

# Sharing in Typed Module Assembly Language

Dominic Duggan

Department of Computer Science  
Stevens Institute of Technology  
Hoboken, NJ 07030.  
`dduggan@cs.stevens-tech.edu`

**Abstract.** There is a growing need to provide low-overhead software-based protection mechanisms to protect against malicious or untrusted code. Type-based approaches such as proof-carrying code and typed assembly language provide this protection by relying on untrusted compilers to certify the safety properties of machine language programs. Typed Module Assembly Language (TMAL) is an extension of typed assembly language with support for the type-safe manipulation of dynamically linked libraries. A particularly important aspect of TMAL is its support for shared libraries.

## 1 Introduction

Protection of programs from other programs is an old and venerable problem, given new urgency with the growing use of applets, plug-ins, shareware software programs and ActiveX controls (and just plain buggy commercial code). Historically the conventional approach to providing this protection has been based on hardware support for isolating the address spaces of different running programs, from each other and from the operating system. The OS kernel and its data structures sit in a protected region of memory, and machine instructions are provided to “trap” into the kernel in a safe way to execute kernel code.

While this approach to protection is widely popularized by operating systems such as Windows 2000 and Linux, there is a growing desire to find alternatives. The problem is that this technique is a fairly heavyweight mechanism for providing protection, relying on expensive context switching between modes and between address spaces. Although application designers have learned to program around this expensive context switching (for example, buffering I/O in application space), this approach breaks down very quickly in software systems composed of separately authored subsystems that do not place much trust in each other, and where context switches may occur much more frequently than in an OS/application scenario [36].

In the OS research community, investigation of alternatives has been motivated by the demands of modular micro-kernel operating systems, where OS modules outside the kernel might not be trusted. Software fault isolation (where the loader inserts software sandboxing checks into machine code [35]) and the SPIN project (where type-safe OS modules are compiled by a trusted compiler

[5]) are examples of approaches to providing protection in software rather than hardware. Sandboxing in Java VMs has also been motivated by the expense of hardware-based memory protection for applets [36]. The commercial world is seeing an explosion in the use of component technology, exemplified by Java Beans and ActiveX controls. This use of component technology again motivates the need to find some relatively lightweight software-based approach to protection in running programs.

Proof-carrying code [30,29] and typed assembly language [28,27] are approaches to providing this protection at low run-time cost. These approaches are examples of *self-certifying code*. A compiler produces a certificate that a program satisfies some property of interest, for example, that the program is well-typed. The user of a compiled program can check that the certificate supplied with a program is valid for that program. If the check succeeds, the program can be run without run-time checks for the safety properties verified by the certificate. This approach has the advantage of moving the compiler out of the trusted computing base, while reducing the need for run-time checks in the code.

Typed assembly language (TAL) enforces a type discipline at the assembly language level, ensuring that malicious or carelessly written components cannot use “tricks” such as buffer overflows or pointer arithmetic to corrupt the data structures in a running program. Unlike the typed machine language underlying the JVM, TAL is not tied to a particular language’s type system or interpreter architecture. Instead the type system is a moderately generic high-level type system with procedures, records and parametric polymorphism, while the target assembly language is any modern RISC or CISC assembly language. The type system is designed to be rich enough that it can serve as a target for compilers for many different languages, while at the same time having as much freedom as possible in its choice of code optimizations, parameter-passing conventions, and data and environment representations [10].

Given the importance of component technology as a motivating factor for TAL, it is clear that there should be support for manipulating components in a type-safe but flexible manner. Modular Typed Assembly Language (MTAL) extends TAL to typed object files and type-safe linking [14]. However this is limited by the assumption that all of a program is linked together before the program is run, with linking happening outside of the program itself. Dynamic linking may be used to avoid loading an entire library when only a small part of the library will be needed. For example, the Linux kernel uses dynamic linking to load in kernel modules on an as-needed basis. While static linkers do a good job of only linking those parts of a library that a program references, they cannot predict in advance what of the referenced modules a program might actually use. Dynamic linking is also useful for shared libraries, allowing several processes to share a commonly used library in memory. Indeed one can consider the operating system itself as a shared library, one that is made available in a protected region of memory to all running programs.

Our interest is in extending TAL with support for dynamic linking and shared libraries. Glew and Morrisett [14] consider some alternative approaches to ex-

tending MTAL with dynamic linking<sup>1</sup>, but this consideration is only informal. One issue that they do not consider, which is central to our work, is what model dynamic linking should use for software components and for linking components.

An obvious candidate is the ML module system [26], which provides fairly sophisticated support for type-safe linking as a language construct [23,6]. Indeed this is the philosophy underlying MTAL, which relies on a phase-splitting transformation to translate ML modules to TAL object files. However the problem with this approach is that it leads to two different models of linking:

1. At the source language level, linking is based on applying parameterized modules. Higher-order parameterized modules may be useful for separate compilation [16,15,18,24,19], but there are still problems with supporting recursive modules [8] (as are found in Java and C).
2. At the assembly language level, linking is based on a type-safe version of the Unix `ld` command. Circular imports present no problem at this level, but much of the sophistication of the type system for ML modules is lost. This is unfortunate, since there are many lessons to be learned from ML that could fruitfully be applied to develop rich linking operations for languages such as Java.

This article describes Typed Module Assembly Language (TMAL), an extension of TAL with run-time support for loading, linking and running modules. Work on dynamic linking has focussed on class loading in the Java virtual machine [21]. Java has the problem of a weak MIL. On the other hand, ML has a powerful MIL but no support for dynamic linking. The current work was originally motivated by the desire to bridge this gap. TMAL pursues a model of linking that is closer to the MTAL approach than the ML approach, because it is closer to the form of linking used by popular languages such as Java. TMAL enriches the MTAL approach in several ways, drawing lessons from the ML experience, but also limiting the ML approach in some ways that are not limiting for Java applications, but do avoid problems with extending ML modules to support Java.

We make the following contributions to TAL:

1. We enrich TAL with coercive interface matching, which allows a module to be coerced to an expected type that makes some fields of the module “private.” This is present in for example the ML module system, but not in MTAL. On the other hand, ML does not provide the same linking primitives as MTAL.
2. We enrich TAL with support for shared libraries. This is supported in the ML module language but not in MTAL. On the other hand, ML does not

---

<sup>1</sup> Glew and Morrisett refer to “dynamic linking” as the process of linking an executable with libraries when it is first invoked, while they refer to “dynamic loading” as the linking and loading of libraries at an arbitrary point during execution. Our use of the generic term *dynamic linking* is meant in the latter sense. We provide separate operations for “loading” a module (reflecting it from the core language to the module language) and for “linking” (linking together two modules).

support recursive modules, which are present in MTAL and which complicate the definition of shared libraries.

3. We extend TAL with primitives for type-safe dynamic linking. Our approach resolves some open problems with dynamic linking and abstract data types. In particular, because types exported by modules are named but opaque (wholly or partially), it is not possible for run-time type checking to discern the underlying representation type for an abstract type.

TMAL arises out of work on a high-level module language, incorporating ideas from ML but with application to languages such as Java, including support for recursive DLLs and shared libraries [12]. It can be viewed as a demonstration of how the semantics of that module language can be incorporated into typed assembly language. A central aspect of this scheme is the proper treatment of shared libraries, an important issue that is addressed in the ML module language but not in more low-level typed module languages [7,14]. A related issue is a *phase distinction* in module languages<sup>2</sup>, between the link-time phase of a module and the run-time phase of a module. The link-time phase is characterized by the application of linking operations to combine a library with other libraries. The run-time phase is initiated by the execution of the initialization code for a library, during or after which the definitions in the library are available to a running client. In static linking the client is always another software component with which the library is linked. With dynamic linking, the client is the running program that loads and initializes the library. This issue is not often explicitly acknowledged in the literature. In TMAL it is recognized by an explicit initialization operation, `dlopen`, that provides the demarcation point between these two phases in the lifetime of a module.

In Sect. 2 we give a brief review of TAL and MTAL. In Sect. 3 we reconsider the approach used in MTAL to represent abstract types that are exported by typed object files, and in particular how type equality and type definitions are handled. In Sect. 4 we give an overview of TMAL. The next four sections describe the operations of TMAL in more detail. In Sect. 5 we describe TMAL's support for coercive interface matching. In Sect. 6 we describe how types and values can be dynamically obtained from a module in TMAL. In Sect. 7 we describe how shared libraries can be constructed in TMAL. In Sect. 8 we describe how DLLs are loaded in a type-safe manner in TMAL. Finally Sect. 10 provides our conclusions.

For reasons of space, we are unable to provide a comprehensive discussion of the various issues in module type systems that motivate some of the design choices presented here. The reader desirous of more contextual discussion than that presented here, is invited to consult [12]. The TMAL type system is based on the linking calculus  $\lambda^{\text{link}}$ , itself intended as a compilation target language

---

<sup>2</sup> This should not be confused with the phase distinction between compile-time and run-time explicated by Harper et al [16]. The phase distinction between link-time and run-time does not exist in the latter calculus, because it translates module-level linking (functor application) to core language operations (function application and generic instantiation).

for a more high-level module language  $\lambda^{\text{mod}}$ , described in [12]. The correctness properties of  $\lambda^{\text{link}}$  carry over to TMAL. There may be some interest in verifying a translation from  $\lambda^{\text{link}}$  to TMAL.

## 2 Modular Typed Assembly Language

In this section, we review Typed Assembly Language (TAL) and Modular Typed Assembly Language (MTAL). This review is largely based on descriptions in the literature [28,10,14]. The syntax of MTAL is given in Fig. 1. Typed Assembly Language can be explained as the result of carrying types in a high-level language through the compilation process, all the way to the final output of assembly language in the backend. Starting with a high-level language, say with procedures

$K \in \text{Kind} ::= ty \mid (K_1 \rightarrow K_2)$ $A, B \in \text{Type Cons} ::= t \mid \text{int} \mid \forall[t_1, \dots, t_m] \Gamma \mid \langle A_1^{j_1}, \dots, A_k^{j_k} \rangle$ $j \in \text{Initialization Flag} ::= 0 \mid 1$ $\Phi \in \text{Type Heap Interface} ::= \{t_1 : K_1, \dots, t_k : K_k\}$ $\Psi \in \text{Value Heap Interface} ::= \{x_1 : A_1, \dots, x_k : A_k\}$ $\Gamma \in \text{Register File Type} ::= \{r_1 : A_1, \dots, r_k : A_k\}$ $\Delta \in \text{Type Var Context} ::= \{t_1 : K_1, \dots, t_k : K_k\}$ $h \in \text{Heap Value} ::= \text{code}[t_1, \dots, t_m] \Gamma.I \mid \langle w_1, \dots, w_k \rangle$ $r \in \text{Register Name} ::= \mathbf{r0}, \mathbf{r1}, \dots$ $w \in \text{Word Value} ::= n \mid x \mid w[A_1, \dots, A_k] \mid \dots$ $v \in \text{Small Value} ::= w \mid r$ $TH \in \text{Type Heap} ::= \{t_1 \mapsto A_1, \dots, t_k \mapsto A_k\}$ $VH \in \text{Value Heap} ::= \{x_1 \mapsto h_1, \dots, x_k \mapsto h_k\}$ $R \in \text{Register File} ::= \{r_1 \mapsto w_1, \dots, r_k \mapsto w_k\}$ $I \in \text{Instruction Sequence} ::= i_1; \dots; i_k$ $i \in \text{Instruction} ::= \text{add } r_1, r_2, v \mid \text{malloc } r[\bar{A}] \mid \text{jmp } v \mid \dots$ $\text{Int} \in \text{Interface} ::= (\Phi, \Psi)$ $O \in \text{Object File} ::= [\text{Int}_I \Rightarrow (TH, VH) : \text{Int}_E]$ $E \in \text{Executable} ::= (TH, VH, x)$ $P \in \text{Program State} ::= (TH, VH, R, I)$
---

Fig. 1. Syntax of MTAL

and records, programs are translated to an assembly language where procedures have been translated to code segments (with code for environment-handling) and records have been translated to heap blocks. Thus for example the procedure definition:

```
int fact (int x) {
    int y = 1;
    while (x != 0) y = (x--) * y;
    return y;
}
```

is translated to the code segment:

```
fact: code[] {a0:int,ra:∇[] {v0:int}}.
    mov    v0,1
    jmp    loop
loop: code[] {a0:int,v0:int,ra:∇[] {v0:int}}.
    bz     a0,ra
    mul   v0,a0,v0
    sub   a0,a0,1
    jmp   loop
```

The register `ra` is the continuation or return address register, pointing to the code to be executed upon return. The `fact` procedure expects an integer in the argument register `a0`, and returns to its caller with an integer in the value return register `v0`. We use MIPS gcc calling conventions to name the registers in examples.

