

Typing a Multi-Language Intermediate Code

Andrew D. Gordon

Don Syme

Microsoft Research

Abstract

The Microsoft .NET Framework is a new computing architecture designed to support a variety of distributed applications and web-based services. .NET software components are typically distributed in an object-oriented intermediate language, Microsoft IL, executed by the Microsoft Common Language Runtime. To allow convenient multi-language working, IL supports a wide variety of high-level language constructs, including class-based objects, inheritance, garbage collection, and a security mechanism based on type safe execution.

This paper precisely describes the type system for a substantial fragment of IL that includes several novel features: certain objects may be allocated either on the heap or on the stack; those on the stack may be boxed onto the heap, and those on the heap may be unboxed onto the stack; methods may receive arguments and return results via typed pointers, which can reference both the stack and the heap, including the interiors of objects on the heap. We present a formal semantics for the fragment. Our typing rules determine well-typed IL instruction sequences that can be assembled and executed. Of particular interest are rules to ensure no pointer into the stack outlives its target. Our main theorem asserts type safety, that well-typed programs in our IL fragment do not lead to untrapped execution errors.

Our main theorem does not directly apply to the product. Still, the formal system of this paper is an abstraction of informal and executable specifications we wrote for the full product during its development. Our informal specification became the basis of the product team's working specification of type-checking. The process of writing this specification, deploying the executable specification as a test oracle, and applying theorem proving techniques, helped us identify several security critical bugs during development.

ACM COPYRIGHT NOTICE. Copyright ©2001 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept, ACM Inc., fax +1 (212) 869-0481, or permissions@acm.org.

1 Introduction

This paper describes typing and evaluation rules, and a type safety theorem, for a substantial fragment of the intermediate language (IL) executed by Microsoft's Common Language Runtime. The rules are valuable because they succinctly and precisely account for some unusual and subtle features of the type system.

Background: IL The Common Language Runtime is a new execution environment with a rich object-oriented class library through which software components written in diverse languages may interoperate. Using the Visual Studio .NET development environment, .NET components can be written in the new object-oriented language C# [HW00], as well as Visual Basic, Visual C++, and the scripting languages VBScript and JScript. Furthermore, prototype .NET compilers exist for COBOL, Component Pascal, Eiffel, Haskell, Mercury, Oberon, Ocaml, and Standard ML.

Type-checking of .NET components implemented in IL has already proved useful for finding code generation bugs. Moreover, the .NET security model assumes type-safe behaviour; type-checking is therefore useful for handling untrusted components. Given these and other applications, the IL type system is worthy of formal specification.

Background: Executable Specifications This paper is one outcome of a research project to evaluate and develop formal specification techniques for describing and analyzing type-checkers in general. Specifically, we applied these techniques to the study of IL. We began by writing a detailed specification of type-checking method bodies. This was an informal document in the style of most language references. Eventually, this document was adopted by the product team as the basis of their detailed specification of type-checking. In parallel, following a methodology advocated by Syme [Sym98], we wrote formal specifications for various IL subsets suitable for comparative testing and formal proof. The executable part of these specifications is in a functional fragment of ML, the rest in higher order logic (HOL). We can compile and run the executable part as an IL type-checker. Since it is purely functional code, we may also interpret it as HOL and use it for theorem proving in DECLARE [Sym98]. In principle, this strategy allows the same source code to serve both as an oracle for testing actual implementations and as a model for formal validation. So far, we have built an ML type-checker for a largely complete subset of the IL type system, but have formally verified only a rather smaller fragment.

As is well known [Coh89], even formal proof cannot guarantee the absence of implementation defects, simply because one has to abstract from details of the environment when writing formal models. We found that developing a test suite that used our formal model as an oracle was an important way of making our model consistent with the runtime. Our suite included about 30,000 automatically generated tests. Our experience was that testing remains the only viable way of relating a specification to software of the complexity we were considering. One of our slogans: *if you specify, you must test*. Writing a formal specification without generating tests may be viable once a design has been frozen, but is simply not effective during the design of a new system. Eventually, we handed over our suite to the test team, who maintain it, and who have found bugs using it.

This Paper: An IL Fragment The main part of the paper concerns an IL fragment based on reference, value, and pointer types.

At its core, the fragment is a class-based object-oriented language with field update and simple imperative control structures. This core is comparable to the imperative object calculus [AC96, GHL99] and to various fragments of Java [DE97, IPW99]. An item of a reference type is a pointer to a heap-allocated object.

Moreover, our fragment includes value and pointer types:

- An item of a value type is a sequence of machine words representing the fields of the type. Value types support the compilation of C-style structs, for instance. Value types may be stack-allocated and passed by value. A `box` instruction turns a value type into a heap-allocated object by copying, and an `unbox` instruction performs the inverse coercion. Hence, when convenient, value types may be treated as ordinary heap-allocated objects.
- An item of pointer type is a machine address referring either to a heap-allocated object or to a variable in the call stack or to an interior field of one of these. The main purpose of pointer types is to allow methods to receive arguments or return results by reference.

We selected these types because they are new constructs not previously described by formal typing rules, and because their use needs to be carefully limited to avoid type loop-holes. In particular, we must take care that stack pointers do not outlive their targets.

For the sake of clarity, our presentation of the semantics differs from the ML code in our executable specifications in two significant ways:

- First, we adopt the standard strategy of presenting the type system as logical inference rules. Such rules are succinct, but not directly executable; we found it better to write executable ML when we initially wrote our specifications in order to help with testing. Still, typing rules are better than code for presenting a type system and for manual proof.
- Second, we adopt a new, non-standard strategy of assuming that each method body has been parsed into a tree-structured applicative expression. Each expression consists of an IL instruction applied to the subexpressions that need to be evaluated to compute the instruction's arguments. This technique allows us to

concentrate on specifying the typing conditions for each instruction, and to suppress the algorithmic details of how a type-checker would compute the types of the arguments to each instruction. These algorithmic details are important in any implementation, but they are largely irrelevant to specifying type safety.

Finally, in the spirit of writing specifications to support testing, our applicative expressions use the standard IL assembler syntax. Hence, any method body that is well-typed according to our typing rules can be assembled and tested on the running system.

In summary, the principal technical contributions made by this paper are the following:

- New typing and evaluation rules for value and pointer types, together with a type safety result, Theorem 1.
- The idea that the essence of a low-level intermediate language can be presented in an applicative notation.

Future Challenges: As we have discussed, this project is a successful demonstration of the value of writing executable, formal specifications during product development.

On the other hand, the main theorem of this paper does not apply to the full product; type safety bugs may well be discovered. An unfulfilled ambition of ours is to prove soundness of the typing rules for the full language through mechanized theorem-proving. So a future challenge is to further develop scalable and maintainable techniques for mechanized reasoning. A soundness proof for the whole of IL would be an impressive achievement. To apply theorem proving during product development, scalability and maintainability of proof scripts are important. Scripts should be scalable in the sense that human effort is roughly linear in the size of the specification (with a reasonable constant factor), or else proof construction cannot keep up with new features as they are added. Scripts should be maintainable in the sense that they are robust in the face of minor changes to the specification, or else proof construction cannot keep up with the inevitable revisions of the design.

In the meantime, another challenge is to develop systematic techniques for test case generation.

A third challenge is to integrate executable specifications, such as our ML type-checker, into the product itself. The .NET Framework, like other component models, itself contributes to this goal, in that its support for multi-language working would easily allow a critical component to be written in ML, say, even if the rest of the product is not.

The remainder of the paper proceeds as follows. Section 2 presents the typing and evaluation rules for our IL fragment, and states our main theorem. Section 3 explains a potentially useful liberalisation of the type system. Section 4 summarizes the omissions from our IL fragment. Section 5 discusses related work. Section 6 concludes.

