\ nset = \text{fns}(v^{eqs}) \cup \text{fn}(\text{ran}(s \downarrow L)) \\ \{\mathbf{n}_1, \dots, \mathbf{n}_k\} = \text{filter}(\lambda \mathbf{n}. \text{typeof}(\mathbf{n}) = T) nset \\ \forall i \neq j. \mathbf{n}_i \neq \mathbf{n}_j \\ v'^{eqs} = \mathbf{n}_1 :: \dots :: \mathbf{n}_k :: []_T \text{ name} \end{array}}{E_n ; \langle E_s, s, \text{definitions}, \text{support}_T v^{eqs} \rangle \rightarrow_{eqs} E_n ; \langle E_s, s, \text{definitions}, v'^{eqs} \rangle}$$

hash ($T', \underline{s}, v^\emptyset$)	\rightarrow_{eqs}	hash ($T', \underline{s}, \mathbf{n}$)	$\text{erase_brackets}(v^\emptyset) = \mathbf{n} \wedge v^\emptyset \neq \mathbf{n}$
compare_name $_T v_1^\emptyset v_2^\emptyset$	\rightarrow_{eqs}	0	$\text{erase_brackets}(v_1^\emptyset) = \text{erase_brackets}(v_2^\emptyset)$
compare_name $_T v_1^\emptyset v_2^\emptyset$	\rightarrow_{eqs}	-1	$\text{erase_brackets}(v_1^\emptyset) < \text{erase_brackets}(v_2^\emptyset)$
compare_name $_T v_1^\emptyset v_2^\emptyset$	\rightarrow_{eqs}	1	$\text{erase_brackets}(v_1^\emptyset) > \text{erase_brackets}(v_2^\emptyset)$
name_of_tie $\text{TIECON}(v_1^{eqs}, v_2^{eqs})$	\rightarrow_{eqs}	v_1^{eqs}	
val_of_tie $\text{TIECON}(v_1^{eqs}, v_2^{eqs})$	\rightarrow_{eqs}	v_2^{eqs}	

Comment: Note that (in contrast to FreshOCaml) reachability here does go through the store.

Comment: Note that this makes the semantics of **support** potentially surprising: **support** $_T e$ is the set of names in e that were *constructed at type T*, not all those that have type T in the present context. A positive consequence is that the reduction rule for **support** is simple to implement because it is independent of the presence of brackets, thus the same reductions are obtained after bracket erasure.

A negative consequence is that the rule fails to account for the type equalities introduced by the brackets surrounding a name, thus possibly not collecting all the relevant names; the rule can also see through abstraction boundaries, thus possibly collecting too many names.

Comment: With module initialisation, one has to decide whether reachability goes through *definitions*, e.g. if you have a store location containing **function** $() \rightarrow M_{M.x}$ and $M_{M.x}$ has either a store location or a **function** $() \rightarrow M'_{M'.x'}$. Here we choose not — note that this is a different notion of reachability from that used in marshalling.

Comment: Note that we treat fresh and hash-generated names uniformly here, allowing swapping etc. over (and between) both.

Comment: We can (indirectly) send and receive modules, but we have no way of swapping over them. This is clearly suspicious and is one more point in favour of more first-class modules.

Comment: Note that the polytypic **swap**, **support** and **freshfor** can see through abstraction boundaries.

Comment: One might want **support** to return a duplicate-free list w.r.t. the embedded simple names, not just w.r.t. **n** equality.

16.8.8 Concurrency

Basic thread operations: termination, create_thread, self, kill

Below we write $\mathbf{n} =_{\text{erased}} \mathbf{n}'$ for $\text{erase_brackets}(\mathbf{n}) = \text{erase_brackets}(\mathbf{n}')$, $\mathbf{n} \in_{\text{erased}} \text{nset}$ for $\text{erase_brackets}(\mathbf{n}) \in \text{erase_brackets}(\text{nset})$ and similarly for \notin_{erased} .

$P|\mathbf{n} : v^\emptyset \rightarrow P|0$

$P|\mathbf{n} : TCC_{eqs}.\text{raise } v^\emptyset \rightarrow P|0$

$P|\mathbf{n} : TCC_{eqs}.\text{op}(\text{create_thread})^3 \mathbf{n}' v_1^\emptyset v_2^\emptyset \rightarrow P|\mathbf{n} : TCC_{eqs}.\text{op}(\text{create_thread})^3 \mathbf{n}' : (v_1^\emptyset v_2^\emptyset) \quad \mathbf{n}' \notin_{\text{erased}} \{\mathbf{n}\}, \text{dom}(P)$

$P|\mathbf{n} : TCC_{eqs}.\text{op}(\text{self})^1 () \rightarrow P|\mathbf{n} : TCC_{eqs}.\mathbf{n}$

$P|\mathbf{n}' : e_1|\mathbf{n} : TCC_{eqs}.\text{op}(\text{kill})^1 \mathbf{n}' \rightarrow P|0|\mathbf{n} : TCC_{eqs}.\text{op}(\text{kill})^1 \mathbf{n}'$

$P|\mathbf{n} : TCC_{eqs}.\text{op}(\text{kill})^1 \mathbf{n} \rightarrow P|0$

$P|\mathbf{n} : TCC_{eqs}.\text{op}(\text{create_thread})^3 \mathbf{n}' v_1^\emptyset v_2^\emptyset \rightarrow P|\mathbf{n} : TCC_{eqs}.\text{raise EXISTENT_NAME } \mathbf{n}' \in_{\text{erased}} \{\mathbf{n}\}, \text{dom}(P)$

$P|\mathbf{n} : TCC_{eqs}.\text{op}(\text{kill})^1 \mathbf{n}' \rightarrow P|\mathbf{n} : TCC_{eqs}.\text{raise NONEXISTENT_THREAD } \mathbf{n}' \notin_{\text{erased}} \{\mathbf{n}\}, \text{dom}(P)$

Comment: At present threads terminate silently, in both value and raised-exception cases. An alternative to the latter would be $P|\text{raise } v^\emptyset \rightarrow 0$, which is more fail-stop (a good thing, in principle) but possibly more annoying. At very least, a thread dying by with a raised exception should currently generate a warning to the console. Ultimately we should perhaps arrange some way to have exception handlers around threads.

Comment: We allow a thread can kill itself — may not make much difference, but this seems slightly more intuitive than the alternative of raising an exception. This is a difference from the **thunkify** semantics.

Comment: **kill** and **thunkify** are both dangerous operations in that they can remove threads which hold mutexes, which will then never be released. We expect them to be used only within the implementations of libraries that provide both **kill**- or **thunkify**-like operations together with safe thread interaction constructs. Otherwise, the preferred idiom for killing a thread should be to ask it to kill itself; it can then exit cleanly.

Thunkify

Say an e is an *atomic internal blocked form* in P if e is either $\text{op}(\text{lock})^1 \mathbf{n}'$ with $\mathbf{n}'' : \text{MX}(\text{true})$ also in P for $\mathbf{n}'' =_{\text{erased}} \mathbf{n}'$, or $\text{op}(\text{waiting})^2 \mathbf{n}' \mathbf{n}''$.

Say an e is an *atomic blocked form* in P if e is either $\text{SLOWRET}_T, \text{op}(\text{lock})^1 \mathbf{n}'$ with $\mathbf{n}'' : \text{MX}(\text{true})$ also in P for some $\mathbf{n}'' =_{\text{erased}} \mathbf{n}'$, or $\text{op}(\text{waiting})^2 \mathbf{n}' \mathbf{n}''$, or $\text{resolve_blocked}(M_M.x, M'_{M'}, \text{resolvespec})$.

Say \mathbf{n} is *internally blocked* in P if there is an $\mathbf{n} : TCC'_{eqs_2}{}^\emptyset.e$ in P where e is an atomic internal blocked form in P .

Say \mathbf{n} is *blocked* in P if there is a $\mathbf{n} : TCC'_{eqs_2}{}^\emptyset.e$ in P where e is an atomic blocked form in P .

Say \mathbf{n} is in a *fast call* in P if there is a $\mathbf{n} : TCC_{eqs_2}^{\prime\emptyset}.e$ in P where e is \mathbf{RET}_T or $\mathbf{resolve}(M_M.x, M'_{M'}, \mathit{resolvespec})$.

$P|\mathbf{n} : TCC_{eqs}.\mathbf{op}(\mathbf{thinkify})^1 \mathit{tks} \rightarrow P'|\mathbf{n} : TCC_{eqs}.e$ if $\mathbf{Thinkify} \mathit{tks} P = (e, P')$

$P|\mathbf{n} : TCC_{eqs}.\mathbf{op}(\mathbf{thinkify})^1 \mathit{tks} \rightarrow P|\mathbf{n} : TCC_{eqs}.\mathbf{raise} e$ if $\mathbf{Thinkify} \mathit{tks} P = \mathbf{FAIL}(e)$

$P|\mathbf{n} : TCC_{eqs}.\mathbf{op}(\mathbf{unthinkify})^1 \mathit{thks} \rightarrow P|\mathbf{n} : TCC_{eqs}().|P'$ if $\mathbf{Unthinkify} \mathit{thks} (\{\mathbf{n}\} \cup \mathit{dom}(P)) = P'$

$P|\mathbf{n} : TCC_{eqs}.\mathbf{op}(\mathbf{unthinkify})^1 \mathit{thks} \rightarrow P|\mathbf{n} : TCC_{eqs}.\mathbf{raise} e$ if $\mathbf{Unthinkify} \mathit{thks} (\{\mathbf{n}\} \cup \mathit{dom}(P)) = \mathbf{FAIL}(e)$

The auxiliaries are defined in ML-like pseudocode below. $\mathbf{Thinkify}$ takes a list tks of thunkkeys and the process state. It gives either \mathbf{BLOCK} , if this $\mathbf{thinkify}$ cannot execute now (there is no transition rule in this case, thus blocking progress until thunkification is possible), or $\mathbf{FAIL}(e)$, if it should raise an exception, or the abstract syntax of an Acute function that takes a list of names (of the right shape) and has a body that unthunkifies the thunked mutexes, cvars and threads with those names and the remaining (non-thunkified) objects. It uses $\mathbf{DoThinkify}$, which is defined recursively on tks , P_1 and i , building up the pattern and body of Acute function as it goes.

$\mathbf{Unthinkify}$ calculates an Acute process to be added to the running program, or returns $\mathbf{FAIL}(e)$ if an exception should be raised.

There is no syntactic distinction between the pseudocode and object-language constructors; hopefully the context makes it clear.

```

Thinkify  $\mathit{tks} P =$ 
  match  $\mathbf{DoThinkify} \mathit{tks} P P 0$  with
     $\mathbf{BLOCK} \rightarrow \mathbf{BLOCK}$ 
  |  $\mathbf{FAIL}(e) \rightarrow \mathbf{FAIL}(e)$ 
  |  $(p, e, P_2) \rightarrow$ 
    (function  $x \rightarrow$ 
      match  $x$  with
         $p \rightarrow \mathbf{unthinkify} e$ 
      |  $\_ \rightarrow \mathbf{raise} \mathbf{THUNKIFY\_KEYLISTS\_MISMATCH}$ ),  $P_2$ 

DoThinkify  $\mathit{tks} P P_1 i =$ 
  match  $\mathit{tks}$  with
     $[] \rightarrow [], (), P_1$ 
  |  $tk :: \mathit{tks}_0 \rightarrow$ 
    match  $tk$  with
       $\mathbf{MUTEX} \mathbf{n}_0 \rightarrow$ 
        if  $\exists \mathbf{n}, P_0, \underline{b}. P_1 \equiv \mathbf{n} : \mathbf{MX}(\underline{b}) | P_0 \wedge \mathbf{n}_0 =_{\text{erased}} \mathbf{n}$  then
          let  $p, e, P_2 = \mathbf{DoThinkify} \mathit{tk}_0 P P_0 (i + 1)$  in
            ( $\mathbf{MUTEX}(x_i : \mathit{mutex} \ \mathit{name}) :: p, \mathbf{THUNKED\_MUTEX}(x_i, \underline{b}) :: e, P_2$ )
        else
           $\mathbf{FAIL}(\mathbf{NONEXISTENT\_MUTEX})$ 
      |  $\mathbf{CVAR} \mathbf{n}_0 \rightarrow$ 
        if  $\exists \mathbf{n}, P_0, v^\emptyset. P_1 \equiv \mathbf{n} : \mathbf{CV} | P_0 \wedge \mathbf{n}_0 =_{\text{erased}} \mathbf{n}$  then
          let  $p, e, P_2 = \mathbf{DoThinkify} \mathit{tk}_0 P P_0 (i + 1)$  in
            ( $\mathbf{CVAR}(x_i : \mathit{cvar} \ \mathit{name}) :: p, \mathbf{THUNKED\_CVAR}(x_i) :: e, P_2$ )
        else
           $\mathbf{FAIL}(\mathbf{NONEXISTENT\_CVAR})$ 
      |  $\mathbf{THREAD}(\mathbf{n}_0, \mathit{tmode}) \rightarrow$ 

```

```

if  $\exists \mathbf{n}, P_0, e_0. P_1 \equiv \mathbf{n} : e_0 | P_0 \wedge \mathbf{n}_0 =_{\text{erased}} \mathbf{n}$  then
  let  $p, e, P_2 = \text{DoThunkify } tk_0 P P_0 (i + 1)$  in
  if  $\mathbf{n}$  not blocked in  $P$  and not in a fast call in  $P$  then
    (THREAD( $x_i$  : thread name, thunkifymode)) ::  $p$ ,
    THUNKED_THREAD( $x_i$ , function ()  $\rightarrow e_0$ ) ::  $e, P_2$ 
  else if  $e_0 = CC'_{eqs_2}{}^\emptyset.e_1$  and  $e_1$  is an atomic blocked form in  $P$ 
    and  $tmode = \text{INTERRUPTING}$  then
      (THREAD( $x_i$  : thread name, thunkifymode)) ::  $p$ ,
      THUNKED_THREAD( $x_i$ , function ()  $\rightarrow CC'_{eqs_2}{}^\emptyset.\text{raise EINTR}$ ) ::  $e, P_2$ 
  else
    BLOCK
  else if  $\exists P_0, \text{definition}, \text{definitions}, e_0. P_1 \equiv \mathbf{n} : \text{definition definitions } e_0 | P_0$  then
    FAIL(THUNKIFY_THREAD_IN_DEFINITION)
  else
    FAIL(NONEXISTENT_THREAD)

```

where the x_i are all fresh.

Unthunkify $thks\ ns = \text{match } thks \text{ with}$

```

  []  $\rightarrow 0$ 
|  $thk :: thks_0 \rightarrow$ 
  (Unthunkify  $thks_0\ ns$ )
  |
  (match  $thk$  with
    THUNKED_MUTEX( $\mathbf{n}, \underline{b}$ )  $\rightarrow$  if  $\mathbf{n} \notin_{\text{erased}} ns$  then  $\mathbf{n} : \text{MX}(\underline{b})$  else FAIL(EXISTENT_NAME)
  | THUNKED_CVAR( $\mathbf{n}$ )  $\rightarrow$  if  $\mathbf{n} \notin_{\text{erased}} ns$  then  $\mathbf{n} : \text{CV}$  else FAIL(EXISTENT_NAME)
  | THUNKED_THREAD( $\mathbf{n}, v^\emptyset$ )  $\rightarrow$  if  $\mathbf{n} \notin_{\text{erased}} ns$  then  $\mathbf{n} : (v^\emptyset ())$  else FAIL(EXISTENT_NAME)
  )

```

Comment: There is a stylistic choice as to how a thunkified value is expressed. In principle it might just be a normal function in the language so far, but this requires non-trivial coding to ensure atomicity, e.g. to ensure one thread does not start before all the other threads are spawned and mutexes recreated. We therefore have a single semantic step, using some non-source-internal-language constructors to code the thunked state.

Comment: Maybe one would want a thunkifymode to apply also to mutexes and condition variables, eg to block until they reach a certain state.

Comment: If a thread tries to thunkify itself the NONEXISTENT_THREAD exception is raised.

Comment: If you feed the wrong things to a think you just get a match failure, not an unthunkify failure, which is slightly unpleasant.