In general heap values  $h$  have the form:

1. A code segment  $\text{code}[t_1, \dots, t_m] \Gamma.I$ , with register file type  $\Gamma = \{r_1 : A_1, \dots, r_n : A_n\}$ . This is a code segment parameterized over  $m$  type variables  $t_1, \dots, t_m$  and expecting its  $n$  arguments in the argument registers  $r_1, \dots, r_n$ . The types of the values in the argument registers are specified in the register file type.  $I$  is the sequence of assembly instructions for the code sequence. This segment has the code type  $\forall[t_1, \dots, t_m] \Gamma$ .
2. A heap block  $\langle w_1, \dots, w_k \rangle$  where the  $k$  values  $w_1, \dots, w_k$  are word values. Such a heap block has a heap block type  $\langle A_1^{j_1}, \dots, A_k^{j_k} \rangle$ , where each  $j_h \in \{0, 1\}$  indicates if the  $h$ th slot has been initialized. Note that the tuple type  $\langle A_1^{j_1}, \dots, A_k^{j_k} \rangle$  should not be confused with tuples of types; we do not therefore have tuple kinds, although they could be added straightforwardly.

Parametric polymorphism is used in an essential way to abstract over the call stack in typing a procedure definition. For example the most general definition of `fact` is:

```
fact: code[EnvT] {a0:int,sp:EnvT,ra:∇[] {v0:int,sp:EnvT}}. ...
```

where the `sp` register points to the environment of the calling procedure. The type parameter `EnvT` ensures that the continuation is passed the calling procedure's environment upon return.

An operational semantics is specified using program states of the form  $(VH, R, I)$ , where

1.  $VH = \{\bar{x} \mapsto \bar{h}\}$  is a value heap, a mapping from labels to heap values  $h$ ;
2.  $R = \{\bar{r} \mapsto \bar{h}\}$  is a register file, a mapping from register names to values; and
3.  $I$  is a sequence of typed assembly instructions.

Program states are typed using register file types  $\Gamma = \{\bar{r} : \bar{A}\}$  and heap types  $\Psi = \{\bar{x} : \bar{A}\}$ , where the latter maps from labels to types. The heap contents are unordered and may contain circular references.

Modular TAL (MTAL) extends these concepts to object files for independent compilation and type-safe linking. An untyped object file imports some values and exports some values, identified by labels pointing into the object file heap. A MTAL object file places types on the imported and exported labels. Furthermore, to support the exportation of abstract data types, an MTAL object file imports and exports types and type operators, identified by labels pointing into a type heap in the object file. An object file  $O$  in MTAL has the form

$$[(\Phi_I, \Psi_I) \Rightarrow (TH, VH) : (\Phi_E, \Psi_E)]$$

where  $\Phi_I$  and  $\Phi_E$  are type interfaces mapping labels to kinds,  $\Psi_I$  and  $\Psi_E$  are value interfaces mapping labels to types,  $TH = \{\bar{t} \mapsto \bar{A}\}$  is a type heap mapping labels to type and type operator definitions and  $VH = \{\bar{x} \mapsto \bar{h}\}$  is a value heap mapping labels to initial values.  $\Phi_I$  and  $\Psi_I$  provide the interfaces for imported types and values, while  $\Phi_E$  and  $\Psi_E$  provide the interfaces for exported types and values. An interface is a pair  $Int = (\Phi, \Psi)$  of type and value heap interfaces.

There are three operations in the MTAL module language:

1. Linking:  $O_1 \text{ link } O_2 \rightsquigarrow O$  combines the object files  $O_1$  and  $O_2$  into the single object file  $O$ . Imports in  $O_1$  and  $O_2$  may be resolved during linking. Interface checking ensures that resolved imports have the correct type.
2. Executable formation:  $(O, x) \xrightarrow{\text{ptg}} E$  identifies the label for executing the code of the object file. Type-checking ensures that this label is bound in the value heap, and that all imports have been resolved.
3. Execution of an executable:  $E \xrightarrow{\text{exec}} P$  produces a program state of the operational semantics from an executable. Program states are extended to include a type heap, and have the form  $(TH, VH, R, I)$ .

### 3 Type Heap Reconsidered

Before giving a description of TMAL, it is useful to explain how our treatment of the type heap and type identity differs from that of MTAL. In MTAL there are two views of a type:

1. Within an object file, a type exported by that object file is completely transparent. The definition of a type label is given by its binding in the type heap,  $TH$ . Because the type heap may contain circular bindings, there are word

value operations `unroll`( $w$ ) and `roll` <sup>$t$</sup> ( $w$ ) that unfold and fold the definition of a type label in the type of  $w$ , respectively. For example if a file system module defines a file abstract type  $t$  as  $\langle \text{int}^t, \text{int}^t \rangle$ , and  $w$  is a word value with this type, then `roll` <sup>$t$</sup> ( $w$ ) gives a word value with type  $t$ , that is with the concrete type folded to the defined type. This means that all types defined in object files are datatypes. In other words, there is no equality theory for implicitly ununfolding the definitions of type identifiers exported by object files, so type equivalence for such type identifiers is based on name equivalence rather than structural equivalence.

2. Outside of an object file, a type exported by that object file may be transparent or opaque. The interface only provides the kind, and the type heap is only visible within the module. Hicks et al [9] use a module system similar to MTAL, except that they also allow an object file to export some of its type definitions, so types may be made transparent to clients.

The advantage of requiring all defined types to be datatypes is that recursive types are assured to be iso-recursive types<sup>3</sup>, thus greatly simplifying the problem of type-checking. The problem with this approach is that it does not adequately handle type sharing for shared libraries. This is explained in more detail in [12]. Consider for example the following Objective ML code [31]:

```
module type S = sig type t; val x:t end
module S1 : S = struct type t = C; val x = C end
module S2 : (S where type t = S1.t) = S1
if true then S1.x else S2.x
```

The module `S1` defines a datatype `t` with single constructor `C`, and binds the field `x` to this constructor. The last conditional type-checks because `S2.x` has type `S2.t`, and the type of `S2` includes the constraint `t=S1.t`, which is also the type of `S1.x`. The structure `S1` is an example of a *shared library*, in the sense that the identity of its (abstract) type component `S1.t` is shared with `S2.t`. The datatype restriction, on the other hand, requires the insertion of marshalling and unmarshalling code at the interface of a shared library, severely curtailing its usability. An example is provided in [12].

It is informally mentioned in the description of MTAL that the implementation includes singleton kinds to expose type definitions to clients of object files. However this is not formalized in the type system and therefore several important issues are left unresolved. For example it is not hard, using singleton kinds, to define two mutually recursive types in separate object files, and linking those files then results in equ-recursive types. This problem can be avoided by only allowing singleton kinds to contain type labels, where the definitions remain encapsulated in the type heap in the object file. In terms of the type

<sup>3</sup> Harper, Crary and Puri [8] make the distinction between *iso-recursive* and *equ-recursive* types. The latter require an equality theory for types that includes a rule for implicitly unrolling a recursive type. The former do not require this equality, and instead rely on operations in the language for explicitly folding and unfolding recursive types.

system presented here, this amounts to only allowing type sharing constraints in the interface, and not allowing type definitions to be exposed.

In our type system we allow both exposure of type definitions, and type sharing, to be expressed in module interfaces. This is done without allowing equ-recursive types in the type system. This is done by separating these two uses of type information in the interface:

1. Exposure of type definitions is expressed using *box kinds*. Box kinds differ from singleton kinds in the following way: whereas singleton kinds allow implicit equality of a type identifier with the type in its singleton kind, box kinds require explicit coercions in the term language between a type identifier and the type in its box kind.
2. Type sharing is expressed using *type sharing constraints*. The type system includes an equality theory that is merely the congruence closure of an equality between type identifiers defined by a context of type sharing constraints. Since equality is only between identifiers, there is no problem with analysing recursive constraints. This is particularly important when we consider dynamic type-checking of DLLs.

TMAL replaces the  $\text{roll}^t$  and  $\text{unroll}$  operations of MTAL, with operations for constructing and deconstructing values of types with box kind:

	Introduction	Elimination
MTAL Expression	$\text{roll}^t(w)$	$\text{unroll}(w)$
MTAL Side-Condition	$w : A, TH(t) = A$	$w : t, TH(t) = A$
TMAL Expression	$\text{fold}_t(w)$	$\text{unfold}_t(w)$
TMAL Side-Condition	$t : \boxtimes A, w : A$	$t : \boxtimes A, w : t$

Because the TMAL operations are typed independently of the type heap, box kinds can be used to expose type definitions in the interface of an object file. In contrast with singleton kinds, because explicit coercions are required between a type with box kind and the type in its kind, recursive types are guaranteed to be iso-recursive types.

## 4 Typed Module Assembly Language

Fig. 2 provides the syntax of Typed Module Assembly Language. In comparison with MTAL, the major changes in module interfaces are:

1. We enrich kinds with box kinds  $\boxtimes A$ . For simplicity we only consider simple types in this account. Box kinds generalize to type operators with some care [12].
2. We enrich import and export interfaces  $Int$  with a type sharing context  $\Xi$ . This is a set of equality constraints between type identifiers.