Proofs omitted from this conference paper appear in a technical report [GS00].

2 A Formal Analysis of BIL, a Baby IL

This section makes the main technical contributions of the paper. We present a substantial fragment of IL that includes enough detail to allow a formal analysis of reference, value, and pointer types, but omits many features not related to these. We name this fragment Baby IL, or BIL for short.

Section 2.1 describes the type structure of BIL. In Section 2.2, we specify the instructions that may appear in method bodies of BIL, and explain their informal semantics. In Section 2.3, we specify a formal memory model for BIL, and a formal semantics for the evaluation of method bodies. In Section 2.4, we specify a formal type system for type-checking method bodies. Section 2.5 introduces conformance relations that express when intermediate states arising during evaluation are type-correct. Finally, Section 2.6 concludes this analysis by stating our Type Safety Theorem.

2.1 Type Structure and Class Hierarchy

All BIL methods run in an execution environment that contains a fixed set of classes. Each class specifies types for a set of field variables, and signatures for a set of methods. Each object belongs to a class. The memory occupied by each object consists of values for each field specified by its class. Methods are shared between all objects of a class (and possibly other classes). Objects of all classes may be stored boxed in a heap, addressed by heap references. Objects of certain classes—known as value classes—may additionally be stored unboxed in the stack or as fields embedded in other objects.

Formally, we assume three sets, *Class*, *Field*, and *Meth*, the sets of class, field, and method names, respectively, and a set *ValueClass* \subseteq *Class* of value class names. We assume a distinguished class name `System.Object` such that `System.Object` \notin *ValueClass*.

Classes, Fields, Methods:

$c \in \textit{Class}$	class name
$vc \in \textit{ValueClass} \subseteq \textit{Class}$	value class name
<code>System.Object</code> $\in \textit{Class} - \textit{ValueClass}$	root of hierarchy
$f \in \textit{Field}$	field name
$\ell \in \textit{Meth}$	method name

Types describe objects, the fields of objects, the arguments and results of methods, and the intermediate results arising during evaluation of method bodies.

Types:

$A, B \in \textit{Type} ::=$	type
<code>void</code>	no bits
<code>int32</code>	32 bit signed integer
<code>class c</code>	boxed object
<code>value class vc</code>	unboxed object
<code>A&</code>	pointer to <i>A</i>

The type `void` describes the absence of data, no bits; `void` is only used for the results of methods or parts of method bodies that return no actual result.

The type `int32` describes a 32 bit integer; BIL uses integers to represent predicates for conditionals and while-loops but includes no primitive arithmetic operations. (IL features a rich selection of numeric types and arithmetic operations.)

A *reference type class* *c* describes a pointer to a boxed object (heap-allocated, subject to garbage collection).

A *value type value class* *vc* describes an unboxed object—a sequence of words representing the fields of the value class *vc*, akin to a C struct. The associated reference type, `class vc` describes a pointer to a boxed object—a heap-allocated representation of the fields.

Finally, a *pointer type* *A&* describes a pointer to data of type *A*, which may be stored either in the heap or the stack.

To avoid dangling pointers—pointers that outlive their targets—our type system restricts pointers as follows. An important use for pointers in IL is to allow arguments and results to be passed by reference. The following are sufficient conditions to type-check this motivating usage while preventing dangling pointers. The following are not necessary conditions; we explain a useful and safe liberalisation in Section 3.

BIL Pointer Confinement Policy:

- (1) No field may hold a pointer.
- (2) No method may return a pointer.
- (3) No pointer may be stored indirectly via another pointer.

(IL itself follows a slightly stricter policy that bans pointers to pointers altogether.) Each of the conditions prevents a way of creating a dangling pointer. If a field could hold a pointer, a method could store a pointer into its stack frame in an object boxed on the heap. If a method could return a pointer, a method could simply return a pointer into its stack frame. If a pointer could be stored indirectly, a method could store a pointer into its stack frame through a pointer to an object boxed on the heap or to an earlier stack frame. In each case, the pointer would outlive its target as soon as the method had returned.

The following predicate identifies types containing no pointers.

Whether a Type Contains No Pointer:

$$\textit{pointerFree}(A) \Leftrightarrow \neg(A = B\& \text{ for some } B)$$

Next, a *method signature* $B \ell(A_1, \dots, A_n)$ refers to a method named ℓ that expects a vector of arguments with types A_1, \dots, A_n , and whose result has type *B*. No two methods in a given class may share the same signature, though they may share the same method name.

Method signature:

$$\textit{sig} \in \textit{Sig} ::= B \ell(A_1, \dots, A_n) \quad \text{method signature}$$

We assume the execution environment organises classes into an inheritance hierarchy. We write *c inherits c'* to mean that *c* inherits from *c'*. We induce a *subtype relation*, $A <: B$, from the inheritance hierarchy. Our type system supports *subsumption*: if $A <: B$ an item of type *A* may be used in a context expecting an item of type *B*. The only non-trivial subtyping is between reference types. The subtype relation is the least to satisfy the following rules.

Subtype Relation: $A <: B$

$$\frac{}{A <: A} \quad \frac{\textit{c inherits c'}}{\textit{class c} <: \textit{class c'}}$$

We assume that the relation *c inherits c'* is transitive, and therefore so is the relation $A <: B$.

The IL assembler recognises a fairly standard notation for single inheritance that allows a class to inherit methods and fields from a single superclass. One might define the inheritance relation by formalizing such a syntax and type-checking rules. Instead, since our focus is type-checking the

BIL instruction set, it is easier and more concise to simply axiomatize the intended properties of the hierarchy. (Although the IL syntax disallows multiple inheritance, it happens that our axioms allow a class to inherit from two super-classes that are incomparable according to the inheritance relation.)

Formally, we assume there is an *execution environment* consisting of three components—a function $fields(c)$, a function $methods(c)$, and an inheritance relation $c inherits c'$ —that satisfy the following axioms:

Execution Environment: ($fields, methods, inherits$)	
$fields \in Class \rightarrow (Field \xrightarrow{fin} Type)$	fields of a class
$methods \in Class \rightarrow (Sig \xrightarrow{fin} Body)$	methods of a class
$inherits \subseteq Class \times Class$	class hierarchy
$c inherits c$	(Hi Refl)
$c inherits c' \wedge c' inherits c'' \Rightarrow c inherits c''$	(Hi Trans)
$c inherits c' \wedge c' inherits c \Rightarrow c = c'$	(Hi Antisymm)
$c inherits System.Object$	(Hi Root)
$c inherits d \wedge f \in dom(fields(d)) \Rightarrow f \in dom(fields(c)) \wedge fields(c)(f) = fields(d)(f)$	(Hi fields)
$c inherits d \Rightarrow dom(methods(d)) \subseteq dom(methods(c))$	(Hi methods)
$c inherits vc \Rightarrow c = vc$	(Hi Val)
$pointerFree(fields(c)(f))$	(Good fields)
$B \ell(A_1, \dots, A_n) \in dom(methods(c)) \Rightarrow pointerFree(B)$	(Good methods)

For any class c , $fields(c) \in Field \xrightarrow{fin} Type$, the set of finite maps from field names to types. If $fields(c) = f_i \mapsto A_i^{i \in 1..n}$, the class c has exactly the set of fields named f_1, \dots, f_n with types A_1, \dots, A_n , respectively.

(The notation $f_i \mapsto A_i^{i \in 1..n}$ exemplifies our notation for finite maps in general. We let $dom(f_i \mapsto A_i^{i \in 1..n}) = \{f_1, \dots, f_n\}$. We assume that the f_i are distinct. Let $(f_i \mapsto A_i^{i \in 1..n})(f) = A_i$ if $f = f_i$ for some $i \in 1..n$, and otherwise be undefined.)