Mutexes: create_mutex, lock, try_lock, unlock.

$$\begin{aligned}
 P|n : TCC_{eqs}.op(\mathbf{create_mutex})^1 n' &\rightarrow P|n : TCC_{eqs}().|n' : \mathbf{MX}(\mathbf{false}) \quad n' \notin_{\text{erased}} \text{dom}(P), \{n\} \\
 P|n' : \mathbf{MX}(\mathbf{false})|n : TCC_{eqs}.op(\mathbf{lock})^1 n'_1 &\rightarrow P|n' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}(). \quad n'_1 =_{\text{erased}} n' \\
 P|n' : \mathbf{MX}(\underline{b})|n : TCC_{eqs}.op(\mathbf{try_lock})^1 n'_1 &\rightarrow P|n' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}(\neg \underline{b}) \quad n'_1 =_{\text{erased}} n' \\
 P|n' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}.op(\mathbf{unlock})^1 n'_1 |n'' : TC'_{eqs'}.op(\mathbf{lock})^1 n'_2 &\rightarrow P|n' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}().|n'' : TC'_{eqs'}.() \quad n'_1 =_{\text{erased}} n' \wedge n'_2 =_{\text{erased}} n'' \\
 P|n' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}.op(\mathbf{unlock})^1 n'_1 &\rightarrow P|n' : \mathbf{MX}(\mathbf{false})|n : TCC_{eqs}(). \quad (*) \quad n'_1 =_{\text{erased}} n' \\
 P|n' : \mathbf{MX}(\mathbf{false})|n : TCC_{eqs}.op(\mathbf{unlock})^1 n'_1 &\rightarrow P|n' : \mathbf{MX}(\mathbf{false})|n : TCC_{eqs}(). \quad n'_1 =_{\text{erased}} n' \\
 P|n : TCC_{eqs}.op(\mathbf{create_mutex})^1 n' &\rightarrow P|n : TCC_{eqs}.raise \text{EXISTENT_NAME} \quad n' \in_{\text{erased}} \text{dom}(P), \{n\} \\
 P|n : TCC_{eqs}.op(\mathbf{lock})^1 n' &\rightarrow P|n : TCC_{eqs}.raise \text{NONEXISTENT_MUTEX} \quad n' \notin_{\text{erased}} \text{dom}(P), \{n\} \\
 P|n : TCC_{eqs}.op(\mathbf{unlock})^1 n' &\rightarrow P|n : TCC_{eqs}.raise \text{NONEXISTENT_MUTEX} \quad n' \notin_{\text{erased}} \text{dom}(P), \{n\} \\
 P|n : TCC_{eqs}.op(\mathbf{try_lock})^1 n' &\rightarrow P|n : TCC_{eqs}.raise \text{NONEXISTENT_MUTEX} \quad n' \notin_{\text{erased}} \text{dom}(P), \{n\}
 \end{aligned}$$

$$(*) \neg \exists (n'' : TC'_{eqs'}.op(\mathbf{lock})^1 n'_2) \in P . n'_2 =_{\text{erased}} n'$$

Comment: These rules give an error if one **lock**, **unlock**, or **try_lock** for a nonexistent mutex — which situation couldn't arise in the single-machine case, but can in ours. The simplest thing to do seems to be to have mutex and condition variable names global (which seems perfectly sensible, really), and to raise exceptions if one tries to use a nonexistent one.

Comment: Do we want to distinguish in the semantics between a **lock** m that has actually blocked and one that has not yet attempted to execute (introducing explicit “slow” states for them)? Can not see any need.

Condition variables: create_cvar, wait, signal, broadcast.

$$\begin{aligned}
 P|n : TCC_{eqs}.op(\mathbf{create_cvar})^1 n' &\rightarrow P|n : TCC_{eqs}().|n' : \mathbf{CV} \\
 P|n' : \mathbf{CV}|n'' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}.op(\mathbf{wait})^2 n'_1 n''_1 |n''' : TC'_{eqs'}.op(\mathbf{lock})^1 n''_2 &\rightarrow P|n' : \mathbf{CV}|n'' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}.op(\mathbf{waiting})^2 n' n'' |n''' : TC'_{eqs'}.() \\
 &\quad n'_1 =_{\text{erased}} n' \wedge n''_1 =_{\text{erased}} n'' \wedge n''_2 =_{\text{erased}} n''' \\
 P|n' : \mathbf{CV}|n'' : \mathbf{MX}(\mathbf{true})|n : TCC_{eqs}.op(\mathbf{wait})^2 n'_1 n''_1 &\rightarrow P|n' : \mathbf{CV}|n'' : \mathbf{MX}(\mathbf{false})|n : TCC_{eqs}.op(\mathbf{waiting})^2 n'_1 n''_1 \\
 &\quad (n'' : TC'_{eqs'}.op(\mathbf{lock})^1 n''_2) \notin P \\
 &\quad n'_1 =_{\text{erased}} n' \wedge n''_1 =_{\text{erased}} n'' \wedge n''_2 =_{\text{erased}} n'' \\
 P|n' : \mathbf{CV}|n : TCC_{eqs}.op(\mathbf{signal})^1 n'_1 &\rightarrow \text{restart_one}(P, n')|n' : \mathbf{CV}|n : TCC_{eqs}(). \quad n'_1 =_{\text{erased}} n' \\
 P|n' : \mathbf{CV}|n : TCC_{eqs}.op(\mathbf{broadcast})^1 n'_1 &\rightarrow \text{restart_all}(P, n')|n : TCC_{eqs}(). \quad n'_1 =_{\text{erased}} n'
 \end{aligned}$$

$$\begin{aligned}
P|n : TCC_{eqs}.\mathbf{op}(\mathbf{create_cvar})^1 n' &\rightarrow P|n : TCC_{eqs}.\mathbf{raise} \text{ EXISTENT_NAME} \quad n' \in_{\text{erased}} \text{dom}(P), \{n\} \\
P|n : TCC_{eqs}.\mathbf{op}(\mathbf{wait})^2 n' n'' &\rightarrow P|n : TCC_{eqs}.\mathbf{raise} \text{ NONEXISTENT_MUTEX} \quad \forall n'_1 =_{\text{erased}} n'' . n'_1 : \text{MX}(\underline{b}) \notin P \\
P|n : TCC_{eqs}.\mathbf{op}(\mathbf{wait})^2 n' n'' &\rightarrow P|n : TCC_{eqs}.\mathbf{raise} \text{ NONEXISTENT_CVAR} \quad \forall n'_1 =_{\text{erased}} n' . n'_1 : \text{CV} \notin P \\
P|n : TCC_{eqs}.\mathbf{op}(\mathbf{signal})^1 n' &\rightarrow P|n : TCC_{eqs}.\mathbf{raise} \text{ NONEXISTENT_CVAR} \quad \forall n'_1 =_{\text{erased}} n' . n'_1 : \text{CV} \notin P \\
P|n : TCC_{eqs}.\mathbf{op}(\mathbf{broadcast})^1 n' &\rightarrow P|n : TCC_{eqs}.\mathbf{raise} \text{ NONEXISTENT_CVAR} \quad \forall n'_1 =_{\text{erased}} n' . n'_1 : \text{CV} \notin P \\
P|n' : \text{CV}|n'' : \text{MX}(\mathbf{false})|n : TCC_{eqs}.\mathbf{op}(\mathbf{wait})^2 n'_1 n'' &\rightarrow P|n' : \text{CV}|n'' : \text{MX}(\mathbf{false})|n : TCC_{eqs}.\mathbf{raise} \text{ MUTEX_EPERM} \quad n'_1 =_{\text{erased}} n'
\end{aligned}$$

The auxiliaries are as follows:

`restart_one(P, n')` gives P but with a single $n : TCC_{eqs_2}^0.\mathbf{op}(\mathbf{waiting})^2 n'_1 n''_1$, if one exists with $n' =_{\text{erased}} n'_1$, replaced by $n : TCC_{eqs_2}^0.\mathbf{op}(\mathbf{lock})^1 n''$. If none exist, then `restart_one(P, n') = P`.

`restart_all(P, n')` gives P but with all $n : TCC_{eqs_2}^0.\mathbf{op}(\mathbf{waiting})^2 n'_1 n''_1$ for $n' =_{\text{erased}} n'_1$ replaced by $n : TCC_{eqs_2}^0.\mathbf{op}(\mathbf{lock})^1 n''$.

Comment: Here `op(wait)2 n' n''` for nonexistent n' and n'' nondeterministically gives one or the other error.

Comment: POSIX specifies that waiting without holding the mutex passed is an error. LinuxThreads appears (from the man page) not to implement this check; replace `MX(true)` with `MX(b)` above and remove the `MX(false)` rule above to mimic this. (If you do this, it is almost certain that your code has a race, so it is nice for the OS to let you know).

Comment: The “mutex handover” rule that atomically performs an unlock together with a lock is necessary for a sane and fair implementation.

Comment: Should a restarted `wait` atomically lock its mutex or not?

Comment: Applications may rely on some fairness property that this semantics does not express. Specifically, the threads waiting on a mutex or condition variable should be woken in FIFO order (i.e., for mutexes and `signal`, the first waiting thread should be woken; for `broadcast`, threads should be woken in such a way that the first waiting thread is first to be scheduled. Scheduling does not have to be this strict, but something like the fairness this implies is assumed by application programmers.

16.8.9 Polymorphism

We have not yet addressed the abstraction-preserving semantics for polymorphism, which will entail adding coloured brackets to the reduction axioms below and adding bracket-pushing rules for these constructs.

Without runtime type names or coloured brackets, the reduction axioms would be as below.

$$\begin{aligned}
(\Lambda t \rightarrow e) T &\rightarrow_{eqs} \{T/t\}e \\
\mathbf{let} \{t, x\} = (\{T, e\} \mathbf{as} T') \mathbf{in} e_2 &\rightarrow_{eqs} \{T/t, e/x\}e_2 \\
\mathbf{namecase} (\{T, (n', e)\} \mathbf{as} T') \mathbf{with} \{t, (x_1, x_2)\} \mathbf{when} x_1 = e \rightarrow e_2 \mathbf{otherwise} &\rightarrow e_3 \\
&\rightarrow_{eqs} \{T/t, e/x\}e_2 \quad \text{if } \text{erase_brackets}(n) = \text{erase_brackets}(n') \\
&\rightarrow_{eqs} e_3 \quad \text{if } \text{erase_brackets}(n) \neq \text{erase_brackets}(n')
\end{aligned}$$

Adding runtime type names, but still without coloured brackets, the unpack rule should be more like the rule below, which generates a fresh type name at each unpack to mirror the static semantics.

$$E_n ; \langle E_s, s, definitions, \mathbf{let} \{t, x\} = (\{T, e\} \mathbf{as} T') \mathbf{in} e_2 \rangle \rightarrow_{eqs} E_{n, n} : EQ(T) ; \langle E_s, s, definitions, \{n/t, e/x\} e_2 \rangle$$

16.9 Type Preservation and Progress

We have not attempted to prove type preservation and progress results, as the definition is of a size where either a hand or machine proof would be a very major undertaking (and a hand proof would probably contain many errors). Indeed, there may well be problems in the definition. However, it is still worth while stating precisely the properties that we believe should hold.

Some confidence in the soundness of the definition comes from running the implementation with runtime typechecking, typechecking the configuration after every reduction step.

The statements of the conjectures are the basis for this runtime typechecking, and are also a useful guide to the intuition while developing the definitions.

Conjecture 16.1 (Typed Compilation)

1. If $\text{compile}_{\Phi}(\text{sourcefilename})E_n \rightsquigarrow (E'_0, E'_1, \text{compiledunit}')$ and $\text{compiledunit}' = E_n'$; $\text{definitions}' \text{ eo}'$ then for some T and for $n \notin \text{dom}(E_n')$ we have $\vdash E_n', n_{\text{thread}}; \langle \text{empty}, \text{empty}, \text{definitions}, n_{\text{thread}} : \text{compiledunit}' \rangle : T$.

Conjecture 16.2 (Type Preservation) If

1. $\vdash E_n; \langle E_s, s, \text{definitions}, P \rangle : T$ and
2. $E_n; \langle E_s, s, \text{definitions}, P \rangle \xrightarrow{\ell} E_n'; \langle E'_s, s', \text{definitions}', P' \rangle$

then $\vdash E_n'; \langle E'_s, s', \text{definitions}', P' \rangle : T$.

Conjecture 16.3 (Progress) If $\vdash E_n; \langle E_s, s, \text{definitions}, P \rangle : T$ and there exists a thread $n : \text{definitions } e$ in P with n neither blocked in P nor in a fast call in P then there exists $E_n'; \langle \text{definitions}', E'_s, s', P' \rangle$ such that $E_n; \langle E_s, s, \text{definitions}, P \rangle \xrightarrow{\ell} E_n'; \langle E'_s, s', \text{definitions}', P' \rangle$.

Comment: One could also formulate a result saying the compilation always succeeds for well-typed source programs that do not include files or involve linking or `with !` etc. It would not be very informative, though.

Comment: The progress property should in principle be strengthened to ensure that in a well-typed configuration every non-blocked thread can make progress. To do so would require more data in the semantics, however, e.g. to track the threads involved in multi-thread reductions, so it is not worth doing now.

16.10 Runtime type checking

The main check is, for each configuration reached by the evaluator, that

$$\vdash E_n; \langle E_s, s, \text{definitions}, P \rangle : \text{unit}$$

(or indicate FAILURE.RUN.TYPECHECK_OF_CONFIGURATION).

We can also run typechecks during compilation on compiled units (or indicate FAILURE.COMPILE.TYPECHECK_OF_COMPILEDUNIT), at unmarshal time (or indicate FAILURE.RUN.TYPECHECK_ON_UNMARSHAL), and when compiled definitions are taken from a URI during module field instantiation (or indicate FAILURE.RUN.TYPECHECK_ON_GET_URI). These are described in §16.7, §16.8.5, and §16.8.5 respectively.

Failure of any of these checks indicates an error in the typesystem or the implementation.

The implementation has switches to control whether these checks are done. They all require structured names to be enabled.

16.11 Vacuous bracket optimization

The semantics above constructs coloured brackets in many circumstances where they are not required — where they do not change the colour. A production implementation would erase all brackets. For our implementation, while we need to keep the colour-changing brackets in order to do runtime typechecking, for execution speed it is useful to optimize away as many vacuous brackets as possible.

Accordingly, we define here an optimized variant semantics, which can, optionally, be used in our implementation.

1. in the rule **Module field instantiation – module case, via import sequence** (page 136), omit the brackets on the rhs if $eqs' = eqs$.
2. in the rule for $[\mathbf{raise} \ v^{eqs'}]_{eqs'}^T$ (page 129), omit the brackets on the rhs if $eqs' = eqs$.
3. in the rule for **marshal** MK $v^{eqs} : T$ (page 129), omit the brackets on the rhs if $eqs = \emptyset$.
4. in each of the 5 rules for pushing brackets through non-nullary constructors (page 130), omit the brackets on the rhs if $eqs' = eqs$.
5. in the rule for pushing brackets through lambda (page 130), omit the outer brackets on the rhs if $eqs' = eqs$.
6. in the rule for bracket type revelation (page 130), omit the brackets on the rhs if $eqs' = eqs$.
7. in the rule for bracket elimination (page 130), omit the brackets on the rhs if $eqs'' = eqs$.
8. in the two rules for op^n and x^n (page 131), omit the brackets on the rhs if $eqs = \emptyset$.

Note that brackets constructed to be used in a substitution (in the definition of `matchsub`, the **function** rule, and the **let rec** rule) cannot be optimized away without analysis of the structure of the expression in which they are being substituted. We do not do this.

The Type Preservation property should still hold.

16.12 Closures

16.12.1 Value closures

For efficiency, the implementation uses an environment instead of substitution. This requires function values to be represented as closures. In this section we extend the small-step semantics given above to model this. We remain in the style of a calculus, rather than an abstract machine.

The reduction arrow now gains another index, the environment ρ , in addition to the colour eqs it already carries:

$$e \rightarrow_{eqs}^{\rho} e'$$

We must also introduce a term form to change the environment – recall that a closure replaces the environment, rather than extending the existing one. This same form can be used also to implement other forms of binding that only extend the environment. We write **inenv** ρ' **do** e to denote that e is evaluated in the environment ρ' . (Recall that colour changes are already handled by brackets.) **inenv** itself is an environment-changing evaluation context.

Evaluation contexts now carry an inner and an outer environment as well as an inner and outer colour; **inenv** ρ' **do** e is an environment-changing context. This has the same effect as the following rule:

$$\frac{e \rightarrow_{eqs}^{\rho'} e'}{\mathbf{inenv} \ \rho' \ \mathbf{do} \ e \rightarrow_{eqs}^{\rho} \mathbf{inenv} \ \rho' \ \mathbf{do} \ e'}$$

Let ρ be an environment, i.e., a list of pairs $\{[v^{eqs}]_{eqs}^T/x\}$, such that earlier pairs scope over later ones. Observe that ρ is both closed and colour-closed: the environment has no free identifiers, and since every value is enclosed in brackets, it is valid at any colour.

Comment: This definition is not entirely correct as stated — below we allow a recursive closure to contain an environment that includes a pair whose second element is the original closure. This is expressible in our implementation language (FreshOCaml, following OCaml, allows recursive value bindings of the form required) but is not well-formed in the naive set-theoretic model of the semantics below. The definition should be adapted.

The new expression and value forms are as follows:

$$\begin{aligned} e & ::= \dots \\ & \quad \mathbf{inenv} \ \rho \ \mathbf{do} \ e \\ & \quad \mathbf{Clos}(\rho, x_2 : T_2, BC_2, e_1, \mathbf{NONE}) \\ & \quad \mathbf{Clos}(\rho, x_2 : T_2, BC_2, e_1, \mathbf{SOME}(x_1 : T_1, BC_1)) \\ v & ::= \dots \text{ except for } \mathbf{function} \ x \rightarrow e \\ & \quad \mathbf{Clos}(\rho, x_2 : T_2, BC_2, e_1, \mathbf{NONE}) \\ & \quad \mathbf{Clos}(\rho, x_2 : T_2, BC_2, e_1, \mathbf{SOME}(x_1 : T_1, BC_1)) \end{aligned}$$

Note that these new expression and value forms appear only in running programs; that is, **function** $x \rightarrow e$ remains a source value (and is allowed to appear as a value in a struct, for example), but is no longer a value in a running program (the corresponding closure is, instead).