$K \in \text{Kind}$	$::= ty \mid \boxtimes A$
$A, B \in \text{Type Cons}$	$::= t \mid int \mid \forall[t_1 : K_1, \dots, t_m : K_m] \Gamma$ $\mid \langle A_1^{j_1}, \dots, A_k^{j_k} \rangle \mid \langle \langle \rangle \rangle \mid OT \mid Int$
$j \in \text{Initialization Tag}$	$::= 0 \mid 1$
$\Phi \in \text{Type Heap Interface}$	$::= \{t_i :: \mathbf{t}_i : K_1, \dots, t_k :: \mathbf{t}_k : K_k\}$
$\Psi \in \text{Value Heap Interface}$	$::= \{x :: \mathbf{x}_1 : A_1, \dots, x :: \mathbf{x}_k : A_k\}$
$\Xi \in \text{Type Sharing Cons}$	$::= \{t_1 \cong t'_1 \in K_1, \dots, t_k \cong t'_k \in K_k\}$
$\Gamma \in \text{Register File Type}$	$::= \{r_1 : A_1, \dots, r_k : A_k\}$
$\Delta \in \text{Type Var Context}$	$::= \{t_1 : K_1, \dots, t_k : K_k\}$
$h \in \text{Heap Value}$	$::= \text{code}[t_1 : K_1, \dots, t_m : K_m] \Gamma. I$ $\mid \langle w_1, \dots, w_k \rangle \mid \langle \langle w, OT \rangle \rangle \mid O \mid ST$
$r, r^m, r^s \in \text{Register Name}$	$::= \mathbf{r0}, \mathbf{r1}, \dots$
$w \in \text{Word Value}$	$::= n \mid x \mid w[A_1, \dots, A_k] \mid \dots$
$v \in \text{Small Value}$	$::= w \mid r$
$TH \in \text{Type Heap}$	$::= \{t_i :: \mathbf{t}_i : K_i \mathcal{B}_i^A, \dots, t_k :: \mathbf{t}_k : K_k \mathcal{B}_k^A\}$
$\mathcal{B}^A \in \text{Type Binding}$	$::= \triangleq A \quad (\text{Type Definition})$ $\mid \cong t \quad (\text{Shared Type Binding})$
$VH \in \text{Value Heap}$	$::= \{x_i :: \mathbf{x}_i : A_i \mathcal{B}_i^e, \dots, x_k :: \mathbf{x}_k : A_k \mathcal{B}_k^e\}$
$\mathcal{B}^e \in \text{Value Binding}$	$::= \triangleq h \quad (\text{Value Definition})$ $\mid \cong x \quad (\text{Shared Value Binding})$
$R \in \text{Register File}$	$::= \{r_1 \mapsto w_1, \dots, r_k \mapsto w_k\}$
$\rho \in \text{Renaming Substitution}$	$::= \{\mathbf{n}_1 \mapsto \mathbf{n}'_1, \dots, \mathbf{n}_k \mapsto \mathbf{n}'_k\}$
$I \in \text{Instruction Sequence}$	$::= i_1; \dots; i_k$
$i \in \text{Instruction}$	$::= \text{add } r_1, r_2, v \mid \text{malloc } r[\overline{A}] \mid \text{jmp } v \mid \dots$
$Int \in \text{Interface}$	$::= (\Phi, \Psi, \Xi)$
$OT \in \text{Object File Type}$	$::= [Int_I \Rightarrow Int_E]$
$O \in \text{Object File}$	$::= [Int_I \Rightarrow (TH, VH) : Int_E]$
$ST \in \text{Symbol Table}$	$::= \{\overline{\mathbf{t}} \mapsto \overline{t}, \overline{\mathbf{x}} \mapsto \overline{y}\}$
$P \in \text{Program State}$	$::= (TH, VH, R, I)$

**Fig. 2.** Syntax of TMAL

Purpose	Instruction	Semantics
Linking, interface matching	<b>dllink</b> $r_1^m, r_2^m, r_3^m$	Link modules
	<b>dlcoerce</b> $r_1^m, r_2^m, OT$	Coerce to interface
	<b>dlrename</b> $r_1^m, r_2^m, \rho$	Rename external labels
Dynamic imports	<b>dlopen</b> $r^s, r^m$	Initialize module
	<b>dlsym.t</b> $[t : K]r_1^s, r_2^s, \mathbf{t}$ <b>dlsym.v</b> $r, r^s, \mathbf{x}$	Import type Import value
Shared definitions	<b>dlsetsym.t</b> $r_1^m, r_2^m, t, \mathbf{t}$	Set shared type
	<b>dlsetsym.v</b> $r_1^m, r_2^m, v, \mathbf{x}$	Set shared value
Dynamic linking	<b>dldynamic</b> $r, v, OT$	Construct DLL
	<b>dlload</b> $r^m, r_1, r_2, OT$	Extract module

**Fig. 3.** Summary of TMAL instructions

3. To support coercive interface matching, we add external labels to type and value heap interfaces. As explained in the next section, this allows some of the fields in a module to be safely made private, whereas allowing private fields in MTAL leads to the possibility of run-time name clashes.

There are two forms of module values in TMAL:

1. Modules or object files  $O \equiv [Int_I \Rightarrow (TH, VH) : Int_E]$ . This defines a type heap  $TH$  and a value heap  $VH$ , that may be linked with other such heaps using the TMAL operations.  $Int_I \equiv (\Phi_I, \Psi_I, \Xi_I)$  is the interface of symbols imported by the module, while  $Int_E \equiv (\Phi_E, \Psi_E, \Xi_E)$  is the interface of symbols exported to clients of the module.
2. Symbol tables  $ST \equiv \{\bar{\mathbf{t}} \mapsto \bar{t}, \bar{\mathbf{x}} \mapsto \bar{x}\}$ . A symbol table arises from the initialization of a module. Initializing a module adds its type and value definitions to the type and value heaps, respectively, of the running program. The symbol table provides mappings from the external labels of the module to the heap addresses of its definitions. TMAL provides operations for dynamically importing these addresses into a running program, using a symbol table to perform a run-time lookup based on external labels.

A type heap definition  $t :: \mathbf{t} : K\mathcal{B}^A$  has one of two forms:

1. A definition of the form  $t :: \mathbf{t} : K \triangleq A$  defines a branded type  $t$  with external name  $\mathbf{t}$  and definition  $A$ . External names are explained in the next section. The most general kind for such a type is  $\boxtimes A$ , revealing the structure of the type definition. This is a subkind of  $ty$ , the kind of simple types that makes type definitions opaque.

2. A definition of the form  $t :: \mathbf{t} : K \cong t'$  defines a shared type  $t$  that is equated to the type  $t'$ . Such a type sharing definition can be exposed in an interface by a type sharing constraint  $t \cong t' \in K$ .

Similarly a value heap definition  $x :: \mathbf{x} : A\mathcal{B}^e$  has one of the two forms  $x :: \mathbf{x} : A \triangleq h$  (analogous to a value heap definition  $x \mapsto h$  in MTAL) or  $x :: \mathbf{x} : A \cong y$  (a value sharing definition). Module initialization transforms a value sharing definition to a value heap definition  $x :: \mathbf{x} : A \triangleq h$  by looking up the definition of  $y$  in the heap. Initialization may detect circular value sharing definitions, which correspond to values with no clearly defined initial values.

Are first-class modules necessary for dynamic linking? In TMAL, modules are manipulated (loaded, coerced and linked) at run-time. This in itself does not necessarily require modules as first-class values, and indeed TMAL is based on a module language where there is a strict separation between module values and simple values [12]. Nevertheless a critical part of the transition from a high-level language to TAL is closure conversion, where environment slots are allocated for local variables in a procedure, and the contents of the register file are saved to the environment on a procedure call. Since some local variables may be bound to module values, it is therefore necessary in TMAL to make modules into first-class values. For example, the kernel language described in [12] includes a `letmod` construct for binding a local module identifier to a module:

$$\text{letmod } s = \text{Mod in Expr}$$

where  $Mod$  is a module language expression and  $Expr$  a core language expression. Closure conversion then requires that an environment slot be allocated for the free module identifier  $s$ , leading to the need for first-class modules.

This potentially has some unpleasant consequences. For example Lillibridge [22] has demonstrated that type-checking is undecidable for a type system with first-class modules. The source of this undecidability is a subtype relation between modules that allows fields to be made private, and allows type definitions to be made opaque. There is no such subtype relation in the core language of TMAL, and therefore no such subtyping for modules. This makes “first-class” modules in TMAL strictly less powerful than general first-class modules. For example with general first-class modules, it is possible for the two arms of a conditional to return modules with different interfaces, by having the result interface contain the intersection of the fields of the two modules. However the weak type system for modules in TMAL is sufficient for the purposes of closure conversion, and avoids the undecidability problems with more general type systems.

Rather than allowing type subsumption for modules, TMAL has a `dlcoerce` instruction for explicitly coercing a module to a required type. This coercion operation requires that the module’s type be a subtype of the required type:

$$\begin{aligned} OT \preceq OT' &\iff OT \equiv [Int_I \Rightarrow Int_E], \quad OT' \equiv [Int'_I \Rightarrow Int'_E], \\ &\quad Int'_I \preceq Int_I \text{ and } Int_E \preceq Int'_E \\ Int \preceq Int' &\iff Int \equiv (\Phi, \Psi, \Xi), \quad Int' \equiv (\Phi', \Psi', \Xi'), \end{aligned}$$

$$\begin{aligned} & \Phi \leq \Phi', \Psi \leq \Psi', \text{ and } \Xi \text{ entails } \Xi' \\ \Phi \leq \Phi' & \iff \Phi \equiv \{\overline{t_k} :: \mathbf{t}_k : \overline{K_k}\}, \Phi' \equiv \{\overline{t_m} :: \mathbf{t}_m : \overline{K_m}\}, k \geq m, \overline{K_m} \leq \overline{K'_m} \\ \Psi \leq \Psi' & \iff \Psi \equiv \{\overline{x_k} :: \mathbf{x}_k : \overline{A_k}\}, \Psi' \equiv \{\overline{x_m} :: \mathbf{x}_m : \overline{A_m}\}, k \geq m, \overline{A_m} = \overline{A'_m} \end{aligned}$$

So interface containment reduces to kind containment (where the only containments are of the form  $\boxtimes A \leq ty$ ) and equality between types. The latter equality relation includes entailment based on sharing constraints in the context. The latter constraints can only relate type identifiers, so the equality relation includes rules for forming the congruence closure of these equalities. Because sharing constraints can only relate type identifiers, it is straightforward to extend the language of types with type operators ( $\lambda$ -abstraction) and  $\beta$ -conversion of types.

The type formation rules for modules (object files) and symbol tables are provided in App. A. These operations are discussed in Sect. 5–8, and formally specified in App. B.

## 5 Coercive Interface Matching

MTAL assumes that all field names are globally defined, and interface matching is based on these global field names. Any “implicit” renaming of an identifier requires it to be rewritten globally. There is no notion (as in our approach) of differentiating between external and internal names, with internal names locally bound, and therefore allowing local renaming of these internal names to avoid name clashes during linking. In the MTAL approach, if two modules have fields with the same name, these names are references to the same global symbol, and any renaming of the symbol must be performed in both modules. As a consequence, if fields of an object file are made private in MTAL, there is no way to rename the private fields in order to avoid name clashes when this object file is linked with other object files.

We want to support run-time linking where a library is loaded from disk into the program address space and linked with other libraries. Type safety requires a run-time type check at some point in this scenario. This type check requires that the labels do not admit implicit renaming (such as alpha-conversion in the lambda-calculus). We do not expect that all labels of the loaded library are known, only those labels specified in the expected interface in the run-time type check. Following the MTAL approach, there is the potential for confusion of labels because some of the “hidden” labels in the loaded library may be the same as labels in the libraries it is linked with.