For any class c , $methods(c) \in Sig \xrightarrow{fin} Body$, the set of finite maps from method signatures to method bodies. We define the set *Body* of method bodies—instruction sequences—in the next section. If $methods(c) = sig_i \mapsto b_i^{i \in 1..n}$, the class c has exactly methods with signatures sig_1, \dots, sig_n , implemented by the bodies b_1, \dots, b_n , respectively.

A binary relation on classes, *inherits*, formalizes the inheritance hierarchy. Axioms (Hi Refl) and (Hi Trans) guarantee it is reflexive and transitive. (Hi Antisymm) asserts it is anti-symmetric, that is, there are no cycles in the hierarchy. According to (Hi Root), every class inherits from `System.Object`, the root of the hierarchy.

Suppose that c is a subclass of d , that is, $c inherits d$. By subsumption, an object of the subclass c may be used in a context expecting an object of the superclass d . Accordingly, (Hi fields) asserts that every field specified by d is also present in the subclass c . The axiom (Hi methods) asserts that every method signature implemented by d is also implemented by the subclass c , though not necessarily by the same method body.

In order to implement a method invocation on an object, we need to know the class of the object. In general, we

cannot statically determine the class of an object from its type, since by subsumption it may in fact be a subclass of the class named in its type. Therefore, each boxed object is tagged in our formal memory model with the name of its class. On the other hand, for the sake of space efficiency, unboxed objects include no type information. Therefore, we must rely on statically determining the class of an unboxed object from its type. For this to be possible, axiom (Hi Val) prevents any other class from inheriting from a value class. So the actual class of any unboxed object is the same as the class named in its type.

Axioms (Good fields) and (Good methods) implement points (1) and (2) of the Pointer Confinement Policy.

We end this section by exemplifying how value and pointer types provide possibly more efficient alternatives to reference types for returning multiple results. Suppose there is a class `Point` $\in ValueClass$ such that $fields(Point) = x \mapsto int32, y \mapsto int32$, that is, a class with two integer fields. Here are three alternative signatures for returning a `Point` from a method named `mouse`:

- As a boxed object: `class Point mouse ()`.
- As an unboxed object: `value class Point mouse ()`.
- In a pre-allocated unboxed object passed by reference: `void mouse (value class Point&)`.

2.2 Syntax of Method Bodies

BIL is a deterministic, single-threaded, imperative, class-based object-oriented language. For the sake of simplicity, we omit constructs for error or exception handling. This section specifies the instruction set as tree-structured applicative expressions, most of which represent an application of an instruction to a sequence of argument expressions. Since each applicative expression is in a postfix notation, it can also be read as a sequence of atomic instructions. We have chosen our syntax carefully so that, subject to very minor editing, this sequence of atomic instructions can be parsed by the IL assembler (as well as our own IL type-checker).

We express the syntax of our conditional and iteration constructs using assembler labels, ranged over by L .

A *method reference* $B c::\ell(A_1, \dots, A_n)$ refers to the method with signature $B \ell(A_1, \dots, A_n)$ in class c .

Inspired by FJ [IPW99], we assume for simplicity that each class has exactly one constructor, whose arguments are the initial values assumed by the fields of the new object. The *constructor reference* for a class c takes the form `void c::ctor(A1, ..., An)`. Constructors are only called to create a new object; `.ctor` $\notin Meth$.

Method and Constructor References:

L	assembler label
$M ::= B c::\ell(A_1, \dots, A_n)$	method reference
$K ::= void c::ctor(A_1, \dots, A_n)$	constructor reference

Applicative Expressions for Method Bodies:

$i4$	32 bit signed integer
$a, b \in Body ::=$	method body
<code>ldc.i4 i4</code>	load integer
<code>a brtrue L₁ b₀ br L₂ L₁:b₁ L₂:</code>	conditional
<code>L₁: a brfalse L₂ b br L₁ L₂:</code>	while-loop
<code>a b</code>	sequencing
<code>a ldind</code>	load indirect

$a\ b\ stind$	store indirect
$ldarga\ j$	load argument address
$a\ starg\ j$	store into argument
$a_1 \dots a_n\ newobj\ K$	create new object
$a_0\ a_1 \dots a_n\ callvirt\ M$	call on boxed object
$a_0\ a_1 \dots a_n\ call\ instance\ M$	call on unboxed object
$a\ ldflda\ A\ c::f$	load field address
$a\ b\ stfld\ A\ c::f$	store into field
$a\ box\ vc$	copy value to heap
$a\ unbox\ vc$	fetch pointer to value

Conditionals and while-loops are not primitive instructions in IL, but it is worthwhile to make them primitive in BIL to allow a simple format for evaluation and typing rules. We have carefully chosen a syntax for these constructs by assembling suitable IL branch instructions and labels. We assume that the assembler labels in these expressions do not appear in any of their subexpressions. The result is a syntax that is a little cryptic but that does produce IL instruction sequences with the appropriate semantics. These abbreviations are more readable:

Abbreviations for Conditionals and While-Loops:

$a\ b_0\ b_1\ cond \triangleq a\ brtrue\ L_1\ b_0\ br\ L_2\ L_1:b_1\ L_2:$
 $a\ b\ while \triangleq L_1:a\ brfalse\ L_2\ b\ br\ L_1\ L_2:$

The technique of representing assembly language in an applicative syntax works for this paper because it can express all the operations on reference, value, and pointer types. We express structured control flow like conditionals or while-loops in this style by treating an assembly of IL branch instructions as a primitive BIL instruction. Still, the technique may not scale well to express control flow such as arbitrary branching within a method or exception handling.

IL includes primitive instructions `ldfld` and `ldarg` to load the contents of an object field or an argument. Instead of taking these as primitives in BIL, we can derive them as follows:

Derived Instructions:

$a\ ldfld\ A\ c::f \triangleq a\ ldflda\ A\ c::f\ ldind$
 $a\ ldarg\ j \triangleq a\ ldarga\ j\ ldind$

2.3 Evaluating Method Bodies

The memory model consists of a heap of objects and a stack of method invocation frames, each of which is a vector of arguments. Our semantics abstracts away from the details of evaluation stacks or registers.

We assume a collection of heap references, p, q , pointing to boxed objects in the heap.

A *pointer* takes one of three forms. A pointer p refers to the boxed object at p . A pointer (i, j) refers to argument j of stack frame i . A pointer $ptr.f$ refers to field f of the object referred to by ptr .

A *result* is either void 0 , an integer $i4$, a pointer ptr , or an unboxed object $f_i \mapsto u_i^{i \in 1..n}$, a finite map consisting of a sequence of results u_1, \dots, u_n corresponding to the fields f_1, \dots, f_n , respectively.

References, Pointers, Results:

p, q heap reference

$ptr ::=$	pointer
p	pointer to boxed object p
(i, j)	pointer to argument j of frame i
$ptr.f$	pointer to field f of object at ptr
$u, v ::=$	result
0	void
$i4$	integer
ptr	pointer
$f_i \mapsto u_i^{i \in 1..n}$	value: unboxed object

Next, we formalize our memory model. A *heap* is a finite map from references to boxed objects, each taking the form $c[f_i \mapsto u_i^{i \in 1..n}]$, where c is the class of the object, and $f_i \mapsto u_i^{i \in 1..n}$ is its unboxed form. A *frame*, fr , is a vector of arguments written as $.args(u_0, \dots, u_n)$: u_0 is the self parameter; u_1, \dots, u_n are the computed arguments. A *stack*, s , is a list of frames $fr_1 \dots fr_n$. Finally, a *store* is a heap paired with a stack.