Identifier lookup uses the environment (this is the delayed substitution in action). We may incorporate the vacuous bracket optimisation, since every binding in the environment has an outermost bracket:

$$\begin{array}{ll} x \rightarrow_{eqs}^{\rho} v^{eqs} & \text{if } \rho(x) = [v^{eqs}]_{eqs}^T \quad \text{identifier lookup, bracket eliminated} \\ x \rightarrow_{eqs}^{\rho} [v^{eqs'}]_{eqs'}^T & \text{if } \rho(x) = [v^{eqs'}]_{eqs'}^T \text{ and } eqs' \neq eqs \quad \text{identifier lookup, bracket required} \end{array}$$

The normal binding constructs (**match**, **try**) use **inenv** in the obvious way:

match v^{eqs} with $p_1 \rightarrow e_1 \dots p_n \rightarrow e_n$	\rightarrow_{eqs}^ρ	inenv $(\rho + \text{matchsub}_{eqs}(v^{eqs}, p_k))$ do e_k	match success (a)
match v^{eqs} with $p_1 \rightarrow e_1 \dots p_n \rightarrow e_n$	\rightarrow_{eqs}^ρ	raise MATCH_FAILURE v'	match failure (b)
try	similarly

everything else is straightforward

Elimination of **inenv** occurs when evaluation within it is complete:

inenv ρ' **do** v^{eqs} \rightarrow_{eqs}^ρ v^{eqs} scope exit

Recall the existing rules for functions:

[function $(x_2 : T_2) \rightarrow e]$	$\xrightarrow{eqs}^{T_2 \rightarrow T_3}$	function $(x_2 : T_2') \rightarrow \{[x_2]_{eqs'}^{T_2} / x_2\} e]$	bracket pushing through function
(function $(x_2 : T_2) \rightarrow e)$	\rightarrow_{eqs}	$\{[v^{eqs}]_{eqs'}^{T_2} / x_2\} e$	function application
let rec $x_1 : T = \text{function } (x_2 : T') \rightarrow e_1$ in e_2	\rightarrow_{eqs}	$\{[\text{let rec } x_1 : T = \text{function } (x_2 : T') \rightarrow e_1 \text{ in } x_1]_{eqs'}^T / x_1\} \text{function } (x_2 : T') \rightarrow e_1$	recursive function application

It may seem that a closure should carry its colour as well as its environment. In fact, however, it shouldn't – just as a function doesn't. Colour change is effected by binding, and we take care in substitution (directly or with environments) to insert sufficient brackets to get this right. Therefore we merely have to get bracket pushing correct for closures.

Thus, a closure carries the function argument, function body, and the environment it was defined in. It also carries a bracket context BC , discussed below, and for recursive closures it carries the name and type of the recursive binder, also discussed below. On application, we reintroduce the environment, and bind the argument. On pushing a bracket through a closure, the environment is untouched (it is colour-closed); the bracket is accumulated in the bracket context.

The rules for closures are as follows (notice that the typing rules imply $T_1 \approx T_2 \rightarrow T_3$).

function $(x_2 : T_2) \rightarrow e$	\rightarrow_{eqs}^ρ	Clos $(\rho, x_2 : T_2, -, e, \text{NONE})$	closure formation
let rec $x_1 : T_1 = \text{function } (x_2 : T_2) \rightarrow e_1$ in e_2	\rightarrow_{eqs}^ρ	inenv ρ' do e_2	recursive closure formation
where $\rho' = \rho + \{[\text{Clos}(\rho', x_2 : T_2, -, e_1, \text{SOME}(x_1 : T_1, -))]_{eqs'}^{T_1} / x_1\}$			
Clos $(\rho', x_2 : T_2, BC_2, e, x_0)$	\rightarrow_{eqs}	inenv $(\rho' + \{BC_2.[v^{eqs}]_{eqs'}^{T_2} / x_2\})$ do e	closure application
$[\text{Clos}(\rho', x_2 : T_2, BC_2, e, \text{NONE})]_{eqs'}^{T_2 \rightarrow T_3}$	\rightarrow_{eqs}	Clos $(\rho', x_2 : T_2', BC_2.[e]_{eqs'}^{T_2}, [e]_{eqs'}^{T_3}, \text{NONE})$	bracket pushing through closure
$[\text{Clos}(\rho', x_2 : T_2, BC_2, e, \text{SOME}(x_1 : T_1, BC_1))]_{eqs'}^{T_2 \rightarrow T_3}$	\rightarrow_{eqs}	Clos $(\rho', x_2 : T_2', BC_2.[e]_{eqs'}^{T_2}, \text{SOME}(x_1 : T_2' \rightarrow T_3, BC_1.[e]_{eqs'}^{T_1}))$	bracket pushing through recursive closure

Notice that **let rec** is similar to **function**, and recursive and non-recursive closures share the same application rule. For flattening purposes, however, we must store the name and type $x_1 : T_1$ of the recursive binder so that we are able to reconstruct the appropriate **let rec**; otherwise naïve application of the ρ would fail to terminate.

In the original bracket pushing through function rule, we perform a substitution on the bound variable(s). We would like to delay this substitution as well. In order to do this, we simply accumulate a sequence of pushed brackets within the closure, adding them to the environment only at application time. This means that bindings in the environment may now be to values surrounded by arbitrary bracket contexts, rather than values only; administrative reductions may be required in order to reduce these to a value. We do not consider this an important difficulty.

We may perform the vacuous bracket optimisation when appending to bracket contexts in closure application and bracket pushing through closure, as follows:

$\text{maybe_cons_bs}_{eqs_0} \left[_ \right]_{eqs}^T$	$=$	$_$	if $eqs_0 = eqs$
$\text{maybe_cons_bs}_{eqs_0} \left[_ \right]_{eqs}^T BC.[_]_{eqs_0}^{T'}$	$=$	$BC.[_]_{eqs_0}^{T'}$	if $eqs_0 = eqs$
$\text{maybe_cons_bs}_{eqs_0} \left[_ \right]_{eqs}^T BC$	$=$	$BC.[_]_{eqs}^T$	otherwise

Here we only append the bracket if it differs from the innermost colour of the existing bracket context (or the ambient colour eqs_0 if the bracket context is empty).

To correctly define flattening in the recursive case, we must extend the codomain of ρ to include bracket-context forms $BC.*$ (denoted by a literal $*$), in addition to the usual expressions. Flattening can now be defined as follows:

$$\begin{aligned}
 \text{flattenclos}_\rho(x) &= \text{flattenclos}_\rho \rho(x) && \text{where } \rho(x) \neq BC.* \\
 \text{flattenclos}_\rho(x) &= BC.x && \text{where } \rho(x) = BC.* \\
 \text{flattenclos}_\rho(\mathbf{inenv} \ \rho' \ \mathbf{do} \ e) &= \text{flattenclos}_{\rho'}(e) \\
 \text{flattenclos}_\rho(\mathbf{Clos}(\rho', x_2 : T_2, BC_2, e_1, \mathbf{NONE})) &= \mathbf{function} \ (x_2 : T_2) \rightarrow \text{flattenclos}_{(\rho' + \{BC_2.x_2/x_2\})}(e_1) \\
 \text{flattenclos}_\rho(\mathbf{Clos}(\rho', x_2 : T_2, BC_2, e_1, \mathbf{SOME}(x_1 : T_1, BC_1))) &= \mathbf{let} \ \mathbf{rec} \ x_1 : T_1 = \mathbf{function} \ (x_2 : T_2) \rightarrow \text{flattenclos}_{(\rho' + \{BC_1.* / x_1, BC_2.* / x_2\})}(e_1) \ \mathbf{in} \ x_1 \\
 &= \mathbf{let} \ \mathbf{rec} \ x_1 : T_1 = \mathbf{function} \ (x_2 : T_2) \rightarrow \text{flattenclos}_{(\rho' + \{BC_1.* / x_1, BC_2.* / x_2\})}(e_1) \ \mathbf{in} \ x_1 \\
 &\text{everything else is just recursive descent}
 \end{aligned}$$

The new codomain form is used where a identifier must be wrapped once, rather than recursively expanded.

The new constructs may be typed directly. We make use of an auxiliary function, envenv , defined as follows:

$$\begin{aligned}
 \text{envenv}(\emptyset) &= \emptyset \\
 \text{envenv}(\{[v^{eqs}]_{eqs}^T / x, \rho'\}) &= x : T, \text{envenv}(\rho')
 \end{aligned}$$

Then the type rules are as follows:

$$\begin{array}{c}
 \frac{E_n, E_0, \text{envenv}(\rho) \vdash_{eqs} e : T}{E_n, E_0, E \vdash_{eqs} \mathbf{inenv} \ \rho \ \mathbf{do} \ e : T} \\
 \\
 \frac{\begin{array}{l} E_n, E_0, x_2 : T_2 \vdash_{eqs} BC_2.x_2 : T'_2 \\ E_n, E_0, \text{envenv}(\rho), x_2 : T'_2 \vdash_{eqs} e_1 : T_3 \end{array}}{E_n, E_0, E \vdash_{eqs} \mathbf{Clos}(\rho, x_2 : T_2, BC_2, e_1, \mathbf{NONE}) : T_2 \rightarrow T_3} \\
 \\
 \frac{\begin{array}{l} E_n, E_0, x_1 : T_1 \vdash_{eqs} BC_1.x_1 : T'_1 \\ E_n, E_0, x_2 : T_2 \vdash_{eqs} BC_2.x_2 : T'_2 \\ E_n, E_0, \text{envenv}(\rho) \vdash_{eqs} x_1 : T'_1 \\ E_n, E_0, \text{envenv}(\rho), x_2 : T'_2 \vdash_{eqs} e_1 : T_3 \\ E_n, E_0 \vdash_{eqs} T_1 \approx T_2 \rightarrow T_3 \end{array}}{E_n, E_0, E \vdash_{eqs} \mathbf{Clos}(\rho, x_2 : T_2, BC_2, e_1, \mathbf{SOME}(x_1 : T_1, BC_1)) : T'_1}
 \end{array}$$

where E_0 is that prefix of the environment which arises from E_{const} and the enclosing *definitions*.

16.12.2 Type closures

A naïve implementation of polymorphism would perform a substitution for each instance of type application, negating much of the benefit of value closures. We therefore introduce type closures as well.

The environment ρ now contains pairs T/t as well as $[v^{eqs}]_{eqs}^T / x$. The reduction arrow, value closures, and the \mathbf{inenv} form all remain the same. We add a new form $\mathbf{TClos}(\rho, t, e)$ with the obvious meaning; a type abstraction is no longer a value, and instead reduces to the obvious type closure. \mathbf{inenv} is used everywhere instead of type substitution. flattenclos_ρ is extended in the obvious way. Brackets and type closures may be commuted freely.

Care must be taken to ensure that whenever a type is used, ρ is taken into account; if the type is taken out of its context (as in reduction of a $\mathbf{marshal}$ expression for example) it must be flattened first.

Part IV

Communication Infrastructure Example

Here we give the Acute code for the communication infrastructure example outlined in §11. It consists of modules `Tcp_padded`, `Tcp_connection_management`, `Tcp_string_messaging`, `Local_channel`, `Distributed_channel`, `Npi1`, `Npi2`, and `Npi`, followed by two simple clients of the `Npi` library, `npi-recv` and `npi-mig`.

```
(* tcp.ac *)
(* This file contains Tcp_padded, Tcp_connection_management and Tcp_string_messaging modules *)

(* These use the Sockets API and local concurrency - threads and mutexes. *)
(* Both are hash modules, providing abstract types of handles. *)

includesource "util.ac"

(* ***** *)
(* ** ** *)
(* ** Tcp_padded ** *)
(* ** ** *)
(* ***** *)

(* The Tcp_padded module implements a wire-format send and receive for
arbitrary strings.

The wire format encoding of a string consists of 21 bytes
containing an ASCII pretty-print of its length followed by the
string itself. This is not efficient(!) but is conveniently
human-readable.
*)

module hash Tcp_padded :
sig
  val send : Tcp.fd -> ((Tcp.ip * Tcp.port) option) -> string -> unit
  val recv : Tcp.fd -> string
end =
struct
  let send fd ippo data =
    let pad data n =
      let padding =
        String.make ( n - (String.length data) ) ' ' in
      (data ^ padding) in
    let data_length = String.length data in
    let data_length_string =
      pad (Pervasives.string_of_int data_length) 21 in
    let rec send_all s =
      let no_options = [] in
    let s' = (Tcp.send fd ippo s no_options) in
    if 0 = (String.length s') then () else send_all s'
    in
    send_all (data_length_string ^ data)

  let recv fd =
    let rec recv_n_bytes = function n ->
      let no_options = [] in
```

```

    let (s,_) = Tcp.recv fd n no_options in
    let _ = IO.print_string ("Tcp_padded.recv got " ^ Pervasives.string_of_int (String.length s)
        ^" bytes; expecting " ^ Pervasives.string_of_int (n - String.length s) ^ " more
n") in
    (* let _ = IO.print_string ("in particular, Tcp_padded.recv got ---" ^ s ^ "---
n") in*)
    let l = String.length s in
    if l = 0 then (Tcp.close fd; raise (Failure "socket closed by the other party") )
    else if l >= n then s else s ^ (recv_n_bytes (n-l)) in
    let data_length_string = recv_n_bytes 21 in
    let first_space = String.index data_length_string ' ' in
    let data_length_string' = String.sub data_length_string 0 first_space in
    let data_length = Pervasives.int_of_string data_length_string' in
    recv_n_bytes data_length

end

```

```

(* ***** *)
(* ** ** *)
(* ** Tcp_connection_management ** *)
(* ** ** *)
(* ***** *)

```

(* The Tcp_connection_management module manages collections of TCP connections.

daemon takes a local address (an Tcp.ip option * Tcp.port option) and an incoming-connection-handler function and creates a listening socket on that address, spawning a thread that invokes the supplied function for any incoming connection and then adds the connection to a list. daemon returns a handle which must be passed in to the other functions. (Using handles rather than module state allows a single runtime to have multiple instances with different local addresses.)

establish_to takes a handle and remote address. If there is already a connection to that address it returns its file descriptor, otherwise it tries to establish one (and returns the new file descriptor).

disestablish_to takes a handle and remote address, closing and removing a connection to that address if one exists.

connection_failed takes a handle and remote address (one for which a connection has failed) and removes it from the stored list.

shutdown closes and removes all connections and closes the listening socket.

local_addr takes a handle and returns the local address.

*)

```

(* TODO: Deal more sensibly with TCP errors and the REUSEADDR semantics, here and in the clients *)
(* TODO: Think about efficiency *)
(* TODO: Have shutdown cleanly terminate the associated thread *)

```

```

module hash Tcp_connection_management :
sig
  type fd = Tcp.fd
  type handle
  val daemon : Tcp.ip option * Tcp.port option ->
    ((Tcp.ip option * Tcp.port)->Tcp.addr->fd -> unit) -> handle
  val establish_to : handle -> Tcp.addr -> fd
  val disestablish_to : handle -> Tcp.addr -> unit
  val shutdown : handle -> unit
  val connection_failed : handle -> Tcp.addr -> unit
  val local_addr : handle -> Tcp.ip option * Tcp.port
end =
struct
  type fd = Tcp.fd
  type handle =
    (Tcp.ip option * Tcp.port) (* local address *)
    * fd (* listening socket *)
    * ((Tcp.ip option * Tcp.port)->Tcp.addr->fd->unit) (* incoming conn handler *)
    * mutex name (* current connections mutex *)
    * (Tcp.addr * fd) list ref (* current connections *)

  let daemon (ipo,po) f =
    let conn_mutex = fresh in
    create_mutex conn_mutex;
    Pervasives.print_endline ("Created TCP mutex " ^ name_to_string conn_mutex);
    let conn = ref [] in
    let fd = Tcp.tcp_socket () in
    let _ = Tcp.bind fd ipo po in
    let (ipo,p) = match Tcp.getsockname fd with
      (Some ip, Some p) -> (Some ip, p)
    | (None, Some p) -> (None,p)
    | _ -> raise (Failure "no local port after bind()") in
    let _ = let backlog = 5 in Tcp.listen fd backlog in
    (while true do
      let (fd',(ip',p')) = Tcp.accept fd in
      let p'' = (unmarshal (Tcp_padded.recv fd') as Tcp.port) in
      f (ipo,p) (ip',p'') fd' ; (* note that f terminates before adding this to conn *)
      Uutils.locked_by_stmt conn_mutex
        (function () ->
          conn := ((ip',p''),fd') :: !conn)
    done |||
    (((ipo,p),fd,f,conn_mutex,conn)
    ))

  let establish_to h (ip',p') =
    let ((ipo,p),fd_listen,f,conn_mutex,conn) = h in
    Uutils.locked_by_stmt2 %[fd] conn_mutex (function ()->
      try
        List.assoc %[Tcp.addr] %[[]] (ip',p') !conn
      with
        Not_found ->
          let fd = Tcp.tcp_socket () in
    Tcp.bind fd ipo None;
      Pervasives.print_endline ("Establish connecting to p' = " ^
        Pervasives.string_of_int(Tcp.int_of_port p') );
    Tcp.connect fd ip' Some p';

```

```

    let d = (marshal "StdLib" p : Tcp.port) in
    Pervasives.print_endline ("Establish p = " ^ Pervasives.string_of_int(Tcp.int_of_port p));
    (* Pervasives.print_endline ("Establish string = " ^ d ); *)
    Tcp_padded.send fd None d;
    f (ipo,p) (ip',p') fd;
    conn := ((ip',p'),fd) :: !conn;
fd
)

let disestablish_to h (ip',p') =
  let ((ipo,p),fd_listen,f,conn_mutex,conn) = h in
  Utils.locked_by_stmt conn_mutex (function ()->
    try
      let fd = List.assoc %[] %[] (ip',p') !conn in
      conn := List.remove_assoc %[] %[] (ip',p') !conn;
      Tcp.close fd
    with
      Not_found -> ()
  )

let shutdown h =
  let ((ipo,p),fd_listen,f,conn_mutex,conn) = h in
  Utils.locked_by_stmt conn_mutex (function ()->
    List.iter %[] (function ((ip,p),fd) -> Tcp.close fd) !conn;
    conn := [];
    Tcp.close fd_listen)

let connection_failed h (ip',p') =
  let ((ipo,p),fd_listen,f,conn_mutex,conn) = h in
  Utils.locked_by_stmt conn_mutex (function () ->
    conn := List.remove_assoc %[] %[] (ip',p') !conn )

let local_addr h =
  let ((ipo,p),fd_listen,f,conn_mutex,conn) = h in
  (ipo,p)

end

(* ***** *)
(* ** ** *)
(* ** Tcp_string_messaging ** *)
(* ** ** *)
(* ***** *)

(* The Tcp_string_messaging module provides asynchronous messaging of
strings to TCP addresses, using Tcp_connection_management.

daemon takes a local address (an Tcp.ip option * Tcp.port option) and
a function to handle incoming strings, of type

(Tcp.ip option * Tcp.port) -> Tcp.addr -> string -> unit

and creates a Tcp_connection_management.daemon, returning a handle.

send takes a handle, a remote TCP address and a string, uses

```


Tcp_connection_management.establish_to to ensure there is a connection, and sends the string (encapsulated in a wire format).

shutdown takes a handle and shuts down (calling Tcp_connection_management.shutdown).

local_addr takes a handle and returns the local TCP address.