This is the motivation for generalizing labels in type and value heap interfaces to include external names  $\mathbf{t}$  and  $\mathbf{x}$ . Type and value heap interfaces have the form

$$\Phi = \{\overline{t} :: \mathbf{t} : \overline{K}\} \text{ and } \Psi = \{\overline{x} :: \mathbf{x} : \overline{A}\}$$

The internal names  $t$  and  $x$  represent local (type and value) heap addresses. These names admit implicit renaming or alpha-conversion, corresponding to relocating symbols in a heap. The external names  $\mathbf{t}$  and  $\mathbf{x}$  represent external labels

that allow reference to the internal contents of a heap component of a module from outside. To allow fields of a module to be made private, external type and value names in type and value heaps include the special symbol  $\star$ , the name of a private field. Fields in a module are made private using the `dlcoerce` instruction, that changes the external names of fields made private to  $\star$ . The private external name  $\star$  should never appear in a type or value heap type.

Before the contents of a module can be used by a running program, its heaps must be combined with the program heaps. This combination ensures that the internal labels of the module heaps are distinct from the internal labels in the program heaps.

Following [12], we provide three operations for combining and adapting modules. The choice of these operations is informed by an analogy between module combination and process composition in process algebras such as CCS [25]:

Operation	TMAL	CCS
Linking	<code>dllink</code> $r_1^m, r_2^m, r_3^m$	$(P \mid Q)$
Coercion	<code>dlcoerce</code> $r_1^m, r_2^m, OT$	$(P \setminus x)$
Renaming	<code>dlrename</code> $r_1^m, r_2^m, \rho$	$P[\rho]$

The `dllink` instruction links together two modules, combining the type and value heaps. The modules being linked together are in the source registers  $r_2^m$  and  $r_3^m$ , and the result of linking is left in the destination register  $r_1^m$ . The exports of the resulting module are the union of the exports of the two modules, while the imports are the union of the imports of the linked modules minus any imports that are resolved by linking. To obtain a coherent result, the type rules require that the external labels of the exports of the two linked modules are distinct. To maintain this restriction, the external labels of a module must always be visible in the type of the module. The linking operation also requires that the internal labels of the exports of the modules be distinct. Since internal names are bound within a module, they can be renamed to avoid name clashes when merging the fields of the modules being linked. In a concrete implementation, this renaming is handled straightforwardly by relocating the internal addresses of two object files that are linked together.

The `dlcoerce` instruction is necessary because of the absence of a subsumption rule based on interface containment for modules. This latter subsumption rule is not allowable because of the requirement that the external labels of a module must always be visible in its type. The coercion operation performs a run-time adaptation of a module, removing some of its external labels. The corresponding definitions are no longer visible to external clients of the module, but are still accessible via their internal labels to other definitions within the module. The source module is in register  $r_2^m$ , while the result of coercion is left in the destination register  $r_1^m$ . The type to which the module is coerced is specified by the object file type  $OT$ . This type annotation is mostly only for type-checking, and can be removed before execution. The part of the annotation that must be

preserved during execution is the association between external names and internal names; TMAL includes instructions for looking up a field in an initialized module based on its external name.

The `dlrename` instruction is a second operation for coercive interface matching, and renames some of the external labels in a label. A renaming substitution  $\rho$  is an injective mapping from external labels to external labels. Since external names are used at run-time, this renaming substitution must be applied at run-time.

## 6 Dynamic Imports

The instructions given in the previous section operate on values at the module language level. At the heart of the TMAL approach are the instructions that connect the module language level to the core language level. In the  $\lambda^{\text{mod}}$  module language described in [12], this connection is provided by an `init` operation that initializes a module and introduces its definitions into a local scope in a core language program. In TMAL the `init` operation is realized by three instructions, for initializing a module and for importing its definitions into the scope of a running thread:

Operation	TMAL
Initialize module	<code>dlopen</code> $r^s, r^m$
Import type	<code>dlsym_t</code> $[t : K]r_1^s, r_2^s, t$
Import value	<code>dlsym_v</code> $r, r^s, x$

These operations allow a program to import some of the symbols from a DLL, using the external labels of a DLL to access its definitions.

The `dlopen` instruction expects a pointer to a module in register  $r^m$ . The instruction initializes the module, adds its type and heap bindings to the program heaps, and building a symbol table with mappings from the module fields to their bindings in the program heaps. A pointer to this symbol table is left in register  $r^s$ .

The `dlsym_t` instruction imports a type definition into the local context of the current thread, while the `dlsym_v` instruction imports a value definition. The `dlsym_t` operation imports a type symbol from a DLL into a register, using the external label of the type symbol and the symbol table of the DLL to map to the internal label. Note that the internal label cannot be known statically; the internal label is chosen at the point where the DLL is initialized and its value definitions are added to the program's value heap. This is in contrast with MTAL, where heap locations are referenced by globally bound internal names, and where renaming to avoid name clashes is not possible. In TMAL, the internal label is chosen so that there is no clash with the labels already given to program heap contents. Since the complete contents of the program heap are not known until run-time, there is no way to know the internal label during type-checking.

The `dlsym_t` instruction expects a pointer to a pointer to a symbol table in register  $r_2^s$ , and specifies the external name `t` of the type symbol to be imported in the local thread address space. The instruction binds the type parameter `t` in subsequent instructions to the type heap address corresponding to this external name. The symbol table type must be modified so that references to the global heap address are rebound to this local type parameter, for type-checking subsequent importations of definitions that rely on this type symbol. Register  $r_1^s$  is left with a pointer to a symbol table of this modified type.

The `dlsym_v` operation imports (the heap address of) a value definition from a DLL into a register, using the external label of the value definition and the symbol table of the DLL to map to the internal label. The external name `x` of the definition is specified in the instruction, and the instruction leaves the value heap address of the definition in the value register `r`.

```

// Assume s1 points to loaded file system module
dlopen  s2,s1           // Initialize module
dlsym_t [FileT:ty] s3,s2,File Import file type
dlsym_v s4,s3,open     // Import file open operation
mov     a0,file_name   // Load file name
mov     ra,retpt[FileT] // Load continuation
jmp     s4[EnvT]       // Jump to file open operation

retpt:   code[FileT]{v0:FileT,sp:EnvT} ...
file_name: "/etc/passwd"

```

Fig. 4. Example of dynamic imports

Fig. 4 gives an example of the use of these operations. Assuming the `s1` register points to a module, the `dlopen` instruction initializes that module, adds its type and value heap definitions to those of the running program. The result of initialization is a pointer, in the `s2` register, to a symbol table mapping from the external labels of the module to the addresses of its definitions in the program heaps.

The important proviso in the `dlsym_v` operation is that none of the free type variables in the type of a value definition are bound by the type heap definitions addressed by the symbol table. For example, recalling the example in Fig. 4, assume that the symbol table resulting from initializing the file system module has type:

```

type File::File : ty
val open::open :
  ∀[EnvT:ty]{a0:String,sp:EnvT,ra:∀[] {v0:File},sp:EnvT}

```

The abstract file type `File` occurs free in the type of the `open` operation. Therefore the `dlsym_v` instruction cannot import this definition immediately. The reason is that the register file type resulting from this importation would have no binding for the type identifier `File` in the type of the `v0` register.

In order to import the `open` operation, the type identifier `File` that occurs free in its type must first be imported from the DLL. This is done using the `dlsym_t` operation. In the example in Fig. 4, the `dlsym_t` instruction binds a local type identifier `FileT` to the abstract type `File` defined by the DLL. The `s3` register is bound to a new symbol table with type:

```
val open::open :
  ∀[EnvT:ty] {a0:String, sp:EnvT, ra:∀[] {v0:FileT}, sp:EnvT}
```

The abstract file type in the type of the `open` operation has been relocated to a type bound in the local context of the current thread, therefore it is now possible to import the `open` definition from the DLL.

## 7 Shared Libraries

Heaps in modules may contain shared type bindings  $t :: \mathbf{t} : K \cong t'$  and shared value bindings  $x :: \mathbf{x} : A \cong y$ . If all linking is performed before a program runs, then shared bindings are unnecessary. However shared bindings become crucial in an environment where modules are initialized at run-time.

For example, consider a module implementing a network protocol. This implementation requires some operations and types that are only provided by the operating system. Module linking can be used to combine these modules into a single module implementing the operating system with that protocol:

```
// Assume s1 points to loaded OS module
// Assume s2 points to loaded protocol module
dllink   s3,s1,s2           // Link OS, protocol modules
dlopen   s4,s3             // Initialize module
dlsym_t  [Conn:ty] s5,s4,Conn // Import connection type
dlsym_v  s6,s5,open        // Import conn open operation
```

However there is a difficulty with this approach: the operating system will have already been initialized when the program runs. In fact the operating system is really the first module to be initialized, and a running program is just another module that has been loaded and initialized by code defined in the operating system module.

Similar remarks apply to access to OS operations from a process. The process must somehow have access to labels into the OS type and value heaps<sup>4</sup>, but it is unrealistic to expect a program to be linked with its own copy of the OS module

<sup>4</sup> As mentioned in Sect. 1, approaches such as typed assembly language should be regarded as an alternative to current heavyweight protection mechanisms such as hardware-based memory protection and the use of library stubs to trap to the OS.

before execution can begin. The OS is one example of a shared library, a library that is loaded and initialized once, and that is subsequently available to other libraries as they are loaded.

The following instructions allow a program to construct a shared library:

Operation	TMAL
Set shared type	<code>dlsetsym_t <math>r_1^m, r_2^m, t, \mathbf{t}</math></code>
Set shared value	<code>dlsetsym_v <math>r_1^m, r_2^m, r, \mathbf{x}</math></code>

The `dlsetsym_t` instruction allows a reference to a type to be added to the export list of a module, while `dlsetsym_v` instruction allows a reference to a value to be added. These are not the only way that shared value and heap definitions can be constructed. For example, the initial value heap in a module may contain the definition of another module (manipulated by the parent module at run-time) that has aliases for value and type bindings, where the child module definitions that are shared are bound in the parent module heaps, or are imported or exported by the child module. However the aforesaid instructions are the only way to introduce aliases into a module, for shared bindings that are not available until run-time. For example, they are the only way to add bindings for OS types and operations into a module that requires those OS definitions. Once such a shared library has been constructed, the `dllink` instruction allows it to be linked with other modules.

The `dlsetsym_t` instruction expects a pointer to a module in register  $r_2^m$ , and a pointer into the type heap in type “register”  $t$ . The module should import a type definition with external label  $\mathbf{t}$  and with a kind compatible with  $t$ . The instruction moves the type field with label  $\mathbf{t}$  from the import list of the module to the export list, binding to the field to the type heap pointer given by  $t$ , and the resulting module is given in register  $r_1^m$ .

The `dlsetsym_v` instruction expects a pointer to a module in register  $r_2^m$ , and a heap address in register  $r$ . The field labelled with  $\mathbf{x}$  in the module is moved from the import list to the export list, bound to the value heap pointer in  $r$ , and a pointer to the resulting module is left in register  $r_1^m$ .