Memory Model:

$o ::= c[f_i \mapsto u_i^{i \in 1..n}]$	boxed object
$h ::= p_i \mapsto o_i^{i \in 1..n}$	heap
$fr ::= .args(u_0, \dots, u_n)$	frame: vector of arguments
$s ::= fr_1 \dots fr_n$	stack (grows left to right)
$\sigma ::= (h, s)$	store

The example heap $h = p \mapsto c[f_1 \mapsto 0, f_2 \mapsto (g \mapsto 1)]$ consists of a single boxed object $c[f_1 \mapsto 0, f_2 \mapsto (g \mapsto 1)]$ at heap reference p . The boxed object is of class c and consists of fields named f_1 and f_2 . The first field contains the integer 0. The second field contains the unboxed object $g \mapsto 1$, which itself consists of a field named g containing the integer 1.

The example stack $s = .args(p, p.f_2.g).args(p, (1, 1))$ consists of two frames. The bottom of the stack is the frame $.args(p, p.f_2.g)$, consisting of two arguments, a reference to the boxed object at p , and a pointer to field g of field f_2 of the same object. The top of the stack is the frame $.args(p, (1, 1))$, consisting of two arguments, a reference to the boxed object at p , and the pointer $(1, 1)$, which refers to argument 1 of frame 1, that is, the pointer $p.f_2.g$.

We rely on two auxiliary partial functions for dereferencing and updating pointers in a store:

Auxiliary Functions for Lookup and Update:

$lookup(\sigma, ptr)$	lookup ptr in store σ
$update(\sigma, ptr, v')$	update store σ at ptr with result v'

We explain the intended meaning of store lookup and update by example. Let store $\sigma = (h, s)$ where h and s are the heap and stack examples introduced above. Then $lookup(\sigma, (1, 0))$ is the reference p stored in argument 0 of frame 1, and $lookup(\sigma, p.f_2.g)$ is the integer 1 stored in field g of the unboxed object stored in field f_2 of the boxed object at p . The outcome of $update(\sigma, (2, 0), 1)$ is to update σ by replacing the reference p in argument 0 of frame 2 with 1. Similarly, the outcome of $update(\sigma, p.f_1.g, 0)$ is to update σ by replacing the integer 1 in field g of field f_1 of the boxed object at p with the integer 0.

A little functional programming suffices to define these two functions; we give the full definitions in the Appendix.

Our operational semantics of method bodies is a formal judgment $\sigma \vdash b \rightsquigarrow v \cdot \sigma'$ meaning that in an initial store

σ , the body b evaluates to the result v , leaving final store σ' . (A “judgment” is simply a predicate defined by a set of inference rules.)

Evaluation Judgment:

$$\sigma \vdash b \rightsquigarrow v \cdot \sigma' \quad \text{given } \sigma, \text{ body } b \text{ returns } v, \text{ leaving } \sigma'$$

Our semantics takes the form of an interpreter. The rest of this section presents the formal rules for deriving evaluation judgments, interspersed with informal explanations.

Evaluation Rules for Control Flow:

(Eval ldc)	(Eval Seq)
$\sigma \vdash \text{ldc.i4 } i_4 \rightsquigarrow i_4 \cdot \sigma$	$\frac{\sigma \vdash a \rightsquigarrow u \cdot \sigma' \quad \sigma' \vdash b \rightsquigarrow v \cdot \sigma''}{\sigma \vdash a b \rightsquigarrow v \cdot \sigma''}$

(Eval Cond) (where $j = 0$ if $i_4 = 0$, otherwise $j = 1$)
$\frac{\sigma \vdash a \rightsquigarrow i_4 \cdot \sigma' \quad \sigma' \vdash b_j \rightsquigarrow v \cdot \sigma''}{\sigma \vdash a b_0 b_1 \text{ cond} \rightsquigarrow v \cdot \sigma''}$

(Eval While 0)
$\frac{\sigma \vdash a \rightsquigarrow 0 \cdot \sigma'}{\sigma \vdash a b \text{ while} \rightsquigarrow 0 \cdot \sigma'}$

(Eval While 1) (where $i_4 \neq 0$)
$\frac{\sigma \vdash a \rightsquigarrow i_4 \cdot \sigma' \quad \sigma' \vdash b \rightsquigarrow v \cdot \sigma'' \quad \sigma'' \vdash a b \text{ while} \rightsquigarrow u \cdot \sigma'''}{\sigma \vdash a b \text{ while} \rightsquigarrow u \cdot \sigma'''}$

The expression `ldc.i4 i_4` evaluates to the integer i_4 .

The expression `$a b$` evaluates a , returning void (that is, nothing). The result of the whole expression is then the result of evaluating b .

The expression `$a b_0 b_1 \text{ cond}$` evaluates a to an integer i_4 . The result of the whole conditional is then the result of evaluating b_0 if $i_4 = 0$, and evaluating b_1 otherwise.

The expression `$a b \text{ while}$` evaluates a to an integer i_4 . If $i_4 = 0$ evaluation terminates, returning void. Otherwise, the body b is evaluated, returning void, and then evaluation of `$a b \text{ while}$` repeats.

Evaluation Rules for Pointer Types:

(Eval ldind)
$\frac{\sigma \vdash a \rightsquigarrow \text{ptr} \cdot \sigma'}{\sigma \vdash a \text{ ldind} \rightsquigarrow \text{lookup}(\sigma', \text{ptr}) \cdot \sigma'}$

(Eval stind)
$\frac{\sigma \vdash a \rightsquigarrow \text{ptr} \cdot \sigma' \quad \sigma' \vdash b \rightsquigarrow v \cdot \sigma''}{\sigma \vdash a b \text{ stind} \rightsquigarrow 0 \cdot \text{update}(\sigma'', \text{ptr}, v)}$

The expression `$a \text{ ldind}$` evaluates a to a pointer, and then returns the outcome of dereferencing the pointer.

The expression `$a b \text{ stind}$` evaluates a to a pointer, stores the result of evaluating b in the (heap or stack) location addressed by the pointer, and returns void.

Evaluation Rules for Arguments:

(Eval ldarga)
$\frac{\sigma = (h, fr_1 \dots fr_i)}{\sigma \vdash \text{ldarga } j \rightsquigarrow (i, j) \cdot \sigma}$

(Eval starg)
$\frac{\sigma \vdash a \rightsquigarrow u \cdot \sigma' \quad \sigma' = (h', fr_1 \dots fr_i)}{\sigma \vdash a \text{ starg } j \rightsquigarrow 0 \cdot \text{update}(\sigma', (i, j), u)}$

The expression `$\text{ldarga } j$` returns a pointer to argument j in the current stack frame.

The expression `$a \text{ starg } i$` evaluates a , stores the result in argument i in the current stack frame, then returns void.

Evaluation Rules for Reference Types Only:

(Eval newobj) (where $K = \text{void } c::\text{ctor}(A'_1, \dots, A'_m)$) $c \notin \text{ValueClass}$
$\frac{\text{fields}(c) = f_i \mapsto A_i \quad i \in 1..n \quad \sigma_i \vdash a_i \rightsquigarrow v_i \cdot \sigma_{i+1} \quad \forall i \in 1..n \quad \sigma_{n+1} = (h, s) \quad p \notin \text{dom}(h) \quad h' = h, p \mapsto c[f_i \mapsto v_i \quad i \in 1..n]}{\sigma_1 \vdash a_1 \dots a_n \text{ newobj } K \rightsquigarrow p \cdot (h', s)}$

(Eval callvirt) (where $M = B c::\ell(A_1, \dots, A_n)$)
$\frac{\sigma_0 \vdash a_0 \rightsquigarrow p_0 \cdot (h_1, s_1) \quad h_1(p_0) = c[f_i \mapsto u_i \quad i \in 1..m] \quad (h_i, s_i) \vdash a_i \rightsquigarrow v_i \cdot (h_{i+1}, s_{i+1}) \quad \forall i \in 1..n \quad \text{methods}(c')(B \ell(A_1, \dots, A_n)) = b \quad (h_{n+1}, s_{n+1}. \text{args}(p_0, v_1, \dots, v_n)) \vdash b \rightsquigarrow v' \cdot (h', s' \text{ fr}')}{\sigma_0 \vdash a_0 a_1 \dots a_n \text{ callvirt } M \rightsquigarrow v' \cdot (h', s')}$

The expression $a_1 \dots a_n \text{ newobj } K$, where K is the constructor for a class $c \notin \text{ValueClass}$, allocates a boxed object whose fields contain the results of evaluating a_1, \dots, a_n , and returns the new reference.