The wire format is implemented by Tcp_padded.

*)

(* TODO: handle send/recv errors and call connection_failed as required *)
(* TODO: need more locking to stop different send/recvs interleaving *)
(* TODO: one might want to pass the handle as another argument to the function argument to daemon *)

```
module hash Tcp_string_messaging :
sig
  type handle
  val daemon : Tcp.ip option * Tcp.port option ->
    ((Tcp.ip option * Tcp.port) -> Tcp.addr -> string -> unit) -> handle
  val send : handle -> Tcp.addr -> string -> unit
  val shutdown : handle -> unit
  val local_addr : handle -> Tcp.ip option * Tcp.port
end =
struct
  type handle = Tcp_connection_management.handle

  let daemon (ipo,po) f =
    let g ipop addr' fd =
      create_thread fresh (function () ->
        while true do
          let data = Tcp_padded.recv fd in
            f ipop addr' data
        done
      ) ()
    in
      Tcp_connection_management.daemon (ipo,po) g

  let send h (ip,p) data =
    let fd = Tcp_connection_management.establish_to h (ip,p) in
      Tcp_padded.send fd (Some(ip,p)) data

  let shutdown h = Tcp_connection_management.shutdown h

  let local_addr h = Tcp_connection_management.local_addr h
end
```

(* ***** *)
(* ** ** *)

```
(* ** Local_channel                                ** *)
(* **                                              ** *)
(* ***** *)
```

(* Module Local_channel provides simple typed asynchronous local channels.

Function send : forall t. t name -> t -> unit sends a message on the specified name, returning immediately.

Function recv : forall t. t name -> (t -> unit) -> unit registers a receiver on the specified name, returning immediately.

As soon as there is both a message and a receiver for a name the receiver is applied to the message. The receiver is then removed.

The interface uses (T name) as the type of channels carrying values of type T. Exposing the fact that this is a name type allows clients to use any of the methods for constructing shared typed names that Acute provides.

One might instead think of using ML-style references as channel 'names'. For a local implementation that would be fine, but one one marshalled values mentioning channels the whole channel data structure would be copied, which is not our desired semantics.

Internally, the pending messages and receivers on the channels are stored in a list of existential packages, of type

```
(exists t. t name * (t list ref * (t->unit) list ref)) list
```

with the Acute namecase operation used in lookups.

This is a hash! module. There is module state: the handle h consists of a mutex name and a pointer to the channel data structure. (Here h is exposed, abstractly, in the interface, purely to work around the current lack of width subsignaturing.) Nonetheless, rebinding to local instances of Local_channel.send and Local_channel.recv should just work, so we use the hash! mode to give an exact-hash version (and, as part of that workaround, to make the abstract type of h hash-generated).

One might think of passing the handle explicitly as an argument to send and recv, doing without module state. That again would lead to the wrong semantics for marshalling values that use this library.

*)

```
(* NB: fields marked by (*A*) will be removed from the interface *)
(* when width subsignaturing is added *)
```

```
module hash! Local_channel :
sig
  type handle      (*A*)
  val h : handle   (*A*)
  val send : forall t. t name -> t -> unit
```

```

    val recv : forall t. t name -> (t -> unit) -> unit
end

=
struct

  type handle = mutex name * (exists t. t name * (t list ref * (t->unit) list ref)) list ref

  let h = (let n = fresh in create_mutex n; n, ref [] )

  (* A handle consists of a mutex and a reference to a list of
     channel structures. Each channel structure is an existential
     package containing a name, a reference to a list of pending
     messages and a reference to a list of pending receptors. We
     maintain the invariant that at most one of those two is
     nonempty.

     Channel structures are added to the list as necessary. At
     present they are never removed; we could remove them when they
     become empty.

     The pending messages and pending receptors are kept with the
     oldest at the heads of the lists.
  *)

  (* Note the use of namecase below *)

  let send = Function t -> fun (cn: t name) (v: t) ->
    let (m,csr) = h in
    Utils.locked_by_stmt m
    (function () ->
      let rec lookup cs' = match cs' with
        [] -> csr :=
          ( {t,(cn,(ref (v::[]),ref []))} as exists t. t name * (t list ref * (t->unit) list ref) )
          :: !csr
        | (c: exists t'. t' name * (t' list ref * (t'->unit) list ref))::cs0 ->
          namecase c with
            {t',(cn',xyz)} when cn'=cn ->
              let ((msgs: t list ref),rcvrs)=xyz in
              match !rcvrs with (* in this branch the typechecker needs to know t=t' *)
                [] -> msgs := (!msgs @ (v::[]))
                | rcvr::rcvrs0 -> (
rcvrs:=rcvrs0; (* could remove this whole channel if it's become empty*)
create_thread fresh rcvr v)
              otherwise ->
                lookup cs0
            in lookup !csr
    )

  let recv = Function t -> fun (cn: t name) (f: t -> unit) ->
    let (m,csr) = h in
    Utils.locked_by_stmt m
    (function () ->
      let rec lookup cs' = match cs' with
        [] -> csr := ({t, (cn,(ref [],ref (f::[])))} as
          exists t. t name * (t list ref * (t->unit) list ref)) :: !csr
        | (c: exists t'. t' name * (t' list ref * (t'->unit) list ref))::cs0 ->

```

```

    namecase c with
      {t,(cn',x)} when cn'=cn ->
        let ((msgs: t list ref),rcvrs)=x in
          match !msgs with
            [] -> rcvrs := !rcvrs @ (f::[])
          | v::vs0 -> (
              msgs:=vs0;
              create_thread %[t] fresh f v)
          otherwise ->
            lookup cs0
        in lookup !csr
    )

end

mark "LChan"

(* TODO: Extend with replicated input and with blocking receive. *)

includesource "tcp.ac"
includesource "local_channel.ac"

(* ***** *)
(* ** ** *)
(* ** Distributed_channel ** *)
(* ** ** *)
(* ***** *)

(* Distributed_channel provides simple typed asynchronous distributed
channels, above Tcp_string_messaging and Local_channel.

Function init : Tcp.ip option * Tcp.port option -> unit initialises
a Tcp_string_messaging daemon with the specified port and IP
address.

Function send : forall t. string -> (Tcp.addr * t name) -> t -> unit
sends a message (marshalled wrt the mark specified) to the
specified channel at the specified TCP address, returning
immediately. It does a case split depending on whether the target
is local or not, for efficiency.

Function rcv : forall t. t name -> (t -> unit) -> unit registers
a receiver on the specified name, returning immediately.

Function local_addr : unit -> Tcp.ip option * Tcp.port option
returns the registered local address.

As soon as there is both a message and a receiver for a name the
receiver is applied to the message. The receiver is then removed.
These are _non-mobile_ distributed channels: the receivers cannot
be moved from one Tcp.addr to another. See npi.ac for a mobile
extension.

```

Similarly to Local_channel, this is a hash! module. The module state consists of an ho field, recording the Tcp_string_messaging handle in use. To allow client code to determine the local TCP address this is set by the init function (it is stored as an option reference and can be set at most once). Use of the hash! mode gives the module an exact-hash version. Use of module state (rather than explicitly-passed handles) ensures the right semantics when marshalling client code.

Internally, the wire format consists of marshalled values of type

```
exists t'.t' name * t'
```

marshalled with respect to whatever mark is supplied to the send function. This mark should usually be at or below the mark "DChan" just below the module, so that the Distributed_channel code itself is not marshalled.

*)

(* NB: fields marked by (*A*) will be removed from the interface *)
 (* when width subsignaturing is added *)

```
module hash! Distributed_channel :
```

```
sig
```

```
  type tf                (*A*)
  type tho               (*A*)
  val f : tf             (*A*)
  val ho : tho           (*A*)
  val init : Tcp.ip option * Tcp.port option -> unit
  val send : forall t. string -> (Tcp.addr * t name) -> t -> unit
  val recv : forall t. t name -> (t -> unit) -> unit
  val local_addr : unit -> Tcp.ip option * Tcp.port
```

```
end
```

```
=
```

```
struct
```

```
  type tf = (Tcp.ip option * Tcp.port) -> Tcp.addr -> string -> unit
  type tho = Tcp_string_messaging.handle option ref
```

```
  let f ipop_local addr_remote data =
```

```
    let {t,x} = unmarshal data as exists t'. t' name * t' in
    let (c,v) = x in
    Pervasives.prerr_endline("Got v: " ^ (marshal "StdLib" (v) : t));
    Local_channel.send [%t] c v
```

```
  let ho = ref None
```

```
  let init (ipo,po) =
```

```
    match !ho with
    | Some _ -> raise (Failure "Distributed_channel already initialised")
    | None -> ho := Some (Tcp_string_messaging.daemon (ipo,po) f)
```

```
  let send = Function t -> fun mk -> fun (addr,(c: t name)) (v: t) ->
```

```
    let h = Utils.the [%[]] !ho in
    let (ip, port) = addr in
    if (Some ip, port) = Tcp_string_messaging.local_addr h then
      Local_channel.send [%t] c v
```

```

else
(Pervasives.prerr_endline("marshalling");
  let data = marshal mk ({t, (c,v)} as exists t'.t' name * t')
    : exists t'.t' name * t' in
(Pervasives.prerr_endline("sending " ^ data);
  Tcp_string_messaging.send h addr data))
let recv = Function t -> fun (c: t name) (f: t -> unit) ->
  Local_channel.recv %[t] c f

  let local_addr () = Tcp_string_messaging.local_addr (Utils.the %[] !ho)
end

```

mark "DChan"

(* TODO: Extend with replicated input and with blocking receive *)
 (* Note that with this code the local-send optimisation will only be
 effective if the local daemon IP was set explicitly, not
 wildcarded. To deal properly with hosts with multiple interfaces one
 should check against getifaddrs. *)

includesource "tcp.ac"

```

(* ***** *)
(* ** ** *)
(* ** Npi, consisting of Npi1 and Npi2 ** *)
(* ** ** *)
(* ***** *)

```

(* The Npi module manages groups of threads in a single acute process,
 implementing the key primitives of the Nomadic Pict language.

A thread can either be registered with the Npi module or not.
 If it is registered, it belongs to exactly one group throughout its
 execution.
 Local communication within a group and inter-group communication
 via typed channels is supported.
 Furthermore, there is a "migrate_group" command, which when called by
 one member of the group, migrates the whole group to a new Tcp address.
 For this to work, the other end also needs to have an initialised
 Npi module running.
 The correct operation of this module depends on the client code not
 using any low-level primitives - thread operations, thunkify, etc.

Most important functions:

```

init : (Tcp.ip option * Tcp.port option) -> unit
  initialise group infrastructure to handle inter-group communication
  and group migrations.

create_group : forall t. (t -> unit) -> t -> unit
  create a new group containing one (new) thread.

create_gthread : forall t. (t -> unit) -> t -> unit
  add a new thread to the current group.

```

```

recv_local : forall t. t name -> t
  receive information from a named typed channel

send_local : forall t. t name -> t -> unit
  send information to a named typed channel (of current group)

send_remote : forall t. string -> (Tcp.addr * group name * t name) -> t -> unit
  send information to a named typed channel of another group at a
  known Tcp address.

migrate_group : Tcp.addr -> unit
  migrate current group to a new Tcp address.

```

As is Local_channel and Distributed_channel, (T name)s are used for channels carrying values of type T, allowing any of the Acute methods for establishing shared typed names to be used.

Internally, migration uses thunkify. Migration and send_remote both use marshal, with a wire format of marshalled values of type

```
(group name * (exists t. t name * t)) + migration
```

for the message and migration cases, where

```

type migration = group name
  * group
  * mutex name * cvar name
  * (thunkkey list -> unit)

```

The recv_local and send_local use namecase (as in Local_channel). Marshalling of migrations is with respect to the mark "Npi_end" set below; marshalling for send_remote is with respect to the supplied mark, which should usually be below "Npi_end". There is some delicate use of local concurrency with mutexes and cvars.

*)

```
(* NB: fields marked by (*A*) will be removed from the interface *)
(* when width subsignaturing is added *)
```

```
(* Note the use of hash! (instead of fresh), as we need to rebind to
this interface on migration with type "group" being compatible *)
```

```

module hash! Npil :
sig
  type tf = (Tcp.ip option * Tcp.port) -> Tcp.addr -> string -> unit
  type tho = Tcp_string_messaging.handle option ref

  type channel = (exists t. t name * (t list ref * cvar name))

  type group = thread name list ref          (* threads in group *)
    * mutex name list ref                    (* mutexes in group *)
    * cvar name list ref                     (* cvars in group *)

```

```

        * channel list ref                (* local channels *)

type migration = group name
    * group
    * (thunkkey list -> unit)

val groups_mutex : mutex name
val groups : (group name * group) list ref
val threadmap : (thread name * group name) list ref
val ho: tho

end
=
struct

type tf = (Tcp.ip option * Tcp.port) -> Tcp.addr -> string -> unit
type tho = Tcp_string_messaging.handle option ref

type channel = (exists t. t name * (t list ref * cvar name))

type group = thread name list ref          (* threads in group *)
    * mutex name list ref                 (* mutexes in group *)
    * cvar name list ref                   (* cvars in group *)
    * channel list ref                     (* local channels *)

(* The group data structure is more generous than its usage:
   it allows also mutexes and condition variables to be associated
   with a group (and be migrated propely).
   At the moment there is no create_gmutex/create_gcvar, although
   their implementation would be trivial.
*)

type migration = group name
    * group
    * (thunkkey list -> unit)

let groups_mutex = hash(mutex, "Npi global mutex") %[mutex] (* fresh *) (* global mutex *)

(* Locking strategy:
   - There is a global mutex ("groups_mutex") at each running acute process.
   - Functions acting on the group data structures are all protected by this
     global lock.
   - When a thread wants to receive a message and there are none in the
     channel, the thread waits on the channel's condition variable.
   - When a new message is sent on empty channel, its condition variable is
     signalled so that a waiting receiver is unblocked.
NB: This does not in principle guarrantee a FIFO delivery order, but will in
    fact have a FIFO ordering with the current version of Acute as threads
    in a condition variable are stored in a FIFO queue.

The locking strategy is quite coarse; a more fine-grained scheme would be
possible, where besides the global lock, a lock per group is also kept.
*)

let groups = ref []                (* group name -> group *)

```



```

let threadmap = ref []      (* thread name -> group name *)

(* threadmap exists to find in which group a thread belongs to
   These maps are simply implemented as linked lists, but a production
   implementation would use a hashtable instead.
   Similarly the list of channels should really be a hashtable.
   *)

let ho = ref None

end

mark "Npi1"

module hash! Npi2 :
sig

  val find_my_group : unit -> Npi1.group name * Npi1.group

  val gthread_wrapper : forall t. (t->unit) -> t -> unit
  val create_group : forall t. (t -> unit) -> t -> unit
  val create_gthread : forall t. (t->unit) -> t -> unit

  val recv_local : forall t. t name -> t

  val my_send_local : forall t. Npi1.group -> t name -> t -> unit
  val send_local : forall t. t name -> t -> unit

  val f : Npi1.tf

  val init : (Tcp.ip option * Tcp.port option) -> unit

  val send_remote : forall t. string -> (Tcp.addr*Npi1.group name*t name) -> t -> unit

  val migrate_group : Tcp.addr -> unit

  val local_addr : unit -> Tcp.ip option * Tcp.port

end
=
struct

  (* returns which group the calling thread belongs to *)
  let find_my_group () = Uutils.locked_by_stmt2 %[] Npi1.groups_mutex
    (function () ->
      Pervasives.print_endline "In find_my_group lock...";
      let gn =
        try List.assoc %[] %[] (self ()) !Npi1.threadmap
        with Not_found ->
raise (Failure "find_my_group:assoc")
      in
      let group_info =
        try List.assoc %[] %[] gn !Npi1.groups
        with Not_found -> raise (Failure "find_my_group:assoc[2]")
      in
      (gn, group_info)
    )