Returning to the example above of a protocol module, suppose that this module requires a type `ProtId` of protocol identifiers and an operation `deliver` from the OS. The latter operation is used by this protocol module to deliver a protocol data unit to the next protocol above it in the protocol stack.

```
// Assume s1 points to initialized OS module
// Assume s2 points to loaded protocol module (PM)
dlsym_t      [ProtId:ty] s1,s1,ProtId // Import prot id type
dlsym_v      s3,s1,deliver           // Import deliver operation
dlsetsym_t   s2,s2,ProtId,ProtId    // Export protocol id to PM
dlsetsym_v   s2,s2,s3,deliver        // Export deliver to PM
dlopen       s4,s2                   // Initialize PM
```

Alternatively, if the code for initializing the protocol module is in the OS itself, then this code can be defined as:

```
// Assume s2 points to loaded protocol module (PM)
dlsetsym_t s2,s2,ProtId,ProtId // Export protocol id to PM
dlsetsym_v s2,s2,deliver,deliver // Export deliver to PM
dlopen s4,s2 // Initialize PM
```

where `ProtId` and `deliver` are direct references into the type and value heaps, respectively, in the module implementing the OS.

For example, considering the example above of assigning the `ProtId` and `deliver` fields of a protocol module, assume that the protocol module has type:

```
import type ProtId::ProtId
import val deliver::deliver :  $\forall[\text{EnvT}]\{a0:\text{ProtId}, \dots\}$ 
export type Conn::Conn
export val open::open :  $\forall[\text{EnvT}]\{a0:\text{String}, \dots\}$ 
```

Then setting the `ProtId` field with the `ProtId` type defined in the OS module results in a module with type:

```
export type ProtId1::ProtId
import val deliver::deliver :  $\forall[\text{EnvT}]\{a0:\text{ProtId1}, \dots\}$ 
export type Conn::Conn
export val open::open :  $\forall[\text{EnvT}]\{a0:\text{String}, \dots\}$ 
sharing type ProtId1  $\cong$  ProtId
```

In this case the internal type name `ProtId1` is a renaming of the internal name for the type of protocol identifiers, so as to avoid a name clash with the internal type name `ProdId` in the OS module. If the OS module has a value heap label `deliver` with type

$$\forall[\text{EnvT}]\{a0:\text{ProtId}, \dots\}$$

then the type sharing constraint allows this type to be equated with the type of the `deliver` heap label in the protocol module. This allows the `dlsetsym_v` instruction to be used to assign this value field.

## 8 Dynamic Loading

The final set of instructions are used to attach run-time type information to a DLL. This type information is used in a run-time type check, to ensure that a DLL that is loaded from disk or from the network has the required module type. There is an instruction `dldynamic` for bundling a value with a type description, and another instruction `dllload` for checking that a DLL has a specified type.

Operation	TMAL
Construct DLL	<code>dldynamic <math>r, v, OT</math></code>
Extract module	<code>dllload <math>r^m, r_1, r_2, OT</math></code>

The type expression  $\langle\langle \rangle\rangle$  denotes the type of a DLL. The `dldynamic` instruction associates a type descriptor  $OT$  with the heap address of a module in a DLL value  $\langle\langle w, OT \rangle\rangle$ , of DLL type. To be completely accurate, object files should be stripped of unnecessary type information before run-time. Then the only places where type information is required are (a) the external labels of import and export lists (since these labels are used by various instructions to look up fields in modules and symbol tables, and (b) in the `dldynamic` and `dload` instructions above, and DLL values. We forgoe specifying this type-erasure semantics for lack of space.

The `dload` instruction extracts a module from a DLL. This instruction also requires the value representation of a module type, the type that is expected of the module in the DLL. The instruction performs a run-time interface containment check, and if this succeeds it coerces the module in the DLL to the required type. If the interface check fails, control transfers to the failure continuation in register  $r_2$ .

The interface check includes a check for entailment of type sharing constraints. The simple form of type sharing constraints, only relating type identifiers, and the fact that the bindings in the type heap are opaque, facilitate this entailment check. The fact that type heap bindings are opaque also has the benefit that the dynamic type check cannot violate encapsulation of abstract types; this is explained in more detail in [12].

## 9 Related Work

There has been a great deal of work on the semantics of module interconnection languages, particularly in the context of the ML module system [16,15,18,19,33]. The notion of separating external and internal field names, with the latter allowing renaming to avoid name clashes, originated with Harper and Lillibridge [15]. A related idea is used by Riecke and Stone to allow fields of an object to be made private, and the object then extended with a field with the same external name. Similar notions of internal and external names appear in the module calculi of Ancona and Zucca [4] and Wells and Vestergaard [38].

Cardelli [7] gives a semantics for Unix-style linking in terms of a simple  $\lambda$ -calculus, ensuring that all symbols in a program are resolved before it is executed. Flatt and Felleisen [13] and Glew and Morrisett [14] extend this work to consider typed module contents and circular import dependencies. It is not clear what the type of a module is in these approaches (linking simply resolves imports against exports in a type-safe way). Glew and Morrisett do not support shared libraries (type sharing) or dynamic linking. Flatt and Felleisen allow dynamic linking of units. However the *invoke* operation for initializing a unit returns a single core language value; there is no other way for a program to access the contents of a unit. The *invoke* operation takes as arguments types and values from the running program that can be provided as imports to a library before initialization. So there are really two linking operations with units: the linking operation for merging units and the more limited linking that is implicitly part of

the semantics of initialization. Our approach provides a single linking operation, and addresses the problem of sharing type (and value) identity that is not considered by these other approaches. In our system, the *invoke* operation of units would translate into a sequence of `dlsetsym.t` and `dlsetsym.v` instructions, to build the imports for the unit, followed by a `dlopen` instruction to initialize the unit, followed by a `dlsym.v` instruction to retrieve the single value returned by unit initialization.

Crary et al [8] give an explanation of recursive modules in terms of the structure calculus [16]. Their work is predicated on the assumption that module linking is based on functor instantiation, and phase-splitting allows this to be transformed to core-language function application. As discussed in [12], it is difficult to generalize this model of linking to the kinds of module operations we consider.

Work on dynamic linking in ML has focused on dynamic types [2,20,1,34,11]. With these approaches a dynamic value tags a value with a runtime type tag, of type *Dynamic*. This is similar to our approach to dynamic linking, but extended to modules rather than simple values, as a way of reifying modules into the core language.

A perennial problem with dynamics is that they violate encapsulation, in the sense that the underlying representation type of a value with abstract type can be exposed, by first bundling the value as a dynamic and then using runtime type checks to examine the representation type. This is an artifact of the fact that types are bound at runtime using beta-reduction. As mentioned in Sect. 8, our approach to DLLs avoids this problem, because the bindings in the type heap remain opaque during program execution. A similar approach is possible in the system of Hicks et al [9,17]

Russo [32] considers an approach to adding first-class modules to ML, based on converting module values to core language values and back again. Explicit type annotations for modules ensure there are no unpleasant interactions with type inference. Russo avoids the undecidability of type-checking with first-class modules by omitting type subsumption for modules converted to core language values. This is similar to our approach to ensuring decidability with first-class modules. Our reflective treatment of DLLs is different from Russo's treatment of first-class modules. A module reified into the core language in Russo's approach retains its type, though reified to a core language type. In contrast, our reification operation (for building a DLL) masks the type entirely, and there must then be a reflection operation (with a dynamic type check) that extracts a module from a DLL. Dynamic typing is not necessary with Russo's approach, since his purpose is not to provide DLLs.

Ancona and Zucca [4], building on earlier work in mixin modules [3], provide a primitive calculus of modules that supports circular dependencies. Types are restricted to branded types, that is, types where equivalence is based on name equivalence rather than structural equivalence. They do not consider dynamic linking or shared libraries (and the resulting issues with recursive type constraints).

Wells and Vestergaard [38] present a calculus for equational reasoning about first-class modules. They do not place any restrictions on circular import dependencies (including dependencies between value components), allowing circular definitions that lazily unwind. They verify strong normalization and confluence for their calculus, relying on a lazy reduction semantics. They do not consider typing aspects of their calculus. So for example they do not consider the problem of equ-recursive versus iso-recursive types, and they provide no support for shared libraries. Finally as with Russo’s work there is no consideration of narrowing a DLL to a specific interface, an important practical facility for dynamic linking.

Crary, Hicks and Weirich [9,17] extend TAL with primitive operations for building type-safe DLLs, on top of which more expressive dynamic linking mechanisms can be constructed. For example they are able to provide a type-safe implementation of the Unix dynamic linking API, as well as an implementation of units. Their approach amounts to extending the TAL kernel with type-safe checked casting [37]. Although their approach is type-safe, it is also more low-level than the approach described here, and so some errors that are caught statically in our type system are only caught dynamically by checked casting in their approach. The single type failure point in our calculus is the `dlload` operation, that reflects a DLL from the core language into the module language. The difference is really one of level; their approach could for example be the basis for an implementation of TMAL.

The module type system underlying that of Crary et al is MTAL, and therefore it shares the limitations of MTAL: the absence of coercive interface matching, and the absence of sharing. There are no operations for linking modules together at run-time, rather modules are loaded into a running program and their imports resolved against bindings in the global program heaps. Crary et al allow a module’s contents to be accessed before all of its imports have been resolved, allowing “lazy” resolution. In our approach a continuation can specify (as a module) the definitions it requires, and the continuation argument can be linked with other modules. To ensure that a module is initialized (opened) no more than once, a module cache can be implemented: the first time a module is initialized, a shared library is constructed (using `dlsetsym_t` and `dlsetsym_v`) with the same interface, and this shared library saved in the cache with the same module name used to load the original module. Subsequent searches for this library will find the cached version, and it can be used for example to resolve the imports of subsequent DLLs. In this way a form of “lazy loading” as found in Java class loaders can be implemented on top of our module system.

## 10 Conclusions

We have described Typed Module Assembly Language (TMAL), an extension of typed assembly language with instructions for manipulating modules at run-time. These instructions include support for coercive interface matching, dynamically importing definitions from a library, constructing shared libraries, and

using DLLs in a type-safe manner. A possible application of these mechanisms is in component-based programming environments, as demonstrated by commercial platforms based on COM or Java. The mechanisms described here can be used to enrich such environments with flexible but type-safe operations for interconnecting modules under program control.

It is plausible that this is not the final word on the choice of instruction set for TMAL. Although the instructions for dynamic imports and shared libraries are fairly RISC-y, this is not true of the `dlink`, `dlcoerce` and `dlrename` instructions, nor is it true of the `dlload` and `dldynamic` instructions. We are considering how these instructions could be decomposed into simpler instructions, to weaken the atomicity requirements of the current instruction set.

**Acknowledgements.** Thanks to Michael Hicks and J. Gregory Morrisett for helpful discussions. Thanks to the anonymous reviewers for their excellent feedback, comments and suggestions for improvement.