The expression $a_0 a_1 \dots a_n \text{ callvirt } M$, where M refers to $B \ell(A_1, \dots, A_n)$ in class c , evaluates a_0 to a reference to a boxed object of class c' (expected to inherit from c), locates the method body for $B \ell(A_1, \dots, A_n)$ in class c' , and returns the result of evaluating this method body in a new stack frame whose argument vector consists of the reference to the boxed object (the self pointer) together with the results of a_1, \dots, a_n . The result of this evaluation is the store $(h', s' \text{ fr}')$, where fr' is the final state of the new stack frame. Once evaluation of the method is complete, the stack is popped, to leave (h', s') as the final store.

Evaluation Rules for Reference and Value Types:

(Eval ldflda)
$\frac{\sigma \vdash a \rightsquigarrow \text{ptr} \cdot \sigma'}{\sigma \vdash a \text{ ldflda } A c::f \rightsquigarrow \text{ptr}.f \cdot \sigma'}$

(Eval stfld)
$\frac{\sigma \vdash a \rightsquigarrow \text{ptr} \cdot \sigma' \quad \sigma' \vdash b \rightsquigarrow v \cdot \sigma''}{\sigma \vdash a b \text{ stfld } A c::f \rightsquigarrow 0 \cdot \text{update}(\sigma'', \text{ptr}.f, v)}$

The expression `$a \text{ ldflda } A c::f$` evaluates a to a pointer to a boxed or unboxed object, then returns a pointer to field f of this object.

The expression `$a b \text{ stfld } A c::f$` evaluates a to a pointer to a boxed or unboxed object, updates its field f with the result of evaluating b , and returns void.

Evaluation Rules for Value Types Only:

(Eval newobj) (where $K = \text{void } vc::\text{ctor}(A'_1, \dots, A'_m)$)
$\frac{\text{fields}(vc) = f_i \mapsto A_i \quad i \in 1..n \quad \sigma_i \vdash a_i \rightsquigarrow v_i \cdot \sigma_{i+1} \quad \forall i \in 1..n}{\sigma_1 \vdash a_1 \dots a_n \text{ newobj } K \rightsquigarrow (f_i \mapsto v_i \quad i \in 1..n) \cdot \sigma_{n+1}}$

(Eval call) (where $M = B vc::\ell(A_1, \dots, A_n)$)
$\frac{\sigma_0 \vdash a_0 \rightsquigarrow \text{ptr} \cdot (h_1, s_1) \quad (h_i, s_i) \vdash a_i \rightsquigarrow v_i \cdot (h_{i+1}, s_{i+1}) \quad \forall i \in 1..n \quad \text{methods}(vc)(B \ell(A_1, \dots, A_n)) = b \quad (h_{n+1}, s_{n+1}. \text{args}(\text{ptr}, v_1, \dots, v_n)) \vdash b \rightsquigarrow v' \cdot (h', s' \text{ fr}')}{\sigma_0 \vdash a_0 a_1 \dots a_n \text{ call instance } M \rightsquigarrow v' \cdot (h', s')}$

$$\frac{\text{(Eval box) (where } p \notin \text{dom}(h'))}{\sigma \vdash a \rightsquigarrow ptr \cdot (h', s') \quad \text{lookup}((h', s'), ptr) = f_i \mapsto v_i \text{ }^{i \in 1..n}}}{\sigma \vdash a \text{ box } vc \rightsquigarrow p \cdot ((h', p \mapsto vc[f_i \mapsto v_i \text{ }^{i \in 1..n}]), s)}$$

$$\frac{\text{(Eval unbox)}}{\sigma \vdash a \rightsquigarrow p \cdot \sigma'}$$

The expression $a_1 \cdots a_n \text{ newobj } K$, where K is the constructor for a value class vc , returns an unboxed object whose fields contain the results of evaluating a_1, \dots, a_n .

The expression $a_0 a_1 \cdots a_n \text{ call instance } M$ where M refers to the signature $B \ell(A_1, \dots, A_n)$ in value class vc , evaluates a_0 to a pointer to an unboxed object (expected to be of class vc), locates the method body for $B \ell(A_1, \dots, A_n)$ in class vc , and returns the result of evaluating this method body in a new stack frame whose argument vector consists of the pointer to the unboxed object (the self pointer) together with the results of a_1, \dots, a_n .

The expression $a \text{ box } c$ evaluates a to a pointer to an unboxed object, allocates it in boxed form in the heap, and returns the fresh heap reference.

The expression $a \text{ unbox } c$ evaluates a to a heap reference to a boxed object, and returns this reference as its result.

2.4 Typing Method Bodies

This section describes a type system for method bodies such that evaluation of well-typed method bodies cannot lead to an execution error. What is perhaps most interesting here is the implementation of the Pointer Confinement Policy of Section 2.1.

Let a *type frame*, Fr , take the form $\text{.args}(A_0, \dots, A_n)$, a description of the types of the results in the current (top) stack frame. Our typing judgment, $Fr \vdash b : B$, means if the current stack frame matches Fr , the body b evaluates to a result of type B .

Type Frames and Typing Judgment:

$$\frac{Fr ::= \text{.args}(A_0, \dots, A_n) \quad \text{frame: types of arguments}}{Fr \vdash b : B \quad \text{given } Fr, \text{ body } b \text{ returns type } B}$$

We make the additional assumption about our execution environment that every method body (b below) conforms to its signature:

Additional Assumptions:

$$\frac{c \notin \text{ValueClass} \wedge \text{methods}(c)(B \ell(A_1, \dots, A_n)) = b \Rightarrow \text{.args}(\text{class } c, A_1, \dots, A_n) \vdash b : B}{\text{c} \notin \text{ValueClass} \wedge \text{methods}(c)(B \ell(A_1, \dots, A_n)) = b \Rightarrow \text{.args}(\text{value class } vc\&, A_1, \dots, A_n) \vdash b : B}$$

(Ref methods) (Val methods)

Next, we give typing rules to define $Fr \vdash b : B$.

Typing Rule for Subsumption:

$$\frac{\text{(Body Subsum)}}{Fr \vdash b : B \quad B <: B'}{Fr \vdash b : B'}$$

This standard rule allows an expression of a subtype B to be used in a context expecting a supertype B' .

Typing Rules for Control Flow:

$$\frac{\text{(Body ldc)}}{Fr \vdash \text{ldc.i4 } i4 : \text{int32}} \quad \frac{\text{(Body Seq)}}{Fr \vdash a : \text{void} \quad Fr \vdash b : B}{Fr \vdash a b : B}$$

$$\frac{\text{(Body Cond)}}{Fr \vdash a : \text{int32} \quad Fr \vdash b_0 : B \quad Fr \vdash b_1 : B}{Fr \vdash a b_0 b_1 \text{ cond} : B}$$

$$\frac{\text{(Body While)}}{Fr \vdash a : \text{int32} \quad Fr \vdash b : \text{void}}{Fr \vdash a b \text{ while} : \text{void}}$$

The rule (Body Seq) uses the type `void` to guarantee that the first part of a sequential composition returns no results.