```

```

(* Ensure that thread exits gracefully by unregistering itself from
 * the group data structure.
 *)
let gthread_wrapper = Function t -> fun (f: t -> unit) (v: t) ->
  f v
(*
  let unregister_my_gthread () =
    let tn = self() in
    let rec remove_me xs = match xs with
      [] -> raise Not_found
    | (x::xs) -> if x = tn then xs
                  else x :: remove_me xs in
    let (gn, (ths, _, _, _)) = find_my_group () in
    Utils.locked_by_stmt Np11.groups_mutex
      (function () ->
        Np11.threadmap := List.remove_assoc %[] %[] tn !Np11.threadmap;
        ths := remove_me !ths
      )
  in
  (try f v
   with e -> (try unregister_my_gthread () with _ -> ()); raise e);
  unregister_my_gthread ()
*)
(* create a new group *)
let create_group = Function t -> fun (f: t -> unit) (v : t) ->
  let gn = fresh %[Np11.group] in
  let tn = fresh %[thread] in
  Utils.locked_by_stmt Np11.groups_mutex
    (function () ->
      let group_info = (ref (tn::[]), ref [], ref [], ref []) in
      Np11.groups := (gn, group_info) :: !Np11.groups;
      Np11.threadmap := (tn, gn) :: !Np11.threadmap;
      create_thread tn (gthread_wrapper %[t] f) v )

(* create a new thread in the current group *)
let create_gthread = Function t -> fun (f: t -> unit) (v: t) ->
  let (gn, (ths, _, _, _)) = find_my_group () in
  let tn = fresh %[thread] in
  Utils.locked_by_stmt Np11.groups_mutex
    (function () ->
      Np11.threadmap := (tn, gn) :: !Np11.threadmap;
      ths := tn :: !ths;
      create_thread %[t] tn (gthread_wrapper %[t] f) v
    )

(* receive a value from a local channel, blocking if there is none *)
let rcv_local = Function t -> fun (cn: t name) ->
  let (gn, group_info) = find_my_group () in
  let (_,_,_,csr) = group_info in
  Utils.locked_by_stmt2 %[t] Np11.groups_mutex
    (function () ->
      let rec lookup cs' = match cs' with
        [] -> let my_cvar = fresh %[cvar] in
                create_cvar my_cvar;
                csr := ({t, (cn, (ref [], my_cvar))} as Np11.channel) :: !csr;
                wait my_cvar Np11.groups_mutex;
                lookup !csr
      in
    )

```

```

| (c: Npi1.channel)::cs0 ->
  namecase c with
  {t,(cn',x)} when cn'=cn ->
    let ((msgs: t list ref), my_cvar) = x in
      let rec ww () =
        match !msgs with
        [] -> wait my_cvar Npi1.groups_mutex; ww ()
        | v::vs -> msgs := vs; v
      in
        ww ()
    otherwise ->
      lookup cs0
  in lookup !csr
)

let my_send_local = Function t -> fun group_info (cn: t name) (v: t) ->
  let (_,_,_,csr) = group_info in
  Utils.locked_by_stmt Npi1.groups_mutex
  (function () ->
    let rec lookup cs' = match cs' with
      [] -> let my_cvar = fresh %[cvar] in
        create_cvar my_cvar;
        csr := ({t,(cn,(ref(v::[]),my_cvar))} as Npi1.channel) :: !csr
    | (c: Npi1.channel)::cs0 ->
      namecase c with
      {t,(cn',x)} when cn'=cn ->
        let ((msgs: t list ref), my_cvar) = x in
          (match !msgs with
            [] -> msgs := v :: !msgs; signal my_cvar
            | _ -> msgs := v :: !msgs)
          otherwise ->
            lookup cs0
    in lookup !csr
  )

let send_local = Function t -> fun (cn: t name) (v: t) ->
  let (gn, group_info) = find_my_group () in
  my_send_local %[t] group_info cn v

(* We have a single site daemon listen for messages and migrating things.
- for messages, it uses the group name to look up in the group data structure
to find the appropriate (local) channel handle, then use that to propagate
the message.
- for migrating things, it'll unthunkify and extend the group data structure.
*)
let f ipop_local addr_remote data =
  Utils.locked_by_stmt Npi1.groups_mutex
  (function () ->
    Pervasives.print_endline "npi daemon received something";
    try
      match (unmarshal data) with
      inj 1 %[(Npi1.group name * (exists t. t name * t)) + Npi1.migration] (gn, channel)
      -> (* a normal value *)
        Pervasives.print_endline "npi daemon received a value";
        let group_info = try List.assoc %[] %[] gn !Npi1.groups
          with Not_found -> raise (Failure

```

```

    "Received a value for a group not present at this TCP address")
in
  let {t, x} = channel in
  let (cn, v) = x in
  send_local %[t] cn v
| inj 2 %[(Npi1.group name*(exists t. t name*t))+Npi1.migration] (gn,groupinfo,unthunk)
-> (* a migration *)
  Pervasives.print_endline "npi daemon received a migration";
  let (ths, mtxs, cvs, csr) = groupinfo in

  if List.mem_assoc %[] %[] gn !Npi1.groups then
    (* NB: this should never occur as group names are only created with
       fresh %[group] and the only operation involving group names
       is migration which is linear.
       This check prevents a type of maliciously forged migrations.
    *)
    raise (Failure "A group with this same name is already present at this site")
  else (
    Npi1.groups := (gn, groupinfo) :: !Npi1.groups;
    List.iter %[] (fun tn -> Npi1.threadmap := (tn, gn) :: !Npi1.threadmap) !ths;
    let tks = List.map %[] %[] (fun n -> Thread (n, Blocking)) !ths
        @ List.map %[] %[] (fun n -> Mutex n) !mtxs
        @ List.map %[] %[] (fun n -> CVar n) !cvs
        @ List.map %[] %[] (fun (p: Npi1.channel) ->
            let {t,x} = p in let (_,(_,n)) = x in CVar n) !csr
    in
      unthunk tks;
      Pervasives.print_endline("unthunked")
    )
  with e -> Pervasives.print_endline "An exception was raised in the npi daemon";
    raise e
)

let init (ipo,po) =
  create_mutex Npi1.groups_mutex;
  Pervasives.print_endline ("Created NPI mutex " ^ name_to_string Npi1.groups_mutex);
  match !Npi1.ho with
  Some _ -> raise (Failure "Npi already initialised")
| None -> Npi1.ho := Some (Tcp_string_messaging.daemon (ipo,po) f)

let send_remote = Function t -> fun mk (addr,gn,cn) (v: t) ->
  let h = Utils.the %[] !Npi1.ho in
  let (ip, port) = addr in
  if (Some ip, port) = Tcp_string_messaging.local_addr h then
    (* note this local-send optimisation will only take effect if the
       IP was set explicitly *)
    let group_info = Utils.locked_by_stmt2 %[] Npi1.groups_mutex (function () ->
        try List.assoc %[] %[] gn !Npi1.groups
        with Not_found -> raise (Failure "send_remote:List.assoc")
    ) in
      my_send_local %[t] group_info cn v
  else
    let channel = {t, (cn, v)} as exists t'.t' name * t' in
    let data = inj 1 %[(Npi1.group name*(exists t.t name*t))+Npi1.migration] (gn, channel) in
    let mar_data = marshal mk data in
    Tcp_string_messaging.send h addr mar_data

```

```

(* Migrate the current group to a new Tcp address.
   All threads except for the calling thread are thunkified with Blocking mode.
   The called thread is blocked with a mutex/cvar. As it is marshalled with
   Interrupting mode, it is woken up at the other end with a Thunkify_EINTR
   exception.
*)
let migrate_group = fun addr ->
  Pervasives.print_endline("migrate_group: started");
  let (gn, group_info) = find_my_group () in
  Pervasives.print_endline("migrate_group: found my group");
  let (ths, mtxs, cvs, csr) = group_info in
  let my_cv = fresh in
  create_cvar my_cv;

  lock Npi1.groups_mutex;
  (* First remove the group and its threads from the global data structures *)
  Npi1.groups := List.remove_assoc %[] %[] gn !Npi1.groups; (* remove gn -> group_info mapping *)
  List.iter %[]
    (fun tn -> Npi1.threadmap := List.remove_assoc %[] %[] tn !Npi1.threadmap)
    !ths; (* remove tn -> gn mapping *)
  Pervasives.print_endline("migrate_group: removed gn,tn data");
  let initiating_thread_name = self() in
  (* make new thread to perform thunkify, otherwise will thunkify self *)
  create_thread fresh
    (function () ->
      Pervasives.print_endline("migrate_group: thunkify thread started");
      Utils.locked_by_stmt Npi1.groups_mutex
      (function () ->
        Pervasives.print_endline("migrate_group: thunkify thread got lock");
        let get_tmode tn =
          if compare_name tn initiating_thread_name = 0 then
            Interrupting
          else Blocking in
        let tks = List.map %[] %[] (fun n -> Thread (n, get_tmode n)) !ths
          @ List.map %[] %[] (fun n -> Mutex n) !mtxs
          @ List.map %[] %[] (fun n -> CVar n) !cvs
          @ List.map %[] %[] (fun (p: Npi1.channel) ->
            let {t,x} = p in let (_,(_,n)) = x in CVar n) !csr
        in
        Pervasives.print_endline("migrate_group: thunkify thread going to thunkify");
        let thunked = thunkify tks in
        Pervasives.print_endline("migrate_group: thunkify thread done thunkify");
        let data = inj 2 %[(Npi1.group name * (exists t. t name *
          t)) + Npi1.migration] (gn, group_info, thunked) in
        let mar_data = marshal "Npi_end" data in
        Pervasives.print_endline("migrate_group: going to send marshalled: ... "
          (* ^ mar_data *) );
        let h = Utils.the %[] !Npi1.ho in
        Tcp_string_messaging.send h addr mar_data
      )
    ) ();
  (* must block thread initiating migration, until thunkify has completed *)
  try
    wait my_cv Npi1.groups_mutex (* Block here - thunkify will cause Thunkify_EINTR *)
  with Thunkify_EINTR -> () (* Migration completed -- we can now continue execution *)

let local_addr () = Tcp_string_messaging.local_addr (Utils.the %[] !Npi1.ho)

```

```

end

mark "Npi2"

module hash! Npi :
sig
  type group

  val create_group : forall t. (t -> unit) -> t -> unit
  val create_gthread : forall t. (t->unit) -> t -> unit

  val recv_local : forall t. t name -> t
  val send_local : forall t. t name -> t -> unit

  val init : (Tcp.ip option * Tcp.port option) -> unit

  val send_remote : forall t. string -> (Tcp.addr * group name * t name) -> t -> unit
  val migrate_group : Tcp.addr -> unit

  val local_addr : unit -> Tcp.ip option * Tcp.port

end
=
struct

  type group = Npi1.group

  let create_group = Npi2.create_group
  let create_gthread = Npi2.create_gthread
  let recv_local = Npi2.recv_local
  let send_local = Npi2.send_local
  let init = Npi2.init
  let send_remote = Npi2.send_remote
  let migrate_group = Npi2.migrate_group
  let local_addr = Npi2.local_addr

end

mark "Npi_end"

(* ***** *)
(* ** ** *)
(* ** npi-recv client ** *)
(* ** ** *)
(* ***** *)

(* example npi client, initialising an npi daemon *)

includesource "npi.ac"

```

```
let addr (ip, port) = (Tcp.ip_of_string ip, Tcp.port_of_int port) in
```

```
let _ = Npi.init (Some(Tcp.ip_of_string "127.0.0.1"),  
                Some(Tcp.port_of_int 6401)) in
```

```
Pervasives.prerr_endline("npi-recv done initialising")
```

```
(* ***** *)  
(* ** ** *)  
(* ** npi-mig client ** *)  
(* ** ** *)  
(* ***** *)
```

```
(* example npi client, migrating an npi group there and back *)
```

```
includesource "npi.ac"
```

```
let addr (ip, port) = (Tcp.ip_of_string ip, Tcp.port_of_int port) in
```

```
let _ = Npi.init (Some(Tcp.ip_of_string "127.0.0.1"),  
                Some(Tcp.port_of_int 6400)) in
```

```
Pervasives.prerr_endline("npi-mig done initialising");
```

```
let _ = Npi.create_group %[]  
  (fun () ->  
    Pervasives.prerr_endline("group created");  
    Npi.migrate_group (addr ("127.0.0.1", 6401));  
    Pervasives.prerr_endline("group migrated");  
    Npi.migrate_group (addr ("127.0.0.1", 6400));  
    Pervasives.prerr_endline("group migrated back")  
  ) () in
```

```
()
```

Part V

Implementation

17 Overview

The implementation is written in FreshOCaml [SPG03], currently around 25 000 lines of code. It has been developed together with the language definition. By and large the definition has led, with extensions and changes to the definition being followed by implementation work to match. This exposed many ambiguities and errors in the semantics. In a few cases the implementation led, with changes propagated back into the definition afterwards. An automated testing framework helps ensure the two are in sync, with tests of compilation and execution that can be re-run automatically.

The main priority for the implementation was to be rather close to the semantics, to make it easy to change as the definition changed, and easy to have reasonable confidence that the two agree, while being efficient enough to run moderate examples. The runtime is essentially an interpreter over the abstract syntax, finding redexes and performing reduction steps as in the semantics. For efficiency it uses closures (as described in §16.12) and represents terms as pairs of an explicit evaluation context and the enclosed term (roughly as in [Rém02, §1.3.1, Ex. 1]) to avoid having to re-traverse the whole term when finding redexes. Marshalled values `marshalled(E_n , E_s , s , $definitions$, e , T)` are represented simply by a pretty-print of their abstract syntax. Numeric hashes use a hash function applied to a pretty-print of their body; it is thus important for this pretty-print to be canonical, choosing bound identifiers appropriately. Acute threads are reduced in turn, round-robin. A pool of OS threads is maintained for making blocking system calls. A `genlib` tool makes it easy to import (restricted versions of) OCaml libraries, taking OCaml `.mli` interface files and generating embeddings and projections between the OCaml and internal Acute representations. It does not support higher-order functions, which would be challenging in the presence of concurrency.

To give a *very* crude idea of performance, the initialisation phase of the `blockhead.ac` game performs about 220000 steps (roughly corresponding to reduction steps) in 4.5 seconds, without runtime typechecking and with the vacuous bracket optimisation. The naive Fibonacci function of 25

```
let rec fib:int->int = function (x:int) ->
  if x <=2 then
    1
  else
    (fib (x-1)) + (fib(x-2))
in
let x = fib 25
```

involves about 1.6 million steps and takes 18 seconds, again without runtime typechecking and with vacuous bracket optimisation. Running the same code in the OCaml toplevel takes 0.0056 seconds, so the Acute implementation is around 3000 times slower. Turning on runtime typechecking in Acute (and using `definitions_lib_small.ac`) for Fibonacci of 15 takes the execution time from 0.16 seconds to 495 seconds (11000 steps), a slowdown of another factor of 3000. These figures are all for a 3.20GHz Pentium 4. In practice this level of performance has been reasonable for the examples we have considered to date. The `blockhead` and `minesweeper` games are playable, and three sample communication infrastructures, based on Nomadic Pict, Distributed Join Calculus, and Ambients, all execute tolerably well. Runtime typechecking, while it would be good to have feasible for these larger examples, in fact is mostly useful for more focussed test cases, for which one wishes to observe the individual reduction steps in any case.