## References

1. Martin Abadi, Luca Cardeli, Benjamin Pierce, and Didier Remy. Dynamic typing in polymorphic languages. In Peter Lee, editor, *Proceedings of the ACM SIGPLAN Workshop on ML and its Applications*, San Francisco, California, June 1992. Carnegie-Mellon University Technical Report CMU-CS-93-105.
2. Martin Abadi, Luca Cardelli, Benjamin Pierce, and Gordon Plotkin. Dynamic typing in a statically typed language. *ACM Transactions on Programming Languages and Systems*, 13(2):237–268, 1991.
3. David Ancona and Elena Zucca. A theory of mixin modules: Basic and derived operators. *Mathematical Structures in Computer Science*, 8(4):401–446, 1998.
4. David Ancona and Elena Zucca. A primitive calculus for module systems. In *Proceedings of the International Conference on Principles and Practice of Declarative Programming*, Paris, France, September 1999. Springer-Verlag.
5. B. N. Bershad, S. Savage, P. Pardyak, E. G. Sire, M. E. Fiuczynski, D. Becker, C. Chambers, and S. Egger. Extensibility, safety and performance in the SPIN operating system. In *Symposium on Operating Systems Principles*, pages 267–283, Copper Mountain, CO, 1995. ACM Press.
6. Edoardo Biagioni, Robert Harper, Peter Lee, and Brian G. Milnes. Signatures for a network protocol stack: A systems application of standard ML. In *Proceedings of ACM Symposium on Lisp and Functional Programming*, pages 55–64, Orlando, Florida, January 1994. ACM Press.
7. Luca Cardelli. Program fragments, linking and modularization. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 266–277. ACM Press, January 1997.
8. Karl Cray, Robert Harper, and S. Puri. What is a recursive module? In *Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation*, Atlanta, GA, 1999. ACM Press.
9. Karl Cray, Michael Hicks, and Stephanie Weirich. Safe and flexible dynamic linking of native code. In *Workshop on Types in Compilation*, Lecture Notes in Computer Science, Montreal, Quebec, Canada, September 2000. Springer-Verlag.

10. Karl Crary and Greg Morrisett. Type structure for low-level programming languages. In *Proceedings of the International Conference on Automata, Languages and Programming*, Lecture Notes in Computer Science. Springer-Verlag, 1999.
11. Dominic Duggan. Dynamic typing for distributed programming in polymorphic languages. *ACM Transactions on Programming Languages and Systems*, 21(1):11–45, January 1999.
12. Dominic Duggan. Type-safe dynamic linking with recursive DLLs and shared libraries. Technical report, Stevens Institute of Technology, 2000.
13. M. Flatt and M. Felleisen. Units: Cool modules for HOT languages. In *Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation*, 1998.
14. Neal Glew and Greg Morrisett. Type-safe linking and modular assembly languages. In *Proceedings of ACM Symposium on Principles of Programming Languages*, San Antonio, Texas, January 1999. ACM Press.
15. Robert Harper and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 123–137, Portland, Oregon, January 1994. ACM Press.
16. Robert Harper, John Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 341–354. Association for Computing Machinery, 1990.
17. Michael Hicks and Stephanie Weirich. A calculus for dynamic loading. Technical Report MS-CIS-00-07, University of Pennsylvania, 2000.
18. Xavier Leroy. Manifest types, modules, and separate compilation. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 109–122, Portland, Oregon, January 1994. acmp.
19. Xavier Leroy. Applicative functors and fully transparent higher-order modules. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 154–163, San Francisco, California, January 1995. ACM Press.
20. Xavier Leroy and Michel Mauny. Dynamics in ML. *Journal of Functional Programming*, 3(4):431–463, 1993.
21. Sheng Liang and Gilad Bracha. Dynamic class loading in the Java virtual machine. In *Proceedings of ACM Symposium on Object-Oriented Programming: Systems, Languages and Applications*. ACM Press, October 1998.
22. Mark Lillibridge. *Translucent Sums: A Foundation for Higher-Order Module Systems*. PhD thesis, Carnegie-Mellon University, Pittsburgh, PA, May 1997. Technical Report CMU-CS-97-122.
23. David MacQueen. Using dependent types to express modular structure. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 277–286. ACM Press, 1986.
24. David MacQueen and Mads Tofte. A semantics for higher-order functors. In *European Symposium on Programming*, volume 788 of *Lecture Notes in Computer Science*, pages 409–423. Springer-Verlag, 1994.
25. Robin Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
26. Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Revised Definition of Standard ML*. The MIT Press, 1997.
27. Greg Morrisett, Karl Crary, Neal Glew, Dan Grossman, Richard Samuels, Frederick Smith, David Walker, Stephanie Weirich, and Steve Zdancewic. TALx86: A realistic typed assembly language. In *Workshop on Compiler Support for Software Systems (WCSS)*, Atlanta, GA, May 1999.

28. Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From System F to typed assembly language. In *Proceedings of ACM Symposium on Principles of Programming Languages*, 1998.
29. George Necula. Proof-carrying code. In *Proceedings of ACM Symposium on Principles of Programming Languages*, 1997.
30. George Necula and Peter Lee. Safe kernel extensions without run-time checking. In *Operating Systems Design and Implementation*, 1996.
31. Didier Rémy and Jérôme Vouillon. Objective ML: An effective object-oriented extension to ml. *Theory and Practice of Object Systems*, 4(1):27–50, 1998.
32. Claudio Russo. Adding first-class modules to Standard ML. In *European Symposium on Programming*, Berlin, Germany, April 2000. Springer-Verlag.
33. Zhong Shao. Transparent modules with fully syntactic signatures. In *Proceedings of ACM International Conference on Functional Programming*, Paris, France, September 1999.
34. Mark Shields, Tim Sheard, and Simon Peyton-Jones. Dynamic typing as staged type inference. In *Proceedings of ACM Symposium on Principles of Programming Languages*, pages 289–302, San Diego, California, January 1998. ACM Press.
35. Robert Wahbe, Steven Lucco, Thomas E. Anderson, and Susan L. Graham. Efficient software-based fault isolation. In *Symposium on Operating Systems Principles*, pages 203–216. ACM Press, 1993.
36. Dan S. Wallach, Dirk Balfanz, Drew Dean, and Edward W. Felten. Extensible security architectures for Java. In *Symposium on Operating Systems Principles*. ACM Press, 1997.
37. Stephanie Weirich. Type-safe cast (functional pearl). In *Proceedings of ACM International Conference on Functional Programming*, Montreal, Canada, September 2000. ACM Press.
38. Joseph B. Wells and René Vestergaard. Equational reasoning for linking with first-class primitive modules. In *European Symposium on Programming*, Berlin, Germany, April 2000. Springer-Verlag.

## A Type Rules for Modules and Symbol Tables

This appendix summarizes the type rules for modules and symbol tables. The type rules for values and heaps are specified using judgements of the form given in Fig. 5. The contexts of type and value heap bindings are defined by:

$$\begin{aligned}\tilde{\Phi} &= \{(t : K) \mid (t :: \mathbf{t} : K) \in \Phi\} \\ \tilde{\Psi} &= \{(x : A) \mid (x :: \mathbf{x} : A) \in \Psi\}\end{aligned}$$

The type rules for modules (object files) require that the type heap satisfies the exported type heap interface, that the value heap satisfies the exported value heap interface, and that the exported type sharing constraints are entailed by the type sharing implied by the type heap, the type sharing context, and the type sharing constraints imposed on the imports.

$\tilde{\Phi}; \Delta \vdash \diamond$	Type context formation
$\tilde{\Phi}; \Delta \vdash \Phi'$	Type heap interface
$\tilde{\Phi}; \Delta; \Xi \vdash \Phi' = \Phi''$	Type heap interface equality
$\tilde{\Phi}; \Delta; \Xi \vdash \Phi' \leq \Phi''$	Type heap interface containment
$\tilde{\Phi}; \Delta; \Xi \vdash TH : \Phi'$	Type heap
$\tilde{\Phi}; \Delta \vdash \Xi$	Sharing heap interface
$\tilde{\Phi}; \Delta; \Xi \vdash \Xi'$	Entailment of type sharing constraints
$\tilde{\Phi}; \Delta \vdash \Psi$	Value heap interface
$\tilde{\Phi}; \Delta; \Xi \vdash \Psi' = \Psi''$	Value heap interface equality
$\tilde{\Phi}; \Delta; \Xi \vdash \Psi' \leq \Psi''$	Value heap interface containment
$\tilde{\Phi}; \Delta; \Xi; \tilde{\Psi} \vdash VH : \Psi'$	Value heap
$\tilde{\Phi}; \Delta \vdash K$	Kind formation
$\tilde{\Phi}; \Delta; \Xi \vdash K = K'$	Kind equality
$\tilde{\Phi}; \Delta; \Xi \vdash K \leq K'$	Kind containment
$\tilde{\Phi}; \Delta \vdash A : K$	Type formation
$\tilde{\Phi}; \Delta; \Xi \vdash A = B \in K$	Type equality
$\tilde{\Phi}; \Delta; \Xi \vdash [Int_I \Rightarrow Int_E] \preceq [Int'_I \Rightarrow Int'_E]$	Module type containment
$\tilde{\Phi}; \Delta; \Xi; \tilde{\Psi} \vdash h : A$	Type of heap value
$\tilde{\Phi}; \Delta; \Xi; \tilde{\Psi} \vdash w : A$	Type of word value
$\tilde{\Phi}; \Delta \vdash \Gamma$	Register file type
$\tilde{\Phi}; \Delta; \Xi; \tilde{\Psi} \vdash R : \Gamma$	Register file
$\tilde{\Phi}; \tilde{\Psi}; \Xi \vdash \{\Delta; \Gamma\} I \{\Delta'; \Gamma'\}$	Instruction formation

Fig. 5. Judgement Forms of TMAL

$$\begin{array}{c}
\tilde{\Phi}; \{\}; \Xi \vdash [Int_I \Rightarrow Int_E] : ty \quad Int_I = (\Phi_I, \Psi_I, \Xi_I) \quad Int_E = (\Phi_E, \Psi_E, \Xi_E) \\
\tilde{\Phi}' = \tilde{\Phi} \cup \tilde{\Phi}_I \cup TENV(TH) \quad \Xi' = \Xi \cup \Xi_I \cup SHARE(TH) \\
\tilde{\Phi}'; \Xi' \vdash TH : \Xi_E \quad \tilde{\Phi}'; \Xi' \vdash TH : \Phi_E \\
\hline
\tilde{\Phi}; \tilde{\Phi}'; \Delta; \Xi'; \tilde{\Psi} \cup \tilde{\Psi}_I \cup VENV(VH); \Delta \vdash_{\text{val}} VH : \Psi_E \\
\hline
\tilde{\Phi}; \tilde{\Phi}; \Delta; \Xi; \tilde{\Psi}; \Delta \vdash_{\text{val}} [Int_I \Rightarrow (TH, VH) : Int_E] : [Int_I \Rightarrow Int_E] \\
\text{(VAL OBJECT FILE)} \\
\\
\frac{\tilde{\Phi}; \tilde{\Phi}; \Delta; \Xi; \tilde{\Psi}; \Delta \vdash_{\text{val}} e : A \quad \mathbf{x} \neq \star}{\tilde{\Phi}; \tilde{\Phi}; \Delta; \Xi; \tilde{\Psi}; \Delta \vdash_{\text{val}} \{x :: \mathbf{x} : A \triangleq e\} : \{x :: \mathbf{x} : A\}} \\
\text{(HEAP VAL DEF)} \\
\\
\frac{(y : A') \in \tilde{\Psi} \quad \tilde{\Phi}; \Delta; \Xi \vdash A = A' : ty \quad \mathbf{x} \neq \star}{\tilde{\Phi}; \tilde{\Phi}; \Delta; \Xi; \tilde{\Psi}; \Delta \vdash_{\text{val}} \{x :: \mathbf{x} : A \cong y\} : \{x :: \mathbf{x} : A\}} \\
\text{(HEAP VAL SHARE)}
\end{array}$$

$$\frac{\tilde{\Phi}; \{\}; \Xi \vdash A : K \quad \mathbf{t} \neq \star}{\tilde{\Phi}; \Xi \vdash \{type \ t :: \mathbf{t} :: K \triangleq A\} : \{t :: \mathbf{t} : K\}} \quad (\text{HEAP TYPE DEF})$$

$$\frac{t' \in Names(\tilde{\Phi}) \quad \tilde{\Phi}; \{\}; \Xi \vdash t' : K \quad \mathbf{t} \neq \star}{\tilde{\Phi}; \Xi \vdash \{type \ t :: \mathbf{t} :: K \cong t'\} : \{t :: \mathbf{t} : K\}} \quad (\text{HEAP TYPE SHARE})$$

$$\frac{\tilde{\Phi}; \{\}; \Xi \vdash A : K}{\tilde{\Phi}; \Xi \vdash \{t :: \mathbf{t} : K \triangleq A\} : \{\}} \quad (\text{HEAP SHARE DEF})$$

$$\frac{t' \in Names(\tilde{\Phi})}{\tilde{\Phi}; \Xi \vdash \{t :: \ell : K \cong t'\} : \{t \cong t' \in ty\}} \quad (\text{HEAP SHARE SHARE})$$