The rules (Body Cond) and (Body While) use the type `int32` to guarantee the predicate expression a returns an integer.

Typing Rules for Pointer Types:

$$\frac{\text{(Body ldind)}}{Fr \vdash a : A\&} \quad \frac{\text{(Body stind) (where } pointerFree(A))}{Fr \vdash a_1 : A\& \quad Fr \vdash a_2 : A}{Fr \vdash a_1 a_2 \text{ stind} : \text{void}}$$

The rule (Body `stind`) implements rule (3) of the Pointer Confinement Policy; without the condition $pointerFree(A)$, `stind` could copy a pointer to the current stack frame further back the stack.

Typing Rules for Arguments:

$$\frac{\text{(Body ldarga)}}{j \in 0..n}{\text{.args}(A_0, \dots, A_n) \vdash \text{ldarga } j : A_j\&}$$

$$\frac{\text{(Body starga)}}{\text{.args}(A_0, \dots, A_n) \vdash a : A_j \quad j \in 0..n}{\text{.args}(A_0, \dots, A_n) \vdash a \text{ starga } j : \text{void}}$$

These rules check that the argument index j exists. Since `starga` only writes within the current frame, we can safely allow A_j to be a pointer.

Typing Rules for Reference Types:

$$\frac{\text{(Ref newobj) (where } K = \text{void } c::\text{ctor}(A_1, \dots, A_n) \text{ and } fields(c) = f_i \mapsto A_i \text{ }^{i \in 1..n})}{Fr \vdash a_i : A_i \quad \forall i \in 1..n \quad c \notin \text{ValueClass}}{Fr \vdash a_1 \cdots a_n \text{ newobj } K : \text{class } c}$$

$$\frac{\text{(Ref callvirt) (where } B \ell(A_1, \dots, A_n) \in \text{dom}(\text{methods}(c))}{Fr \vdash a_0 : \text{class } c \quad Fr \vdash a_i : A_i \quad \forall i \in 1..n}{Fr \vdash a_0 a_1 \cdots a_n \text{ callvirt } B c::\ell(A_1, \dots, A_n) : B}$$

$$\frac{\text{(Ref ldflda) (where } fields(c) = f_i \mapsto A_i \text{ }^{i \in 1..n})}{Fr \vdash a : \text{class } c \quad j \in 1..n}{Fr \vdash a \text{ ldflda } A_j c::f_j : A_j\&}$$

$$\frac{\text{(Ref stfld) (where } fields(c) = f_i \mapsto A_i \text{ }^{i \in 1..n} \text{ and } pointerFree(A_j))}{Fr \vdash a : \text{class } c \quad Fr \vdash b : A_j \quad j \in 1..n}{Fr \vdash a b \text{ stfld } A_j c::f_j : \text{void}}$$

These are fairly standard rules for operations on boxed objects. Recall that the axiom (Good *fields*) guarantees every field is pointer-free. So the *pointerFree*($-$) condition on the rule (Ref *stfld*) is redundant. Still, it is not redundant in a variation of our type system considered in Section 3, that allows value classes to include pointers.

Typing Rules for Value Types:

$\frac{\text{(Val newobj)} \quad (\text{where } K = \text{void } vc::\text{ctor}(A_1, \dots, A_n) \text{ and } \text{fields}(vc) = f_i \mapsto A_i^{i \in 1..n})}{Fr \vdash a_i : A_i \quad \forall i \in 1..n}$ $Fr \vdash a_1 \dots a_n \text{ newobj } K : \text{value class } vc$
$\text{(Val call)} \quad (\text{where } B \ell(A_1, \dots, A_n) \in \text{dom}(\text{methods}(vc)))$ $\frac{Fr \vdash a_0 : \text{value class } vc\& \quad Fr \vdash a_i : A_i \quad \forall i \in 1..n}{Fr \vdash a_0 a_1 \dots a_n \text{ call instance } B \text{ } vc::\ell(A_1, \dots, A_n) : B}$
$\text{(Val ldflda)} \quad (\text{where } \text{fields}(vc) = f_i \mapsto A_i^{i \in 1..n})$ $\frac{Fr \vdash a : \text{value class } vc\& \quad j \in 1..n}{Fr \vdash a \text{ ldflda } A_j \text{ } vc::f_j : A_j\&}$
$\text{(Val stfld)} \quad (\text{where } \text{fields}(vc) = f_i \mapsto A_i^{i \in 1..n} \text{ and } \text{pointerFree}(A_j))$ $\frac{Fr \vdash a : \text{value class } vc\& \quad Fr \vdash b : A_j \quad j \in 1..n}{Fr \vdash a \text{ b stfld } A_j \text{ } vc::f_j : \text{void}}$
$\text{(Val box)} \quad (\text{where } \text{pointerFree}(\text{value class } vc))$ $\frac{Fr \vdash a : \text{value class } vc\&}{Fr \vdash a \text{ box } vc : \text{class } vc}$
(Val unbox) $\frac{Fr \vdash a : \text{class } vc}{Fr \vdash a \text{ unbox } vc : \text{value class } vc\&}$

These are similar to the typing rules for operations on boxed objects, except we refer to the object via a pointer type instead of a reference type. Like (Ref *stfld*), the rules (Val *stfld*) and (Val *box*) bear *pointerFree*($-$) conditions that are redundant in the current system, but not in the system of Section 3.

2.5 Typing the Memory Model

In this section, we present predicates, known as conformance judgments, that confer types on our memory model. In the next, we show that these predicates are invariants of computation, that is, are preserved by method evaluation.

We begin by introducing types for the components of our memory model. A *heap type* $p_i \mapsto c_i^{i \in 1..n}$ determines the actual class of each boxed object. A *stack type* $Fr_1 \dots Fr_n$ determines frame types for each frame in the stack. A *store type* $\Sigma = (H, S)$ determines a heap type H and stack type S .

Heap, Stack, and Store Types:

$H ::= p_i \mapsto c_i^{i \in 1..n}$	heap type
$S ::= Fr_1 \dots Fr_n$	stack type
$\Sigma ::= (H, S)$	store type

Our first conformance judgment, $\Sigma \models u : A$, means that in a store matching the store type Σ , the result u is well-formed and has type A . We define what it means for a store

to match a store type through other conformance judgments, defined later.

Conformance Judgment for Results (Including Pointers):

$\Sigma \models u : A$	in Σ , result u has type A
------------------------	---------------------------------------

Conformance Rules for References and Pointers:

$\text{(Res Ref)} \quad \frac{H(p) = c \quad c \text{ inherits } c'}{(H, S) \models p : \text{class } c'}$	$\text{(Ptr Ref)} \quad \frac{H(p) = vc}{(H, S) \models p : \text{value class } vc\&}$
$\text{(Ptr Arg)} \quad \frac{i \in 1..m \quad Fr_i = \text{.args}(A_0, \dots, A_n) \quad j \in 0..n}{(H, Fr_1 \dots Fr_m) \models (i, j) : A_j\&}$	
$\text{(Ptr Field)} \quad (\text{where } A = \text{class } c \text{ or } A = \text{value class } c\&)$ $\frac{\Sigma \models ptr : A \quad \text{fields}(c) = f_i \mapsto A_i^{i \in 1..n} \quad j \in 1..n}{\Sigma \models ptr.f_j : A_j\&}$	

The rule (Res Ref) assigns a reference type *class* c' to a heap reference p , so long as c' is a superclass of the actual class of the object referred to by p .

The rule (Ptr Ref) assigns a pointer type to a heap reference p that refers to a value that is boxed on the heap.

These two rules can assign both a reference type and a pointer type to a heap reference to a value class. If $H(p) = vc$, then we have $(H, S) \models p : \text{class } c$ by (Res Ref), but also $(H, S) \models p : \text{value class } c\&$ by (Ptr Ref). We need (Res Ref) to type references constructed by the *box* instruction. We need (Ptr Ref) to type pointers constructed by the *unbox* instruction.