18 Command line options

acute <options> <filename>

where

the <level>s are:

- 0 none
- 1 expression
- 2 expression, store
- 3 expression, store, userdefns
- 4 expression, store, userdefns, libdefns

and options are:

```
-definitionslib <filename>      semantic: Read the standard definitions from <filename>
(default: definitions_lib_small.ac
but use definitions_lib.ac for full set)
-nodefinitionslib              semantic: No standard definitions
-o <filename>                  phase: Output to <filename> (default: <stdout>)
-df <filename>                 phase: Print final state dump to to <filename> (default: <stdout>)
-err <filename>                phase: Print debug output to <filename> (default: <stderr>)
-writefinal <filename>         phase: Pretty print result to <filename> (default: <stdout>)
-checkfinal <filename>        phase: Check result against contents of <filename> (default: None)
-emitobjectfile <filename>    phase: Emit compiled (object) code after compilation
-emitsourcefile <filename>   phase: Emit source code after compilation
-debug <class>[,<class>...]    output: Which classes of debug output to display
(default: default, flattenclos, desugar, tcopt, mkhash, lexer, evalstep, marshal, hashify, tcquant, linkok, namecase, nameval)
-dumpstepinterval <n>        output: Print the configuration (at dumptrace level) every <n> steps
-dumpfrom <n>                output: Only print the configuration (at dumptrace level) after <n> steps
-printstepinterval <n>      output: Print the reduction step count every <n> steps
-noprintstepinterval        output: Do not print the reduction step count
-production                  rttc: Set options used for a production implementation
<norttc><nomttc><notypecheckcompiled><lithash><nolinkok_sig_typecheck><hack_optimise>
-noproduction                rttc: Set options used for a non-production implementation
<rttc><mttc><typecheckcompiled><nolithash><linkok_sig_typecheck><nohack_optimise>
-tcdepth <depth>            (4) output: Context depth for typechecking errors
-dumpparse <level> (0-4)    (0) output: Dump result of parse
-dumppreinf <level> (0-4)  (0) output: Dump input to inference
-dumppostinf <level> (0-4) (0) output: Dump output of inference
-dumpdesugared <level> (0-4) (0) output: Dump output of desugaring
-dumpcompiled <level> (0-4) (3) output: Dump output of compilation
-dumptrace <level> (0-4)   (1) output: Dump traced execution steps
-dumpfinal <level> (0-4)   (1) output: Dump final state (if no type failure)
-dumptypefail <level> (0-4) (3) output: Dump on type failure (or unmarshal fail)
-[no]showpasses              (*) output: Show names of compilation passes
-[no]showtimes               (*) output: Show time taken per pass
-[no]showprogress           ( ) output: Show progress during type inference
-[no]showlocs                ( ) output: Show locations in dump output
-[no]showtrailer            (*) output: Show trailer information (e.g., hash values) when printing
-[no]suffixall              ( ) output: Always suffix names, even when unshadowed
-[no]shownames               ( ) output: Show internal representation of bound names
-[no]globalhashmap          (*) output: Use a common map for abbreviating hashes and abstract names
-[no]show_options           ( ) output: Show the command line used, including default options
-[no]showtcenv              ( ) output: Show environment in typecheck errors
-[no]emitobject              ( ) output: Emit compiled (object) code after compilation
-[no]printenv                ( ) output: Print runtime environments
-[no]printenvbodies          ( ) output: Print runtime environment bodies (RHSs)
-[no]printclos               ( ) output: Print closures as closures (rather than expanding)
-[no]printererrordeath      (*) output: Print error message when a thread exits with an exception
```

```
-[no]printcleandeath      ( ) output: Print message when a thread exits cleanly
-[no]debug                ( ) output: Generate debug output (on stderr)
-[no]showfocussing       ( ) output: Show focussing process in dumptrace
-[no]dumptex              ( ) output: Dump in tex format
-[no]dumphuman            ( ) output: Dump for humans (no type annotations)
-[no]dumpall              ( ) output: Don't ever abbreviate traces to ...
-[no]parsetest            ( ) phase: Parser - pretty printer identity test
-[no]desugar               (* ) phase: Desugar
-[no]compile               (* ) phase: Compile
-[no]typecheckcompiled     (* ) phase: Typecheck the compiled program
-[no]run                   (* ) phase: Run program
-[no]lithash               ( ) rttc: Emit literal 0#123ABC hashes in certain places
-[no]rttc                  (* ) rttc: Do runtime typechecking
-[no]mttc                  (* ) rttc: Do unmarshaltime typechecking
-[no]terminate_on_tc      (* ) rttc: Terminate if typecheckcompiled or rttc is on and fails
-[no]default               (* ) semantic: Default underspecified types to unit
-[no]disable_import_typecheck ( ) semantic: Disable typechecking of import links
-[no]disable_eqsok_typecheck ( ) semantic: Disable typechecking of |- eqs ok
-[no]internal_weqs        (* ) semantic: Allow use of with! equations inside modules
(not just at boundary)
-[no]linkok_sig_typecheck  (* ) semantic: Do full subsignature typecheck in linkok
(not just syntactic check)
-[no]hack_optimise         (* ) semantic: Perform vacuous-bracket optimisation
-[no]really_hack_optimise  ( ) semantic: Erase all brackets
-[no]abstract_existentials (* ) semantic: Dynamically-abstract existentials
-[no]nonunitthread         ( ) semantic: Threads do not have to evaluate to unit
-[no]marshaltex            ( ) semantic: Marshal in tex format (cannot be unmarshalled)
-help   Display this list of options
--help  Display this list of options
```

19 Concrete user source grammar

This is the concrete source grammar, automatically extracted from the implementation ocaml yacc source.

```

core_type                ::= core_type_pri
compilation_unit_definition ::= [ source_definition | includesource STRING |
                                includecompiled STRING ]
compilation_unit_definitions ::= { compilation_unit_definition semisemis }
nameenv                  ::= { ( ) | nameenv_non_empty } )
nameenv_non_empty        ::= [ { nameenv_entry , } nameenv_entry ]
nameenv_entry             ::= [ ABSTRNAME : ( nmodule modname_extern hmodule_body |
                                nimport modname_extern himport_body | Type | core_type_pri )
                                ]
definitions               ::= { definition }
optional_mode             ::= [ hash | hash! | cfresh! | cfresh | fresh ]
definition                ::= [ cmodule modname_binder cmodule_body | cimport
                                modname_binder cimport_body | module fresh
                                modname_binder module_body | import fresh
                                modname_binder import_body | mark STRING ]
source_definition         ::= [ module optional_mode modname_binder module_body |
                                amodule modname_binder amodule_body | import
                                optional_mode modname_binder import_body | mark STRING ]
module_body               ::= : module_type version_opt = module_expr withspec_opt
valuability               ::= valuable
                            | cvaluable
                            | nonvaluable
valuabilities              ::= ( valuability , valuability )
cmodule_body              ::= hash : eqs module_type valuabilities module_type version_val =
                                module_expr
hmodule_body              ::= : eqs module_type version_nonopt = module_expr
amodule_body              ::= : module_type = modname_use
import_body               ::= : module_type version_constraint_opt likespec resolvespec_opt
                                moo_module_opt
cimport_body              ::= hash : module_type valuabilities module_type
                                version_constraint_val likestr resolvespec_nonopt moo_module
himport_body              ::= : module_type version_constraint_nonopt likestr
hash                       ::= hash ( hmodule modname_extern hmodule_body )
                            | hash ( himport modname_extern himport_body )
                            | LITHASH
                            | ABSTRNAME
hash_or_modname_dot_ident ::= [ ( hash | modname_use ) . ident_extern ]
name_value                 ::= [ name_value ( ( hash ( hash . ident_extern ) app_ty ) | hash (
                                core_type_pri , STRING ) ) | hash ( core_type_pri , STRING ,
                                name_value ) ) | ABSTRNAME app_ty ) ) ]
eqs                        ::= { ( ) | eqs_body_non_empty } )
eqs_body_non_empty         ::= [ eqs_body_item [ , eqs_body_non_empty ] ]
eqs_body_item              ::= [ ( hash | modname_use ) . typename_extern = core_type_pri ]
version_opt                ::= [ version version ]
version_val                ::= version version
version_nonopt             ::= version version
version_constraint_val     ::= version version_constraint
version_constraint_nonopt  ::= version version_constraint
version_constraint_opt     ::= [ version version_constraint ]

```

```

withspec_opt      ::= [ with! weqs ]
weqs_single       ::= modname_use . typename_extern = core_type_pri
weqs              ::= weqs_rev
weqs_rev          ::= [ { weqs_single , } weqs_single ]
likespec          ::= [ like modname_use | likestr ]
likestr           ::= [ like struct structure end ]
resolvespec_nonopt ::= [ by resolvespec_non_empty ]
resolvespec_opt   ::= [ resolvespec_nonopt ]
moo_module_opt    ::= [ [ moo_module ] ]
moo_module        ::= [ = ( unlinked | modname_use ) ]
module_expr       ::= struct structure end
module_type       ::= sig signature end
structure_items   ::= [ structure_item ( ; ; structure_items | structure_items ) ]
structure         ::= [ structure_item ( ; ; structure_items | structure_items ) ]
structure_item    ::= let ident_binder = typed_expr
                    | let ident_binder non_empty_pattern_list = typed_expr
                    | type typename_binder = core_type_pri
signature_items   ::= [ signature_item ( ; ; signature_items | signature_items ) ]
signature         ::= [ signature_item ( ; ; signature_items | signature_items ) ]
signature_item    ::= val ident_binder : core_type_pri
                    | type typename_binder
                    | type typename_binder = core_type_pri
                    | type typename_binder : kind
marshalled_body   ::= marshalled_nameenv_opt , { definitions } , { loctyp_list } , {
                    store } , simple_expr , core_type_pri
marshalled_nameenv_opt ::= -
                    | nameenv
marshalled_value_pri ::= marshalled ( marshalled_body )
store              ::= [ store_non_empty ]
store_non_empty    ::= [ { store_item , } store_item ]
store_item         ::= ( location := expr )
hash_in_version    ::= hash ( hmodule modname_extern hmodule_body )
                    | hash ( himport modname_extern himport_body )
                    | LITHASH
                    | ABSTRNAME
version_literal    ::= INT
                    | hash_in_version
version            ::= atomic_version [ version_dotted_suffix ]
version_dotted_suffix ::= { . atomic_version } . atomic_version
atomic_version     ::= myname
                    | version_literal
atomic_hash_version_constraint ::= [ modname_use | hash_in_version ]
atomic_version_constraint ::= [ atomic_hash_version_constraint | INT ]
atomic_version_constraints_non_empty ::= [ { atomic_version_constraint . } atomic_version_constraint ]
tail_version_constraint ::= atomic_version_constraint
                    | INT - INT
                    | - INT
                    | INT -
                    | *
version_constraint ::= [ name = atomic_hash_version_constraint | tail_version_constraint
                    | atomic_version_constraints_non_empty . tail_version_constraint
                    ]
resolvespec_non_empty ::= [ resolvespec_item [ , resolvespec_non_empty ] ]

```

```

resolvespec_item ::= Static_Link
                  | Here_Already
                  | STRING
seq_expr         ::= [ expr [ ( ; | ||| ) seq_expr ] ]
expr             ::= simple_expr
                  | simple_expr simple_expr_or_app_ty_list
                  | let pattern = typed_expr in seq_expr
                  | let ident_internal_binder non_empty_pattern_list = typed_expr in
                    seq_expr
                  | let rec ident_internal_binder non_empty_pattern_list =
                    typed_expr in seq_expr
                  | let rec ident_internal_binder optional_colon_core_type_pri =
                    function mtch.when_sugary in seq_expr
                  | match seq_expr with mtch
                  | function mtch.when_sugary
                  | fun non_empty_pattern_list -> seq_expr
                  | try seq_expr with mtch
                  | ref opt_ty simple_expr
                  | ref opt_ty
                  | raise simple_expr
                  | if seq_expr then expr else expr
                  | while seq_expr do seq_expr done
                  | expr :: expr
                  | expr && expr
                  | expr || expr
                  | expr := opt_ty expr
                  | expr = opt_ty expr
                  | expr @ opt_ty expr
                  | expr + expr
                  | expr - expr
                  | expr * expr
                  | expr > expr
                  | expr < expr
                  | expr INFIXOP0 expr
                  | expr INFIXOP1 expr
                  | expr INFIXOP2 expr
                  | expr INFIXOP3 expr
                  | expr INFIXOP4 expr
                  | expr freshfor expr
                  | - expr
                  | Function typename_internal_binder -> seq_expr
                  | let { typename_internal_binder , ident_internal_binder } =
                    typed_expr in seq_expr
                  | namecase expr with { typename_internal_binder , (
                    ident_internal_binder , ident_internal_binder ) } when ident_use =
                    expr -> expr otherwise -> expr
typed_expr      ::= seq_expr
                  | seq_expr : loc_core_type
                  | seq_expr as loc_core_type
                  | typed_expr1
                  | seq_expr ; typed_expr1
                  | seq_expr ||| typed_expr1
typed_expr1     ::= { core_type_pri , expr } as core_type_pri

```

```

simple_expr ::= { constr0 | ident_use | econst_use | modname_use . ident_extern |
              hash . ident_extern | location | ( typed_expr ) | (
              expr_comma_list ) | ! opt_ty simple_expr | constr1 simple_expr |
              standalone_infixop | fresh opt_ty | cfresh opt_ty | hash (
              hash_or_modname_dot_ident ) app_ty | hash ( core_type_pri ,
              expr ) app_ty | hash ( core_type_pri , expr , expr ) app_ty |
              name_value | swap expr and expr in simple_expr | support
              opt_ty simple_expr | modname_use @ ident_extern | name_of_tie
              simple_expr | val_of_tie simple_expr | PREFIXOP |
              PREFIXOP.TYP opt_ty | marshal simple_expr simple_expr |
              unmarshal }

simple_expr_or_app_ty_list ::= simple_expr
                          | app_ty
                          | simple_expr simple_expr_or_app_ty_list
                          | app_ty simple_expr_or_app_ty_list

opt_ty ::= [ [ %[ core_type_pri ] ] ]
app_ty ::= [ %[ ( core_type_pri ) | ] ]
optional_colon_core_type_pri ::= [ [ : core_type_pri ] ]
location ::= { < INT > }
loctyp_list ::= [ loctyp_list_non_empty ]
loctyp_list_non_empty ::= { loctyp_pair , } loctyp_pair
loctyp_pair ::= ( location : core_type_pri )
mtch ::= [ [ ] match_cases ]
mtch_when_sugary ::= [ mtch | ( ident_internal_binder : core_type_pri ) match_action ]
match_cases ::= pattern_match_action { | pattern_match_action }
pattern_match_action ::= pattern match_action
match_action ::= -> seq_expr
expr_comma_list ::= ( expr_comma_list | expr ) , expr
standalone_infixop ::= ( ( standalone_infixopstr ) | && ) | || ) | ! opt_ty ) | = opt_ty ) |
                    := opt_ty ) | @ opt_ty ) )

standalone_infixopstr ::= [ + | - | * | < | > | INFIXOP0 | INFIXOP1 | INFIXOP2 |
                          INFIXOP3 | INFIXOP4 ]

pattern ::= pattern_pri
pattern_pri ::= simple_pattern
              | constr1 simple_pattern
              | pattern_pri :: pattern_pri

simple_pattern ::= ident_internal_binder
              | -
              | constr0
              | - INT
              | ( pattern_pri )
              | ( pattern_pri : core_type_pri )
              | ( pattern_comma_list )

pattern_comma_list ::= ( pattern_comma_list | pattern_pri ) , pattern_pri
non_empty_pattern_list ::= non_empty_rev_pattern_list
non_empty_rev_pattern_list ::= { pattern_pri } pattern_pri
kind ::= Type
      | Eq ( core_type_pri )

loc_core_type ::= core_type_pri
core_type_pri ::= fun_core_type
              | forall typename_internal_binder . core_type_pri
              | exists typename_internal_binder . core_type_pri

fun_core_type ::= tup_core_type { -> tup_core_type }

```

```

simple_core_type      ::= ( core_type_pri )
                    |  typename_constr_use0
                    |  modname_use . typename_extern
                    |  hash . typename_extern
                    |  simple_core_type ref
                    |  simple_core_type name
                    |  simple_core_type typename_constr_use1
tup_core_type        ::= simple_core_type [ * core_type_list_tuple | + core_type_list_sum ]
core_type_list_tuple ::= simple_core_type { * simple_core_type }
core_type_list_sum   ::= simple_core_type { + simple_core_type }
constr0              ::= [ ] opt_ty
                    |  None opt_ty
                    |  baseconstr0
baseconstr0          ::= ( )
                    |  INT
                    |  false
                    |  true
                    |  CHAR
                    |  STRING
                    |  BASECON0
constr1              ::= inj INT app_ty
                    |  Some
                    |  tiecon
                    |  NODE
                    |  BASECON1
ident_use            ::= LIDENT
econst_use           ::= ECONST
ident_binder         ::= LIDENT
ident_internal_binder ::= LIDENT
ident_extern         ::= LIDENT
typename_constr_use0 ::= LIDENT
typename_constr_use1 ::= LIDENT
typename_binder      ::= LIDENT
typename_extern      ::= LIDENT
typename_internal_binder ::= LIDENT
modname_use          ::= [ UIDENT ]
modname_binder       ::= UIDENT
modname_extern       ::= UIDENT
semisemis            ::= [ [ semisemis_plus ] ]
semisemis_plus       ::= [ { ; ; } ; ; ]

```

20 Concrete compiled-form grammar

This is the concrete compiled-form grammar, automatically extracted from the implementation `ocamlyacc` source.

```

core_type                ::= core_type_pri
compilation_unit_definition ::= [ source_definition | includesource STRING |
                                includecompiled STRING ]
compilation_unit_definitions ::= { compilation_unit_definition semisemis }
nameenv                  ::= { ( ) | nameenv_non_empty } )
nameenv_non_empty        ::= [ { nameenv_entry , } nameenv_entry ]
nameenv_entry             ::= [ ABSTRNAME : ( nmodule modname_extern hmodule_body |
                                nimport modname_extern himport_body | Type | core_type_pri )
                                ]
definitions               ::= { definition }
optional_mode             ::= [ hash | hash! | cfresh! | cfresh | fresh ]
definition                ::= [ cmodule modname_binder cmodule_body | cimport
                                modname_binder cimport_body | module fresh
                                modname_binder module_body | import fresh
                                modname_binder import_body | mark STRING ]
source_definition         ::= [ module optional_mode modname_binder module_body |
                                amodule modname_binder amodule_body | import
                                optional_mode modname_binder import_body | mark STRING ]
module_body               ::= : module_type version_opt = module_expr withspec_opt
valuability               ::= valuable
                            | cvaluable
                            | nonvaluable
valuabilities              ::= ( valuability , valuability )
cmodule_body              ::= hash : eqs module_type valuabilities module_type version_val =
                                module_expr
hmodule_body              ::= : eqs module_type version_nonopt = module_expr
amodule_body              ::= : module_type = modname_use
import_body               ::= : module_type version_constraint_opt likespec resolvespec_opt
                                moo_module_opt
cimport_body              ::= hash : module_type valuabilities module_type
                                version_constraint_val likestr resolvespec_nonopt moo_module
himport_body              ::= : module_type version_constraint_nonopt likestr
hash                       ::= hash ( hmodule modname_extern hmodule_body )
                            | hash ( himport modname_extern himport_body )
                            | LITHASH
                            | ABSTRNAME
hash_or_modname_dot_ident ::= [ ( hash | modname_use ) . ident_extern ]
name_value                 ::= [ name_value ( ( hash ( hash . ident_extern ) app_ty ) | hash (
                                core_type_pri , STRING ) ) | hash ( core_type_pri , STRING ,
                                name_value ) ) | ABSTRNAME app_ty ) ) ]
eqs                        ::= { ( ) | eqs_body_non_empty } )
eqs_body_non_empty         ::= [ eqs_body_item [ , eqs_body_non_empty ] ]
eqs_body_item              ::= [ ( hash | modname_use ) . typename_extern = core_type_pri ]
version_opt                ::= [ version version ]
version_val                ::= version version
version_nonopt             ::= version version
version_constraint_val     ::= version version_constraint
version_constraint_nonopt ::= version version_constraint
version_constraint_opt     ::= [ version version_constraint ]