There are also rules for typing “private” type and value fields of a modules (private fields have the special external name  $\star$ ).

The type rule for symbol tables is relatively straightforward. A symbol table is a mapping from type and value external names to type and value labels, respectively, in the global type and value heaps. The side-conditions that  $\tilde{\Phi}' \subseteq \tilde{\Phi}$  means that, in checking the well-formedness of types and kinds, global type heap labels are chosen to be consistent with the internal type names used in the interface of the symbol table.

$$\frac{Int_E = (\Phi', \Psi', \Xi') \quad ST = \{\bar{\mathbf{t}} \mapsto \bar{t}, \bar{\mathbf{x}} \mapsto \bar{x}\} \quad \begin{array}{l} \tilde{\Phi}; \Delta \vdash \Phi' \quad \tilde{\Phi}; \Delta \vdash \Psi' \quad \tilde{\Phi}; \Delta; \Xi \vdash \Xi' \\ \Phi' = \{\bar{t} :: \bar{\mathbf{t}} : \bar{K}\} \quad \tilde{\Phi}' \subseteq \tilde{\Phi} \quad \Psi' = \{\bar{x} :: \bar{\mathbf{x}} : \bar{A}\} \end{array}}{\tilde{\Phi}; \Delta; \Xi; \tilde{\Psi} \vdash ST : Int_E} \quad (\text{VAL SYMBOL TABLE})$$

The *VENV*, *TENV* and *SHARE* metafunctions are defined as follows:

$$\begin{aligned} VENV(VH) &= \{(x : A) \mid (x :: \ell : A \mathcal{B}^e) \in VH\} \\ TENV(TH) &= \{(t : K) \mid (t :: \mathbf{t} : K \mathcal{B}^A) \in TH\} \\ SHARE(TH) &= \{(t_1 \cong t_2 \in ty) \mid (t_1 :: \mathbf{t}_1 : K \cong t_2) \in TH\} \end{aligned}$$

## B Semantics of Module Linking Instructions

In this appendix we provide more details of the static and dynamic semantics of the instructions of TMAL. The reduction rules use program states of the form

$$(\widetilde{TH}, \widetilde{VH}, R, I)$$

where  $R$  is a register file and  $I$  an instruction stream, and

$$\begin{aligned} \widetilde{TH} &= \{(t \mathcal{B}^A) \mid (t :: \mathbf{t} : K \mathcal{B}^A) \in TH\} \\ \widetilde{VH} &= \{(x \mathcal{B}^e) \mid (x :: \mathbf{x} : A \mathcal{B}^e) \in VH\} \end{aligned}$$

The global type and value heaps never contain shared bindings; such bindings are removed from an object file's type and value heaps, as part of initialization, before they are merged with the global heaps.

The type rules for `dllink`, `dlcoerce` and `dlrename` are similar to that for similar constructs described in [12]. We omit the rules here for lack of space. The reduction rules for these instructions are given by:

$$\begin{array}{c}
R(r_i^m) = x_i \text{ and } \widetilde{VH}(x_i) = [Int_I^i \Rightarrow (TH_i, VH_i) : Int_E^i], i = 2, 3 \\
Int_E^1 = (Int_E^2 \cup Int_E^3) \quad Int_I^1 = (Int_I^2 \sqcup Int_I^3) \setminus \{n \mid n \in idom(Int_E^1)\} \\
TH_1 = TH_2 \cup TH_3 \quad VH_1 = VH_2 \cup VH_3 \quad x_1 \notin dom(\widetilde{VH}) \\
idom(TH_1) \cap idom(TH_2) = \{\} \quad idom(VH_1) \cap idom(VH_2) = \{\} \\
\widetilde{VH}' = \widetilde{VH} \cup \{x_1 \triangleq [Int_I^1 \Rightarrow (TH_1, VH_1) : Int_E^1]\} \\
\hline
(\widetilde{TH}, \widetilde{VH}, R, (\text{dllink } r_1^m, r_2^m, r_3^m; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}', R[r_1^m \mapsto x_1], I) \\
\text{(RED DL LINK)}
\end{array}$$

$$\begin{array}{c}
R(r_2^m) = x_2 \text{ and } \widetilde{VH}(x_2) = [Int_I' \Rightarrow (TH', VH') : Int_E'] \\
OT = [Int_I \Rightarrow Int_E] \quad x_1 \notin dom(\widetilde{VH}) \\
\widetilde{VH}'' = \widetilde{VH} \cup \{x_1 \triangleq [Int_I \Rightarrow (COERCE(TH', Int_E'), COERCE(VH', Int_E')) : Int_E]\} \\
\hline
(\widetilde{TH}, \widetilde{VH}, R, (\text{dlcoerce } r_1^m, r_2^m, OT; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}'', R[r_1^m \mapsto x_1], I) \\
\text{(RED DL COERCE)}
\end{array}$$

$$\begin{array}{c}
R(r_2^m) = x_2 \text{ and } \widetilde{VH}(x_2) = [Int_I \Rightarrow (TH', VH') : Int_E] \\
x_1 \notin dom(\widetilde{VH}) \quad \widetilde{VH}'' = \widetilde{VH} \cup \{x_1 \triangleq [\rho(Int_I) \Rightarrow (\rho(TH'), \rho(VH')) : \rho(Int_E)]\} \\
\hline
(\widetilde{TH}, \widetilde{VH}, R, (\text{dlrename } r_1^m, r_2^m, \rho; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}'', R[r_1^m \mapsto x_1], I) \\
\text{(RED DL RENAME)}
\end{array}$$

The RED DL LINK type rule for the `dllink` instruction computes the new import list using the join operation *sqcup*, rather than simply unioning the import lists of the two object files being merged. This is because, for an import definition that is imported by both object files, the new import list must constrain the import to one compatible with both of the preceding import lists. For example, if one of the argument object files imports a type field *t* with kind *ty*, while another imports a type field *t* with kind  $\boxtimes$ `int`, then the new object file resulting from merging imports a definition of *t* with kind  $\boxtimes$ `int`. In computing the new import list, the `dllink` instruction removes from the import list any symbols that can be resolved against the combined export list. This operation is defined in [12].

The `dlcoerce` instruction uses the *COERCE* metafunction to hide bindings in the heaps (renaming their external name to  $\star$ ) that are made private by the new object file type and export interface. A definition of this metafunction is provided in [12].

The following type rule and reduction rule explain the semantics of the `dlopen` operation. The metafunction  $dom$  denotes the domain of a mapping, while  $idom(\widetilde{TH}) = dom(\widetilde{TH})$  and  $idom(\widetilde{VH}) = dom(\widetilde{VH})$ .  $\Gamma[r : A]$  denotes the replacement of the type of  $r$  (if any) in the register file type  $\Gamma$  with the new type  $A$ .  $R[r \mapsto w]$  denotes the replacement of the contents of  $r$  in the register file  $R$  with the new contents  $w$ . The `dlopen` operation expects register  $r^m$  to point to a module with type  $[Int_I \Rightarrow Int_E]$ , where  $Int_I = (\{\}, \{\}, \{\})$ . The operation leaves in register  $r^s$  a pointer to a symbol table with interface  $Int_E$ , after adding the heaps of the module to the program heaps:

$$\frac{\widetilde{\Phi}; \Delta; \widetilde{\Psi}; \Xi \vdash r^m : [(\{\}, \{\}, \{\}) \Rightarrow Int_E] \quad \Gamma' = \Gamma[r^s : Int_E]}{\widetilde{\Phi}; \widetilde{\Psi}; \Xi \vdash \{\Delta; \Gamma\} \text{ (dlopen } r^s, r^m) \{\Delta; \Gamma'\}} \quad (\text{INSTR DL OPEN})$$

$$\frac{\begin{array}{l} R(r^m) = x \text{ and } \widetilde{VH}(x) = [(\{\}, \{\}, \{\}) \Rightarrow (TH', VH') : Int_E] \\ ST = \{(t \mapsto t) \mid (t :: t : K) \in Int_E\} \cup \{(x \mapsto x) \mid (x :: x : A) \in Int_E\} \\ x' \notin idom(\widetilde{VH}) \cup idom(\widetilde{VH}') \quad idom(\widetilde{TH}) \cap idom(\widetilde{TH}') = \{\} \quad \widetilde{TH}'' = \widetilde{TH} \cup \widetilde{TH}' \\ idom(\widetilde{VH}) \cap idom(\widetilde{VH}') = \{\} \quad \widetilde{VH}'' = CLOS(\widetilde{VH} \cup \widetilde{VH}' \cup \{x' \triangleq ST\}) \end{array}}{(\widetilde{TH}, \widetilde{VH}, R, \text{ (dlopen } r^s, r^m; I)) \longrightarrow (\widetilde{TH}'', \widetilde{VH}'', R[r^s \mapsto x'], I)} \quad (\text{RED DL OPEN})$$

The  $CLOS(\widetilde{VH})$  operation removes shared value bindings of the form  $x : A \cong y$  from the value heap, by dereferencing  $y$  to its heap value definition:

$$CLOS(\widetilde{VH}) = \{(x \triangleq h) \mid x \in dom(\widetilde{VH}), h = Deref_{\widetilde{VH}}(x)\}$$

$$Deref_{\widetilde{VH}}(x) = \begin{cases} h & \text{if } (x \triangleq h) \in \widetilde{VH} \\ h & \text{if } (x \cong y) \in \widetilde{VH}, h = Deref_{\widetilde{VH}}(y) \end{cases}$$

The result of  $CLOS(\widetilde{VH})$  is undefined if  $\widetilde{VH}$  contains circular shared value bindings. This corresponds to an initialization failure due to cycles in the specification of initial values.