The rule (Ptr Arg) assigns a pointer type to a stack pointer (i, j) that refers to argument j of frame i .

The rule (Ptr Field) assigns a pointer type to a pointer referring to the field f_j of the object referred to by ptr . The base pointer ptr may either be of type *class* c or *value class* $c\&$. The first case is needed for a pointer to a field of a heap object that is not in a value class. The second case is needed for a pointer to a field of a heap or stack object in a value class.

Conformance Rules for Other Results:

$\text{(Res Void)} \quad \frac{}{\Sigma \models 0 : \text{void}}$	$\text{(Res Int)} \quad \frac{}{\Sigma \models i4 : \text{int32}}$
$\text{(Res Value)} \quad \frac{\text{fields}(vc) = f_i \mapsto A_i^{i \in 1..n} \quad \Sigma \models v_i : A_i \quad \forall i \in 1..n}{\Sigma \models f_i \mapsto v_i^{i \in 1..n} : \text{value class } vc}$	

The rules (Res Void) and (Res Int) assign the *void* and *int32* types to void and integer values, respectively.

The rule (Res Value) assigns a value type *value class* vc to a value. By axiom (Hi Val), the inheritance hierarchy is flat for value types. So (Res Value), unlike (Res Ref), does not allow vc to be a proper superclass of the actual class of the value.

Other Conformance Judgments:

$H \models o : c$	in H , object o has class c
$H \models h$	heap h conforms to H

$\Sigma \models fr : Fr$	frame fr conforms to Fr
$\Sigma \models \sigma$	store σ conforms to Σ

Conformance Rule for Objects:

(Con Object) (where $fields(c) = f_i \mapsto A_i \ i \in 1..n$)
$(H, \emptyset) \models v_i : A_i \ \forall i \in 1..n$
$H \models c[f_i \mapsto v_i \ i \in 1..n] : c$

This rule defines when a heap object $c[f_i \mapsto v_i \ i \in 1..n]$ is well-typed. The preconditions $(H, \emptyset) \models v_i : A_i$ require that the fields v_i be typed with an empty stack type. It follows that no field v_i contains a stack pointer, since the rule (Ptr Arg) for typing stack pointers assumes a non-empty stack type.

Conformance Rule for Heaps:

(Con Heap) (where $H = p_i \mapsto c_i \ i \in 1..n$)
$H \models o_i : c_i \ \forall i \in 1..n$
$H \models p_i \mapsto o_i \ i \in 1..n$

This rule defines when a heap $p_i \mapsto o_i \ i \in 1..n$ conforms to the heap type $p_i \mapsto c_i \ i \in 1..n$. The heap type contains the actual class c_i of each object o_i .

Conformance Rule for Frames:

(Con Frame)
$\Sigma \models u_i : A_i \ \forall i \in 0..n$
$\Sigma \models .args(u_0, \dots, u_n) : .args(A_0, \dots, A_n)$

This rule defines when a frame conforms to a frame type.

Conformance Rule for Stores:

(Con Store)
$H \models h \quad (H, Fr_1 \dots Fr_n) \models fr_i : Fr_i \ \forall i \in 1..n$
$(H, Fr_1 \dots Fr_n) \models (h, fr_1 \dots fr_n)$

This rule defines when a store $(H, Fr_1 \dots Fr_n)$ conforms to a store type $(h, fr_1 \dots fr_n)$. It asks that the heap h conform to the heap type H , and that each stack frame fr_i conform to the corresponding frame type Fr_i , but after removing from the store type any higher—shorter lived—stack frames. Hence, there may be pointers from a higher to a lower stack frame, but not the other way round.

2.6 Evaluation Respects Typing

We use standard proof techniques to show the consistency of the BIL evaluation semantics with its type system. The following is the main type safety result of the paper. If a program satisfies the restrictions on type structure imposed in Section 2.1 and the typing rules for method bodies in Section 2.4 then its evaluation according to the rules in Section 2.3 can lead only to conformant intermediate states as defined in Section 2.5. Let $H \leq H'$ mean that $dom(H) \subseteq dom(H')$ and $H(p) = H'(p)$ for all $p \in dom(H)$.

Theorem 1 *If $(H, S Fr) \models \sigma$ and $Fr \vdash b : B$ and $\sigma \vdash b \rightsquigarrow v \cdot \sigma^\dagger$ then there exists a heap type H^\dagger such that $H \leq H^\dagger$ and $(H^\dagger, S Fr) \models v : B$ and $(H^\dagger, S Fr) \models \sigma^\dagger$.*

Proof By induction on the derivation of $\sigma \vdash b \rightsquigarrow v \cdot \sigma^\dagger$. We omit the details. See the Appendix for the main lemmas about the type system needed in the proof. \square

As usual, such a theorem is vacuous if there is no σ^\dagger such that $\sigma \vdash b \rightsquigarrow v \cdot \sigma^\dagger$ holds, which happens either because the computation would diverge, or because it gets stuck (if there is no applicable evaluation rule). Stuck states correspond to execution errors, such as calling a non-existent method, or attempting to de-reference an integer or a dangling pointer. As discussed by Abadi and Cardelli [AC96], we conjecture it would be straightforward to adapt the proof of Theorem 1 to show that no stuck state is reachable.

3 Variation: Allowing Pointers in Fields of Value Classes

To avoid dangling pointers, the IL type system prevents the fields of all objects, whether boxed on the heap or unboxed on the stack, from holding pointers. In fact, as pointed out by Fergus Henderson, a more liberal type system that allows unboxed objects to contain pointers is useful for compiling nested functions.

When compiling a language with nested functions (for example, Pascal or Ada), each invocation of a nested function needs access to the activation records (that is, the arguments and local variables) of the lexically enclosing functions. A standard technique is to pass the function a display [ASU86], an array of pointers to these activation records. One strategy is to implement an activation record (containing those arguments and local variables referred to by nested functions) as a value class on the stack, and to implement the display by pointers to the value classes representing the activation records. Since arguments may be passed by reference, this scheme works only if we allow value classes to hold pointers. Otherwise, we need to pay the cost of boxing these activation records on the heap.

If we allow fields of value classes to hold pointers, the following more liberal policy still avoids dangling pointers.

A More Liberal Pointer Confinement Policy:

- (1) No field of a boxed object may hold a pointer.
- (2) No method may return a result containing a pointer.
- (3) No result containing a pointer may be stored indirectly via another pointer.

Though this policy helps compile nested functions, we lose the possibly useful fact that every value class may be boxed, and hence treated as a subtype of `class System.Object`.

To formalize this policy, we amend BIL as follows.

- Change the definition of $pointerFree(A)$ to hold if and only if (1) A is not itself a pointer type, and (2) if A is a value class then the type B of each field satisfies $pointerFree(B)$. (The only change is the insertion of clause (2).)
- Change axiom (Good fields) to read: $c \notin ValueClass \Rightarrow pointerFree(fields(c)(f))$. (The only change is the insertion of the $c \notin ValueClass$ precondition.)

To see the effect of these changes, recall there are four typing rules that mention the $pointerFree(-)$ predicate: (Ref stfld), (Body stind), (Val stfld), and (Val box).

Previously, any value could be stored via (Body stind), and the pointer-free conditions on the other three rules were

redundant. Now, these rules prevent the export of values containing pointers to the heap or further back the stack. Now, (Ref stfld) prevents a pointer being stored into a boxed value class with a pointer field. In fact, no such boxed value classes can even be allocated, given the *pointerFree*(-) condition on (Val box).

Our proof of Theorem 1, outlined in the Appendix, is in fact for this more liberal system. Type safety for the original system is a corollary of type safety for this more liberal system, since any method body typed by the original system remains typable.