```



```

withspec_opt      ::= [ with! weqs ]
weqs_single       ::= modname_use . typename_extern = core_type_pri
weqs              ::= weqs_rev
weqs_rev          ::= [ { weqs_single , } weqs_single ]
likespec          ::= [ like modname_use | likestr ]
likestr           ::= [ like struct structure end ]
resolvespec_nonopt ::= [ by resolvespec_non_empty ]
resolvespec_opt   ::= [ resolvespec_nonopt ]
moo_module_opt    ::= [ moo_module ]
moo_module        ::= [ = ( unlinked | modname_use ) ]
module_expr       ::= struct structure end
module_type       ::= sig signature end
structure_items   ::= { structure_item }
structure          ::= [ structure_item structure_items ]
structure_item    ::= let ident_binder = typed_expr
                  | type typename_binder = core_type_pri
signature_items   ::= { signature_item }
signature         ::= [ signature_item signature_items ]
signature_item    ::= val ident_binder : core_type_pri
                  | type typename_binder
                  | type typename_binder = core_type_pri
                  | type typename_binder : kind
marshalled_body   ::= marshalled_nameenv_opt , { definitions } , { loctyp_list } , {
                  store } , simple_expr , core_type_pri
marshalled_nameenv_opt ::= -
                  | nameenv
marshalled_value_pri ::= marshalled ( marshalled_body )
store              ::= [ store_non_empty ]
store_non_empty    ::= [ { store_item , } store_item ]
store_item         ::= ( location := expr )
hash_in_version    ::= hash ( hmodule modname_extern hmodule_body )
                  | hash ( himport modname_extern himport_body )
                  | LITHASH
                  | ABSTRNAME
version_literal    ::= INT
                  | hash_in_version
version            ::= atomic_version [ version_dotted_suffix ]
version_dotted_suffix ::= { . atomic_version } . atomic_version
atomic_version     ::= myname
                  | version_literal
atomic_hash_version_constraint ::= [ modname_use | hash_in_version ]
atomic_version_constraint ::= [ atomic_hash_version_constraint | INT ]
atomic_version_constraints_non_empty ::= [ { atomic_version_constraint . } atomic_version_constraint ]
tail_version_constraint ::= atomic_version_constraint
                  | INT - INT
                  | - INT
                  | INT -
                  | *
version_constraint ::= [ name = atomic_hash_version_constraint | tail_version_constraint
                  | atomic_version_constraints_non_empty . tail_version_constraint
                  ]
resolvespec_non_empty ::= [ resolvespec_item [ , resolvespec_non_empty ] ]
resolvespec_item    ::= Static_Link

```

```

| Here_Already
| STRING
seq_expr ::= [ expr [ ( ; | ||| ) seq_expr ] ]
expr ::= simple_expr
| simple_expr simple_expr_or_app_ty_list
| let rec ident_internal_binder optional_colon_core_type_pri =
  function mtch_when_sugary in seq_expr
| match seq_expr with mtch
| function mtch_when_sugary
| try seq_expr with mtch
| ref opt_ty simple_expr
| raise simple_expr
| if seq_expr then expr else expr
| while seq_expr do seq_expr done
| expr :: expr
| expr && expr
| expr || expr
| expr := opt_ty expr
| expr = opt_ty expr
| expr @ opt_ty expr
| expr + expr
| expr - expr
| expr * expr
| expr > expr
| expr < expr
| expr INFIXOP0 expr
| expr INFIXOP1 expr
| expr INFIXOP2 expr
| expr INFIXOP3 expr
| expr INFIXOP4 expr
| expr freshfor expr
| - expr
| Function typename_internal_binder -> seq_expr
| let { typename_internal_binder , ident_internal_binder } =
  typed_expr in seq_expr
| namecase expr with { typename_internal_binder , (
  ident_internal_binder , ident_internal_binder ) } when ident_use =
  expr -> expr otherwise -> expr
typed_expr ::= seq_expr
| seq_expr : loc_core_type
| seq_expr as loc_core_type
| typed_expr1
| seq_expr ; typed_expr1
| seq_expr ||| typed_expr1
typed_expr1 ::= { core_type_pri , expr } as core_type_pri

```

```

simple_expr ::= { constr0 | ident_use | econst_use | modname_use . ident_extern |
              hash . ident_extern | location | ( typed_expr ) | (
              expr_comma_list ) | ! opt_ty simple_expr | constr1 simple_expr |
              standalone_infixop | fresh opt_ty | cfresh opt_ty | hash (
              hash_or_modname_dot_ident ) app_ty | hash ( core_type_pri ,
              expr ) app_ty | hash ( core_type_pri , expr , expr ) app_ty |
              name_value | swap expr and expr in simple_expr | support
              opt_ty simple_expr | modname_use @ ident_extern | name_of_tie
              simple_expr | val_of_tie simple_expr | PREFIXOP |
              PREFIXOP.TYP opt_ty | [ expr ] _ [ ] ^ { core_type_pri } | [ expr
              ] _ [ eqs_body_non_empty ] ^ { core_type_pri } | marshal
              simple_expr simple_expr | marshalz STRING simple_expr |
              unmarshal }

simple_expr_or_app_ty_list ::= simple_expr
                          | app_ty
                          | simple_expr simple_expr_or_app_ty_list
                          | app_ty simple_expr_or_app_ty_list

opt_ty ::= [ %[ core_type_pri ] ]
app_ty ::= [ %[ core_type_pri ] ]
optional_colon_core_type_pri ::= [ : core_type_pri ]
location ::= { < INT > }
loctyp_list ::= [ loctyp_list_non_empty ]
loctyp_list_non_empty ::= { loctyp_pair , } loctyp_pair
loctyp_pair ::= ( location : core_type_pri )
mtch ::= [ match_cases ]
mtch_when_sugary ::= [ mtch | ( ident_internal_binder : core_type_pri ) match_action ]
match_cases ::= pattern_match_action { | pattern_match_action }
pattern_match_action ::= pattern match_action
match_action ::= -> seq_expr
expr_comma_list ::= ( expr_comma_list | expr ) , expr
standalone_infixop ::= ( ( standalone_infixopstr ) | && ) | || ) | ! opt_ty ) | = opt_ty ) |
                    := opt_ty ) | @ opt_ty ) )

standalone_infixopstr ::= [ + | - | * | < | > | INFIXOP0 | INFIXOP1 | INFIXOP2 |
                          INFIXOP3 | INFIXOP4 ]

pattern ::= pattern_pri
pattern_pri ::= simple_pattern
              | constr1 simple_pattern
              | pattern_pri :: pattern_pri

simple_pattern ::= ident_internal_binder
              | -
              | constr0
              | - INT
              | ( pattern_pri )
              | ( pattern_pri : core_type_pri )
              | ( pattern_comma_list )

pattern_comma_list ::= ( pattern_comma_list | pattern_pri ) , pattern_pri
kind ::= Type
       | Eq ( core_type_pri )

loc_core_type ::= core_type_pri
core_type_pri ::= fun_core_type
               | forall typename_internal_binder . core_type_pri
               | exists typename_internal_binder . core_type_pri

fun_core_type ::= tup_core_type { -> tup_core_type }

```

```

simple_core_type      ::= ( core_type_pri )
                    |  typename_constr_use0
                    |  modname_use . typename_extern
                    |  hash . typename_extern
                    |  simple_core_type ref
                    |  simple_core_type name
                    |  simple_core_type typename_constr_use1
tup_core_type        ::= simple_core_type [ * core_type_list_tuple | + core_type_list_sum ]
core_type_list_tuple ::= simple_core_type { * simple_core_type }
core_type_list_sum   ::= simple_core_type { + simple_core_type }
constr0              ::= [ ] opt_ty
                    |  None opt_ty
                    |  baseconstr0
baseconstr0          ::= ( )
                    |  INT
                    |  false
                    |  true
                    |  CHAR
                    |  STRING
                    |  BASECON0
constr1              ::= inj INT app_ty
                    |  Some
                    |  tiecon
                    |  NODE
                    |  BASECON1
ident_use            ::= LIDENT
econst_use          ::= ECONST
ident_binder         ::= [ LIDENT [ LIDENT ] ]
ident_internal_binder ::= LIDENT
ident_extern         ::= LIDENT
typename_constr_use0 ::= LIDENT
typename_constr_use1 ::= LIDENT
typename_binder      ::= [ LIDENT [ LIDENT ] ]
typename_extern      ::= LIDENT
typename_internal_binder ::= LIDENT
modname_use          ::= [ UIDENT [ UIDENT ] ]
modname_binder       ::= [ UIDENT [ UIDENT ] ]
modname_extern       ::= UIDENT
semisemis            ::= [ [ semisemis_plus ] ]
semisemis_plus      ::= [ { ; ; } ; ; ]

```

21 Library interfaces

The following libraries are semi-automatically imported from OCaml – see the OCaml documentation for their semantics. For the moment, for historical reasons, the types are mostly concretized. They are subject to frequent change.

(* Automatically generated by genlib.ml. Do not edit directly! *)

```

module hash!Pervasives : sig
  val min : int -> int -> int
  val max : int -> int -> int
  val not : bool -> bool
  val abs : int -> int
  val lnot : int -> int
  val int_of_char : char -> int
  val char_of_int : int -> char
  val string_of_bool : bool -> string
  val bool_of_string : string -> bool
  val string_of_int : int -> string
  val int_of_string : string -> int
  val print_char : char -> unit
  val print_string : string -> unit
  val print_int : int -> unit
  val print_endline : string -> unit
  val print_newline : unit -> unit
  val prerr_char : char -> unit
  val prerr_string : string -> unit
  val prerr_int : int -> unit
  val prerr_endline : string -> unit
  val prerr_newline : unit -> unit
  val read_line : unit -> string
  val read_int : unit -> int
end

module hash!Agraphics : sig
  val open_graph : string -> unit
  val close_graph : unit -> unit
  val set_window_title : string -> unit
  val clear_graph : unit -> unit
  val size_x : unit -> int
  val size_y : unit -> int
  val rgb : int -> int -> int -> int
  val set_color : int -> unit
  val background : unit -> int
  val foreground : unit -> int
  val black : unit -> int
  val white : unit -> int
  val red : unit -> int
  val green : unit -> int
  val blue : unit -> int
  val yellow : unit -> int
  val cyan : unit -> int
  val magenta : unit -> int
  val plot : int -> int -> unit
  val plots : (int * int) list -> unit
  val point_color : int -> int -> int
  val moveto : int -> int -> unit
  val rmoveto : int -> int -> unit

```

```
val current_x : unit -> int
val current_y : unit -> int
val current_point : unit -> int * int
val lineto : int -> int -> unit
val rlineto : int -> int -> unit
val curveto : int * int -> int * int -> int * int -> unit
val draw_rect : int -> int -> int -> int -> unit
val draw_poly_line : (int * int) list -> unit
val draw_poly : (int * int) list -> unit
val draw_segments : (int * int * int * int) list -> unit
val draw_arc : int -> int -> int -> int -> int -> int -> unit
val draw_ellipse : int -> int -> int -> int -> unit
val draw_circle : int -> int -> int -> unit
val set_line_width : int -> unit
val draw_char : char -> unit
val draw_string : string -> unit
val set_font : string -> unit
val set_text_size : int -> unit
val text_size : string -> int * int
val fill_rect : int -> int -> int -> int -> unit
val fill_poly : (int * int) list -> unit
val fill_arc : int -> int -> int -> int -> int -> int -> unit
val fill_ellipse : int -> int -> int -> int -> unit
val fill_circle : int -> int -> int -> unit
val transp : unit -> int
val wait_next_event : int list -> int * int * bool * bool * char
val mouse_pos : unit -> int * int
val button_down : unit -> bool
val read_key : unit -> char
val key_pressed : unit -> bool
val sound : int -> int -> unit
val auto_synchronize : bool -> unit
val synchronize : unit -> unit
val display_mode : bool -> unit
val remember_mode : bool -> unit
end

module hash!Char : sig
  val code : char -> int
  val chr : int -> char
  val escaped : char -> string
  val lowercase : char -> char
  val uppercase : char -> char
  val compare : char -> char -> int
  val unsafe_chr : int -> char
end

module hash!String : sig
  val length : string -> int
  val get : string -> int -> char
  val create : int -> string
  val make : int -> char -> string
  val copy : string -> string
  val sub : string -> int -> int -> string
  val concat : string -> string list -> string
  val escaped : string -> string
  val index : string -> char -> int
```

```
val rindex : string -> char -> int
val index_from : string -> int -> char -> int
val rindex_from : string -> int -> char -> int
val contains : string -> char -> bool
val contains_from : string -> int -> char -> bool
val rcontains_from : string -> int -> char -> bool
val uppercase : string -> string
val lowercase : string -> string
val capitalize : string -> string
val uncapitalize : string -> string
val compare : string -> string -> int
end

module hash!Sys : sig
  val file_exists : string -> bool
  val remove : string -> unit
  val rename : string -> string -> unit
  val getenv : string -> string
  val command : string -> int
  val chdir : string -> unit
  val getcwd : unit -> string
  val catch_break : bool -> unit
end

module hash!Tcp : sig
  type fd : Type
  type ip : Type
  type port : Type
  type addr : Eq(ip * port)
  type netmask : Type
  type ifid : Type
  type msgbflag : Eq(int)
  type sock_type : Eq(int)
  val ip_of_string : string -> ip
  val string_of_ip : ip -> string
  val port_of_int : int -> port
  val int_of_port : port -> int
  val fd_of_int_private : int -> fd
  val int_of_fd : fd -> int
  val ifid_of_string2 : string -> ifid
  val string_of_ifid2 : ifid -> string
  val netmask_of_int2 : int -> netmask
  val int_of_netmask2 : netmask -> int
  val accept : fd -> fd * (ip * port)
  val bind : fd -> ip option -> port option -> unit
  val close : fd -> unit
  val connect : fd -> ip -> port option -> unit
  val dup : fd -> fd
  val dupfd : fd -> int -> fd
  val getifaddrs2 : unit -> (ifid * ip * ip list * netmask) list
  val getsockname : fd -> ip option * port option
  val getpeername : fd -> ip * port
  val getsockerr : fd -> unit
  val getsocklistening : fd -> bool
  val listen : fd -> int -> unit
  val pselect2 : fd list -> fd list -> fd list -> (int * int) option -> fd list * (fd list * fd list)
  val recv : fd -> int -> msgbflag list -> string * ((ip option * port option) * bool) option
```

```
val send : fd -> (ip * port) option -> string -> msgbflag list -> string
val shutdown : fd -> bool -> bool -> unit
val sockatmark : fd -> bool
val socket : int -> fd
val tcp_socket : unit -> fd
val udp_socket : unit -> fd
end
```

```
module hash!Persist : sig
  val write : string -> unit
  val read : unit -> string
  val write2 : string -> unit
  val read2 : unit -> string
end
```

```
module hash!Digest : sig
  val string : string -> string
  val substring : string -> int -> int -> string
  val file : string -> string
  val to_hex : string -> string
end
```

```
module hash!Filename : sig
  val concat : string -> string -> string
  val is_relative : string -> bool
  val is_implicit : string -> bool
  val check_suffix : string -> string -> bool
  val chop_suffix : string -> string -> string
  val chop_extension : string -> string
  val basename : string -> string
  val dirname : string -> string
  val temp_file : string -> string -> string
  val quote : string -> string
end
```

```
module hash!Unix : sig
  val sleep : int -> unit
end
```


22 The IO module

For writing concise examples we use either the persistent store IO module or the following network IO module, which implements send and receive with loopback TCP and provides brief aliases for `Pervasives.print_string` and `Pervasives.print_int`.

```

module IO :
sig
  val print_int : int->unit
  val print_string : string -> unit
  val print_newline : unit -> unit
  val send : string -> unit
  val receive : unit -> string
end =
struct
  let print_int = function x ->
    Pervasives.print_int x

  let print_string = function s ->
    Pervasives.print_string s

  let print_newline = function () ->
    Pervasives.print_newline ()

  let send = function data ->
    let fdesc = Tcp.tcp_socket () in
    let _ = Tcp.connect fdesc (Tcp.ip_of_string "127.0.0.1") (Some (Tcp.port_of_int 6666)) in
    let pad =
      function s ->
        function n ->
          let padding =
            String.make ( n - (String.length s) ) ' ' in
            (s ^ padding) in
    let data_length = String.length data in
    let data_length_string =
      pad (Pervasives.string_of_int data_length) 21 in
    let rec send_all = function s ->
      let no_options = [] in
      let s' = (Tcp.send fdesc None s no_options) in
      if 0 = (String.length s') then () else send_all s' in
    send_all (data_length_string ^ data);
    Tcp.close fdesc

  let receive = function () ->
    let fdesc = Tcp.tcp_socket () in
    let _ = Tcp.bind fdesc (Some (Tcp.ip_of_string "127.0.0.1"))
      (Some (Tcp.port_of_int 6666)) in
    let _ = Tcp.listen fdesc 5 in
    let (fdesc2,(_,_))= Tcp.accept fdesc in
    let rec rcv_n_bytes = function n ->
      let no_options = [] in
      let (s, _) = Tcp.rcv fdesc2 n no_options in
      let l = String.length s in
      if l >= n then s else s ^ (rcv_n_bytes (n-l)) in
    let data_length_string = rcv_n_bytes 21 in
    let first_space = String.index data_length_string ' ' in
    let data_length_string' = String.sub data_length_string 0 first_space in

```

```
let data_length = Pervasives.int_of_string data_length_string' in
let data = recv_n_bytes data_length in
Tcp.close fdesc;
Tcp.close fdesc2;
data
end
```

Appendix

This appendix gives most of the Acute syntax for reference. This is the fully type-annotated source language, including sugared forms, together with other non-source constructs that are needed to express the semantics. The implementation can infer many of the type annotations, and the *mode*, *withspec*, *likespec*, *vce*, *vne*, and *resolvespec* annotations on **module** and **import** default to reasonable values if omitted. The internal parts M , t and x of identifiers M_M , t_t and x_x are inferred by scope resolution.