The type rule and reduction rule for the `dlsym.v` instruction are as follows:

$$\frac{\widetilde{\Phi}; \Delta; \widetilde{\Psi}; \Xi \vdash r^s : (\Phi', \Psi', \Xi') \quad (x :: \mathbf{x} : A) \in \Psi' \quad FV(A) \cap idom(\Phi') = \{\} \quad \Gamma' = \Gamma[r : A]}{\widetilde{\Phi}; \widetilde{\Psi}; \Xi \vdash \{\Delta; \Gamma\} \text{ (dlsym.v } r, r^s, \mathbf{x}) \{\Delta; \Gamma'\}} \quad (\text{INSTR DL SYMV})$$

$$\frac{R(r^s) = x \text{ and } \widetilde{VH}(x) = ST}{(\widetilde{TH}, \widetilde{VH}, R, \text{ (dlsym.v } r, r^s, \mathbf{x}; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}, R[r \mapsto ST(\mathbf{x})], I)} \quad (\text{RED DL SYMV})$$

$\widehat{R}(v)$  denotes the application of the register file  $R$  to the small value (register or word value)  $v$ :

$$\widehat{R}(v) = \begin{cases} w & \text{if } v = w \\ R(r) & \text{if } v = r \end{cases}$$

The type rule and reduction rule for the `dlsym.t` instruction are as follows:

$$\frac{\begin{array}{c} \widetilde{\Phi}; \Delta; \widetilde{\Psi}; \Xi \vdash r_2^s : (\Phi', \Psi', \Xi') \\ (t :: \mathbf{t} : K) \in \Phi' \quad \widetilde{\Phi} \cup \widetilde{\Phi}'; \Delta; \Xi \cup \Xi' \vdash K \leq K' \quad t \notin \text{dom}(\widetilde{\Phi}) \\ \Delta' = \Delta \cup \{t : K'\} \quad \Gamma' = \Gamma[r_1^s : (\widetilde{\Phi}_1 \cup \widetilde{\Phi}_2, \Psi', \Xi')] \end{array}}{\widetilde{\Phi}; \widetilde{\Psi}; \Xi \vdash \{\Delta; \Gamma\} \text{ (dl}_{\text{sym}}.t \ [t : K']_{r_1^s, r_2^s, \mathbf{t}} \ \{\Delta'; \Gamma'\})} \quad (\text{INSTR DL SYMT})$$

$$\frac{R(r_2^s) = x \text{ and } \widetilde{VH}(x) = ST \quad ST = ST' \uplus \{t \mapsto t\} \quad R' = R[r_1^s \mapsto ST']}{(\widetilde{TH}, \widetilde{VH}, R, \text{(dl}_{\text{sym}}.t \ [t : K']_{r_1^s, r_2^s, \mathbf{t}} \ I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}, R', \{t/t'\}I)} \quad (\text{RED DL SYMT})$$

In the reduction rule, the local type identifier  $t'$  is bound to the global type heap address  $t$  of the type definition pointed to by the symbol table. This allows the remainder of the instruction stream  $I$  to access the value heap definitions, pointed to by the symbol table, that have references to this type heap address.

Type heap addresses and type identifiers serve only to support type-checking of the assembly code, and are stripped for run-time execution. The substitution  $\{t/t'\}I$  is performed only in the abstract reduction semantics. Although we do not elaborate on it further here, the `dlsym.t` instruction can be generalized to import run-time type tags from a DLL, for languages such as Java and Modula-3 that associate type tags with some values.

For the next two instructions, we abuse notation slightly by allowing union and set difference operations to be applied to interfaces. These are to be understood as the operations distributing over the components of the interfaces, for example:

$$\begin{aligned} (\Phi_1, \Psi_1, \Xi_1) \cup (\Phi_2, \Psi_2, \Xi_2) &= (\Phi_1 \cup \Phi_2, \Psi_1 \cup \Psi_2, \Xi_1 \cup \Xi_2) \\ (\Phi, \Psi, \Xi) \cup \Phi' &= (\Phi \cup \Phi', \Psi, \Xi) \end{aligned}$$

The type rule for the `dlsetsym.v` instruction is reasonably straightforward. The only complication is that the type of the value field being assigned may have free type identifiers that are bound in the module. The typing rule relies on type sharing constraints in the module type that relate these locally bound type identifiers to global identifiers bound by the program type heap:

$$\frac{\begin{array}{c} \widetilde{\Phi}; \Delta; \widetilde{\Psi}; \Xi \vdash v : A \quad \widetilde{\Phi}; \Delta; \widetilde{\Psi}; \Xi \vdash r_2^m : [\text{Int}_I \Rightarrow \text{Int}_E] \\ \text{Int}_I = (\Phi_I, \Psi_I, \Xi_I) \quad \text{Int}_E = (\Phi_E, \Psi_E, \Xi_E) \quad (x :: \mathbf{x} : B) \in \Psi_I \\ \widetilde{\Phi} \cup \widetilde{\Phi}_I \cup \widetilde{\Phi}_E; \Delta; \Xi \cup \Xi_I \cup \Xi_E \vdash A = B \in \text{ty} \\ \Gamma' = \Gamma[r_1^m \mapsto [(\text{Int}_I - \{x :: \mathbf{x} : B\}) \Rightarrow (\text{Int}_E \cup \{x :: \mathbf{x} : B\})]] \end{array}}{\widetilde{\Phi}; \widetilde{\Psi}; \Xi \vdash \{\Delta; \Gamma\} \text{ (dl}_{\text{setsym}}.v \ r_1^m, r_2^m, v, \mathbf{x}) \ \{\Delta; \Gamma'\}} \quad (\text{INSTR DL SETSYMV})$$

$$\begin{array}{c}
 R(r_2^m) = x \text{ and } \widetilde{VH}(x) = [Int_I \Rightarrow (TH', VH') : Int_E] \quad (x :: \mathbf{x} : A) \in Int_I \quad \widehat{R}(v) = y \\
 Int'_I = Int_I - \{x :: \mathbf{x} : A\} \quad Int'_E = Int_E \cup \{x :: \mathbf{x} : A\} \\
 z \notin \text{dom}(\widetilde{VH}) \quad \widetilde{VH}'' = \widetilde{VH} \cup \{z \triangleq [Int'_I \Rightarrow (TH', VH' \cup \{x :: \mathbf{x} : A \cong y\}) : Int'_E]\} \\
 \hline
 (\widetilde{TH}, \widetilde{VH}, R, (\text{dlsetsym.v } r_1^m, r_2^m, v, \mathbf{x}; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}'', R[r_1^m \mapsto z], I) \\
 \text{(RED DL SETSYMV)}
 \end{array}$$

The `dlsetsym.t` instruction for assigning a type field in a module similarly relies on type sharing to equate any local type identifiers with global type identifiers in the kind of the type being assigned. Free type identifiers may appear free in the kind of a field with box kind. Once a type field has been assigned, a type sharing constraint is added to the export interface of the module, to allow subsequent value fields to be assigned:

$$\begin{array}{c}
 \widetilde{\Phi}; \Delta \vdash t' : K \quad \widetilde{\Phi}; \Delta; \widetilde{\Psi}; \Xi \vdash r_2^m : [Int_I \Rightarrow Int_E] \\
 Int_I = (\Phi_I, \Psi_I, \Xi_I) \quad Int_E = (\Phi_E, \Psi_E, \Xi_E) \quad (t :: \mathbf{t} : K') \in \Phi_I \\
 \widetilde{\Phi} \cup \widetilde{\Phi}_I \cup \widetilde{\Phi}_E; \Delta; \Xi \cup \Xi_I \cup \Xi_E \vdash K = K' \\
 \Gamma' = \Gamma[r_1^m \mapsto [(Int_I - \{t :: \mathbf{t} : K'\}) \Rightarrow (Int_E \cup \{(t :: \mathbf{t} : K'), (t \cong t' \in K')\})]] \\
 \hline
 \widetilde{\Phi}; \widetilde{\Psi}; \Xi \vdash \{\Delta; \Gamma\} \text{ (dlsetsym.t } r_1^m, r_2^m, t', \mathbf{t}) \{\Delta; \Gamma'\} \\
 \text{(INSTR DL SETSYMT)}
 \end{array}$$

$$\begin{array}{c}
 R(r_2^m) = x \text{ and } \widetilde{VH}(x) = [Int_I \Rightarrow (TH', VH') : Int_E] \quad (t :: \mathbf{t} : K) \in Int_I \\
 Int'_I = Int_I - \{t :: \mathbf{t} : K\} \quad Int'_E = Int_E \cup \{(t :: \mathbf{t} : K), (t \cong t' \in K)\} \\
 z \notin \text{dom}(\widetilde{VH}) \quad \widetilde{VH}'' = \widetilde{VH} \cup \{z \triangleq [Int'_I \Rightarrow (TH' \cup \{t :: \mathbf{t} : K \cong t'\}, VH') : Int'_E]\} \\
 \hline
 (\widetilde{TH}, \widetilde{VH}, R, (\text{dlsetsym.t } r_1^m, r_2^m, t', \mathbf{t}; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}'', R[r_1^m \mapsto z], I) \\
 \text{(RED DL SETSYMT)}
 \end{array}$$

Finally the reduction rules for the instructions for creating a DLL, and for extracting a module from a DLL, are as follows:

$$\begin{array}{c}
 x \notin \text{dom}(VH) \quad \widetilde{VH}' = \widetilde{VH} \cup \{x \triangleq \langle \widehat{R}(v), OT \rangle\} \\
 \hline
 (\widetilde{TH}, \widetilde{VH}, R, (\text{dlldynamic } r, v, OT; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}', R[r \mapsto x], I) \\
 \text{(RED DL DYNAMIC)}
 \end{array}$$

$$\begin{array}{c}
 R(r_1) = x \text{ and } \widetilde{VH}(x) = \langle \langle y, OT \rangle \rangle \\
 OT = [Int_I \Rightarrow Int_E] \quad OT'' = [Int''_I \Rightarrow Int''_E] \quad \widetilde{VH}(y) = [Int'_I \Rightarrow (TH', VH') : Int'_E] \\
 TENV(\widetilde{TH}); \{\}; \text{SHARE}(\widetilde{TH}) \vdash [Int_I \Rightarrow Int_E] \preceq [Int''_I \Rightarrow Int''_E] \\
 z \notin \text{dom}(\widetilde{VH}) \quad \widetilde{VH}'' = \widetilde{VH} \cup \{z \triangleq [Int''_I \Rightarrow (TH', VH') : Int''_E]\} \\
 \hline
 (\widetilde{TH}, \widetilde{VH}'', R, (\text{dlload } r^m, r_1, r_2, OT''; I)) \longrightarrow (\widetilde{TH}, \widetilde{VH}'', R[r^m \mapsto z], I) \\
 \text{(RED DL LOAD SUCC)}
 \end{array}$$

The last rule handles the case where loading a DLL (reflecting a DLL from the core language into the module language) succeeds with a runtime type check.

There is also an associated rule for when the runtime type check fails; in this case, control transfers to the address specified by the address register  $r_2$ , i.e.,  $r_2$  contains a pointer to a failure continuation that should be invoked if the runtime type check fails.