Novel source features are highlighted in green and novel non-source constructs are highlighted in yellow.

Abstract names n **Store locations** l **Standard library constants** (with arity) x^n

Kinds

$K ::= \text{TYPE} \mid \text{EQ}(T)$

Types

$T ::= \text{int} \mid \text{bool} \mid \text{string} \mid \text{unit} \mid \text{char} \mid \text{void} \mid T_1 * .. * T_n \mid T_1 + .. + T_n \mid T \rightarrow T' \mid T \text{ list} \mid T \text{ option} \mid T \text{ ref} \mid \text{exn} \mid M_M.t \mid t \mid \forall t. T \mid \exists t. T \mid T \text{ name} \mid T \text{ tie} \mid \text{thread} \mid \text{mutex} \mid \text{cvar} \mid \text{thunkifymode} \mid \text{thunkkey} \mid \text{thunklet} \mid h.t \mid n$

Constructors $C_0 ::= \dots$ $C_1 ::= \dots$

Operators

$op ::= \text{ref}_T \mid (=)_T \mid (<) \mid (\leq) \mid (>) \mid (\geq) \mid \text{mod} \mid \text{land} \mid \text{lor} \mid \text{lxor} \mid \text{lsl} \mid \text{lsr} \mid \text{asr} \mid (+) \mid (-) \mid (*) \mid (/) \mid - \mid (@_T) \mid (\wedge) \mid \text{create_thread}_T \mid \text{self} \mid \text{kill} \mid \text{create_mutex} \mid \text{lock} \mid \text{try_lock} \mid \text{unlock} \mid \text{create_cvar} \mid \text{wait} \mid \text{signal} \mid \text{broadcast} \mid \text{exit}_T \mid \text{compare_name}_T \mid \text{thunkify} \mid \text{unthunkify}$

Expressions

$e ::= C_0 \mid C_1 e \mid e_1 :: e_2 \mid (e_1, \dots, e_n) \mid \text{function } mtch \mid \text{fun } mtch \mid l \mid op^n e_1 \dots e_n \mid x^n e_1 \dots e_n \mid x \mid M_M.x \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \text{while } e_1 \text{ do } e_2 \text{ done} \mid e_1 \ \&\& \ e_2 \mid e_1 \ || \ e_2 \mid e_1 ; e_2 \mid e_1 e_2 \mid !_T e \mid e_1 :=_T e_2 \mid \text{match } e \text{ with } mtch \mid \text{let } p = e' \text{ in } e'' \mid \text{let } x : T \ p_1..p_n = e' \text{ in } e'' \mid \text{let rec } x : T = \text{function } mtch \text{ in } e \mid \text{let rec } x : T \ p_1..p_n = e' \text{ in } e'' \mid \text{raise } e \mid \text{try } e \text{ with } mtch \mid \Lambda t \rightarrow e \mid e \ T \mid \{T, e\} \text{ as } T' \mid \text{let } \{t, x\} = e_1 \text{ in } e_2 \mid \text{marshal } e_1 \ e_2 : T \mid \text{unmarshal } e \text{ as } T \mid \text{fresh}_T \mid \text{cfresh}_T \mid \text{hash}(X.x)_T \mid \text{hash}(T, e_2)_{T'} \mid \text{hash}(T, e_2, e_1)_{T'} \mid \text{swap } e_1 \text{ and } e_2 \text{ in } e_3 \mid e_1 \text{ freshfor } e_2 \mid \text{support}_T e \mid M_M@x \mid \text{name_of_tie } e \mid \text{val_of_tie } e \mid \text{namecase } e_1 \text{ with } \{t, (x_1, x_2)\} \text{ when } x_1 = e \rightarrow e_2 \text{ otherwise } \rightarrow e_3 \mid e_1 \ || \ e_2 \mid n_T \mid h.x \mid e_1 :='_T e_2 \mid \text{marshalz } \underline{s} \ e : T \mid \text{RET}_T \mid \text{SLOWRET}_T \mid \text{TERM} \mid \text{op}(op^n)^n \ e_1 \dots e_n \mid \text{op}(x^n)^n \ e_1 \dots e_n \mid [e]_{eqs}^T \mid \text{resolve}(M_M.x, M'_{M'}, \text{resolvespec}) \mid \text{resolve_blocked}(M_M.x, M'_{M'}, \text{resolvespec})$

Matches and Patterns

$mtch ::= p \rightarrow e \mid (p \rightarrow e \mid mtch)$
 $p ::= (- : T) \mid (x : T) \mid C_0 \mid C_1 p \mid p_1 :: p_2 \mid (p_1, \dots, p_n) \mid (p : T)$

Signatures and Structures

$sig ::= \text{empty} \mid \text{val } x_x : T \text{ sig} \mid \text{type } t_t : K \text{ sig}$ $Sig ::= \text{sig } sig \text{ end}$
 $str ::= \text{empty} \mid \text{let } x_x : T \ p_1..p_n = e \text{ str} \mid \text{type } t_t = T \text{ str}$ $Str ::= \text{struct } str \text{ end}$

Version and version constraint expressions

$avne ::= \underline{n} \mid \underline{N} \mid h \mid \text{myname}$ $avce ::= ahvce \mid \underline{n}$
 $vne ::= avne \mid avne.vne$ $dvce ::= avce \mid \underline{n-\underline{n}'} \mid \underline{-\underline{n}} \mid \underline{-} \mid * \mid avce.dvce$
 $ahvce ::= \underline{N} \mid h \mid M_M$ $vce ::= dvce \mid \text{name} = ahvce$

Source definitions and Compilation Units

```

sourcedefinition ::= module mode  $M_M : Sig$  version vne = Str withspec
                    import mode  $M_M : Sig$  version vc likespec by resolvespec = Mo
                    mark MK
                    module  $M_M : Sig = M'_M$ 

mode ::= hash | cfresh | fresh | hash! | cfresh!
withspec ::= empty | with !eqs
likespec ::= empty | like  $M_M$  | like Str
resolvespec ::= empty | STATIC_LINK, resolvespec | HERE_ALREADY, resolvespec | URI, resolvespec
Mo ::=  $M_M$  | UNLINKED

compilationunit ::= empty | e | sourcedefinition ;; compilationunit |
                    includesource sourcefilename ;; compilationunit |
                    includecompiled compiledfilename ;; compilationunit

```

Compiled Definitions and Compiled Units

```

definition ::= cmodule... | cimport... | module fresh... | import fresh... | mark MK
compiledunit ::= empty | e | definition ;; compiledunit

```

Marshaled value contents (marshalled values are strings that unmarshal to these)

```

mv ::= marshalled( $E_n$ ,  $E_s$ ,  $s$ , definitions, e,  $T$ )

```

Module names (hashes and abstract names)

```

h ::= hash(hmoduleeqs  $M : Sig_0$  version vne = Str) | hash(himport  $M : Sig_0$  version vc like Str) | n
X ::=  $M_M$  | h

```

Expression name values

```

n ::=  $n_T$  | hash(h.x) $T$  | hash( $T'$ ,  $\underline{s}$ ) $T$  | hash( $T'$ ,  $\underline{s}$ , n) $T$ 

```

(In the implementation all *h* and **n** forms can be represented by a long bitstring taken from \mathbb{H} , ranged over by \underline{N} .)

Type equation sets (the M_M forms occur in the source language)

```

eqs ::=  $\emptyset$  | eqs,  $X.t \approx T$ 

```

Type Environments (for identifiers and store locations — not required at run-time in the implementation)

```

E ::= empty |  $E, x : T$  |  $E, l : T$  ref |  $E, t : K$  |  $E, M_M : Sig$ 

```

Type Environments (for global names — not required in the implementation)

```

 $E_n$  ::= empty |  $E_n, n : nmodule_{eqs} M : Sig_0$  version vne = Str |  $E_n, n : nimport M : Sig_0$  version vc like Str |
         $E_n, n : TYPE$  |  $E_n, n : T$  name

```

Processes

```

P ::=  $0$  | ( $P_1$  |  $P_2$ ) | n : definitions e | n : MX( $\underline{b}$ ) | n : CV

```

Single-Machine Configurations

```

config ::=  $E_n$  ;  $\langle E_s, s, definitions, P \rangle$ 

```

References

- [Ali03] The Alice project, 2003. <http://www.ps.uni-sb.de/alice/>.
- [AVWW96] J. Armstrong, R. Virding, C. Wikstrom, and M. Williams. *Concurrent Programming in Erlang*. Prentice Hall, 1996. 2nd ed.
- [BCF02] N. Benton, L. Cardelli, and C. Fournet. Modern concurrency abstractions for C[#]. In *Proc. ECOOP, LNCS 2374*, 2002.
- [BHS⁺03] G. Bierman, M. Hicks, P. Sewell, G. Stoye, and K. Wansbrough. Dynamic rebinding for marshalling and update, with destruct-time λ . In *Proc. ICFP*, 2003.
- [Bou03] Gérard Boudol. ULM: A core programming model for global computing. Draft, 2003.
- [Car95] L. Cardelli. A language with distributed scope. In *Proc. 22nd POPL*, pages 286–297, 1995.
- [DEW99] S. Drossopoulou, S. Eisenbach, and D. Wragg. A fragment calculus towards a model of separate compilation, linking and binary compatibility. In *Proc. LICS*, pages 147–156, 1999.
- [dot03] Packacking and deploying .net framework applications (.net framework tutorials), 2003. <http://msdn/microsoft.com/library/default.asp?url=/library/en-us/dnanchor/html/netdevanchor.asp>.
- [Fes01] Fabrice Le Fessant. Detecting distributed cycles of garbage in large-scale systems. In *Proc. Principles of Distributed Computing(PODC)*, 2001.
- [FGL⁺96] C. Fournet, G. Gonthier, J.-J. Lévy, L. Maranget, and D. Rémy. A calculus of mobile agents. In *Proc. 7th CONCUR, LNCS 1119*, 1996.
- [GMZ00] D. Grossman, G. Morrisett, and S. Zdancewic. Syntactic type abstraction. *ACM TOPLAS*, 22(6):1037–1080, 2000.
- [HL94] R. Harper and M. Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *Proc. 21st POPL*, 1994.
- [HP] R. Harper and B. C. Pierce. Design issues in advanced module systems. Chapter in *Advanced Topics in Types and Programming Languages*, B. C. Pierce, editor. To appear.
- [HP99] Haruo Hosoya and Benjamin C. Pierce. How good is local type inference? Technical Report MS-CIS-99-17, University of Pennsylvania, June 1999.
- [HRY04] M. Hennessy, J. Rathke, and N. Yoshida. Safedpi: A language for controlling mobile code. In *Proc. FOSSACS, LNCS 2987*, 2004.
- [HS00] R. Harper and C. Stone. A type-theoretic interpretation of standard ML. In *Proof, Language and Interaction: Essays in Honour of Robin Milner*. 2000.
- [JoC] JoCaml. <http://pauillac.inria.fr/jocaml/>.
- [LBR03] Didier Le Botlan and Didier Rémy. MLF: Raising ML to the power of System-F. In *Proceedings of the International Conference on Functional Programming (ICFP 2003), Uppsala, Sweden*, pages 27–38. ACM Press, aug 2003.
- [Ler94] X. Leroy. Manifest types, modules, and separate compilation. In *Proc. 21st POPL*, 1994.
- [LPSW03] J. J. Leifer, G. Peskine, P. Sewell, and K. Wansbrough. Global abstraction-safe marshalling with hash types. In *Proc. 8th ICFP*, 2003.
- [MCHP04] T. Murphy, K. Crary, R. Harper, and F. Pfenning. A symmetric modal lambda calculus for distributed computing. In *Proc. LICS*, 2004.

- [OK93] Atsushi Ohori and Kazuhiko Kato. Semantics for communication primitives in a polymorphic language. In *Proc. POPL*, pages 99–112, 1993.
- [OZZ01] Martin Odersky, Christoph Zenger, and Matthias Zenger. Colored local type inference. *ACM SIGPLAN Notices*, 36(3):41–53, March 2001.
- [PT98] Benjamin C. Pierce and David N. Turner. Local type inference. 1998. Full version in *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 22(1), January 2000, pp. 1–44.
- [PT00] B. C. Pierce and D. N. Turner. Pict: A programming language based on the pi-calculus. In *Proof, Language and Interaction: Essays in Honour of Robin Milner*. 2000.
- [Rém02] Didier Rémy. Using, understanding, and unraveling the ocaml language. In Gilles Barthe, editor, *Applied Semantics. Advanced Lectures. LNCS 2395*, pages 413–537. 2002.
- [Rep99] J. H. Reppy. *Concurrent Programming in ML*. Cambridge Univ Press, 1999.
- [Ros03] A. Rossberg. Generativity and dynamic opacity for abstract types. In *Proc. 5th PPDP*, August 2003.
- [Sew00] Peter Sewell. Applied π – a brief tutorial. Technical Report 498, Computer Laboratory, University of Cambridge, August 2000. An extract appeared as Chapter 9, *Formal Methods for Distributed Processing, A Survey of Object Oriented Approaches*.
- [Sew01] P. Sewell. Modules, abstract types, and distributed versioning. In *Proc. 28th POPL*, 2001.
- [SLW⁺] Peter Sewell, James J. Leifer, Keith Wansbrough, Mair Allen-Williams, Francesco Zappa Nardelli, Pierre Habouzit, and Viktor Vafeiadis. Acute: high-level programming language design for distributed computation. Draft available <http://www.cl.cam.ac.uk/users/pes20/acute>.
- [SPG03] M. R. Shinwell, A. M. Pitts, and M. J. Gabbay. FreshML: Programming with binders made simple. In *Proc. 8th ICFP*, pages 263–274, 2003.
- [SWP99] P. Sewell, P. T. Wojciechowski, and B. C. Pierce. Location-independent communication for mobile agents: a two-level architecture. In *Internet Programming Languages, LNCS 1686*, pages 1–31, 1999.
- [SY97] T. Sekiguchi and A. Yonezawa. A calculus with code mobility. In *Proc. 2nd FMOODS*, pages 21–36, 1997.
- [TLK96] Bent Thomsen, Lone Leth, and Tsung-Min Kuo. A Facile tutorial. In *CONCUR'96, LNCS 1119*, 1996.
- [US01] A. Unyapoth and P. Sewell. Nomadic Pict: Correct communication infrastructure for mobile computation. In *Proc. POPL*, pages 116–127, January 2001.
- [ves] Vesta. <http://www.vestasys.org/>.
- [Wei02] Stephanie Weirich. *Programming With Types*. PhD thesis, Cornell University, August 2002.

Index

abstractin, 93
atomic internal blocked form, 140
atomic blocked form, 140

blocked, 140

canonical type, 44
compile, 45, 116
compiledform, 74

desugar, 109
dom(E), 75
dom(eqs), 77

E₀, 118
E₁, 118
envenv, 150
eqs_of_sign_str, 105
evalcfresh, 118

fast call, 141
filter, 122
flattenclos, 150
fmv, 133
fn, 138
fns, 138

HASH, 76
hashifyCN, 118

internally blocked, 140

limitdom, 93
linkok, 106
locs, 133

makeimports, 134
matchsub, 129
matchty, 110
maybe_cons_bs, 149

namepart, 75

patty, 110

raw_unmarshal, 135
revbc, 138
 ρ^C , 118

selfifysig, 93
sourceinternalform, 74
SSwap, 138

sugaredsourceinternalform, 74
syntacticsubsig, 106

typeflattensig, 119
typeflattenstruct, 119
typeof, 76