Implementation of the new scheme remains future work.

4 IL Features Omitted From BIL

To give a flavour of the full intermediate language, we briefly enumerate the main features omitted from BIL. The IL Assembly Programmer's Reference Manual [Mic00] contains a complete informal description of IL.

We omit all discussion of IL metadata, such as how classes, static data and method headers are described. We omit any discussion of the on-disk format, the specification of linkage information, and assemblies, the unit of software deployment.

Our object model omits null objects, global fields and methods, static fields and methods, non-virtual methods, single dimensional and multidimensional covariant arrays, and object interfaces. Our instruction set omits local variables, arithmetic instructions, arbitrary branching, jumping, and tail calls. Tail calls require care, because the type system must prevent pointers to the current stack frame being passed as arguments. The current IL policy is to prevent the passing of any pointers via a tail call.

We omit delegates (that is, built-in support for anonymous method invocation), typed references (that is, a pointer packaged with its type, required for Visual Basic), attributes, native code calling conventions, interoperability with COM, remoting (object distribution) and multi-threading. We also omit exception handling, a fairly elaborate model that permits a unified view of exceptions in C++, C#, and other high-level languages.

5 Related Work

The principle of formalizing type-checking via logical inference rules is a long-standing topic in the study of programming languages [Car97]. Formal typing rules have been developed for several high-level languages, including SML [MTHM97], Haskell [PW92], and for subsets of Java [DE97, IPW99]. Formal typing rules have also been developed for several low-level languages, including TAL [MWCG99] and for subsets of the JVM [SA98, Qia99, Yel99, FM00]. The properties established by proof-carrying code [Nec97] can be viewed as typing derivations for native code. The idea of formalizing a type system via an executable type-checker has recently been advocated for Haskell [Jon99]. Our use of an executable specification as an oracle is an instance of the standard software engineering principle of multi-version prototyping. Proofs of soundness of several programming language type systems have been partially mechanised in theorem provers [Van96, Nor98, Sym99, vN99].

Several compilers, such as GHC [PHH⁺93], TIL [TMC⁺96], FLINT [Sha97], and MARMOT [FKR⁺00],

use a typed intermediate language internally. One [MWCG99] in particular translates all the way from System F, a polymorphic λ -calculus, down to a typed assembly language, TAL. The idea of writing a type-checker for a textual assembly format (like our type-checker for IL) appears in connection with TAL: the TALx86 type-checker accepts input in a typed form of the IA32 assembly language that can also be processed by the standard MASM assembler.

Reference types for heap-allocated data structures akin to the reference types of the type system of Section 2 appear in all of these intermediate languages. What is new about our type system is its inclusion of value and pointer types.

- Value types describe the unboxed stack-allocated form of a class. The box and unbox instructions coerce between stack and heap forms of a class. Types for boxed and unboxed non-strict data structures [PL91] and automatic type-based coercions between boxed and unboxed forms [Ler92] have been studied previously. Other approaches include region analysis [TT97] and escape analysis [PG92]. Still, the idea and formalization of types to differentiate between unboxed and boxed forms of class-based objects appears to be new.
- Pointer types describe pointers to either stack or heap allocated items. A risk with a stack pointer is that it may dangle, if its lifetime exceeds the lifetime of its target. The stack-based form of TAL [MCGW98] includes a type constructor for describing pointers into the stack; the parameter to the type constructor is a stack type that ensures the target is still live when the pointer is dereferenced. Instead, the Pointer Confinement Policy of Section 2 avoids dangling pointers via various syntactic restrictions. IL's pointer types are easier to integrate with high-level languages like Visual Basic with rather simple type systems than a more sophisticated solution using stack types, as found in TAL.

6 Conclusions

One of the innovations in Microsoft's Common Language Runtime is support for typed stack pointers, for passing arguments and results by reference, for example. We presented formal typing rules and a type safety result for a substantial fragment of the Common Language Runtime intermediate language. Our treatment of value types and pointer types appears to be new. These rules were devised through our writing informal and executable specifications of the full intermediate language. This effort clarified the design and helped find bugs, but further research is needed on machine support for formal reasoning and on test case generation. We exploited our formal model to validate a liberalisation of the IL policy that allows object fields to contain stack pointers.

Acknowledgements This paper is one outcome of a collaborative project between Microsoft Research Cambridge and the .NET Common Language Runtime product team. We thank the Common Language Runtime team, and especially S. Dasan, J. Forbes, J. Miller, and V. Morrison, for their help during our collaboration. We thank our research interns P. Hankin, J. Matthews, and C. Pitcher for their work on the project. M. Abadi, C.A.R. Hoare, A. Jeffrey, and C. Skalka suggested improvements to a draft of this paper.

References

- [AC96] M. Abadi and L. Cardelli. *A Theory of Objects*. Springer Verlag, 1996.
- [ASU86] A.V. Aho, R. Sethi, and J.D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, 1986.
- [Car97] L. Cardelli. Type systems. In A.B. Tucker, editor, *The Computer Science and Engineering Handbook*, chapter 103, pages 2208–2236. CRC Press, 1997.
- [Coh89] A. Cohn. The notion of proof in hardware verification. *Journal of Automated Reasoning*, 5(2):127–139, June 1989.
- [DE97] S. Drossopoulou and S. Eisenbach. Java is type safe—probably. In *Proceedings ECOOP'97*, June 1997.
- [FKR⁺00] R. Fitzgerald, T.B. Knoblock, E. Ruf, B. Steensgaard, and D. Tarditi. Marmot: An optimizing compiler for Java. *Software: Practice and Experience*, 30(3), 2000.
- [FM00] S. Freund and J.C. Mitchell. A type system for object initialization in the Java bytecode language. *ACM Transactions on Programming Languages and Systems*, 2000. To appear.
- [GHL99] A.D. Gordon, P.D. Hankin, and S.B. Lassen. Compilation and equivalence of imperative objects. *Journal of Functional Programming*, 9(4):373–426, 1999.
- [GS00] A.D. Gordon and D. Syme. Typing a multi-language intermediate code. Technical Report MSR-TR-2000-106, Microsoft Research, 2000.
- [HW00] A. Hejlsberg and S. Wiltamuth. C# Language Reference. Available at <http://msdn.microsoft.com/vstudio/nextgen/technology/csharpintro.asp>, 2000.
- [IPW99] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, October 1999.
- [Jon99] M.P. Jones. Typing Haskell in Haskell. In *Proceedings Haskell Workshop, Paris*, 1999. Available at <http://www.cse.ogi.edu/~mpj/thih>.
- [Ler92] X. Leroy. Unboxed objects and polymorphic typing. In *19th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 177–188. ACM Press, 1992.
- [MCGW98] G. Morrisett, K. Crary, N. Glew, and D. Walker. Stack-based typed assembly language. In *Workshop on Types in Compilation*, volume 1473 of *Lecture Notes in Computer Science*, pages 28–52. Springer Verlag, 1998.
- [Mic00] Microsoft Corporation. *Microsoft IL Assembly Programmer's Reference Manual*, July 2000. Part of the *.NET Framework Software Development Kit*, distributed on CD at the Microsoft Professional Developers Conference, Orlando, Florida, July 11–14, 2000.
- [MTHM97] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [MWCG99] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):528–569, 1999.
- [Nec97] G. Necula. Proof-carrying code. In *24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 106–119. ACM Press, 1997.
- [Nor98] M. Norrish. *C formalised in HOL*. PhD thesis, University of Cambridge, 1998.
- [PG92] Y.G. Park and B. Goldberg. Escape analysis on lists. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 116–127. ACM Press, 1992.
- [PHH⁺93] S. Peyton Jones, C. Hall, K. Hammond, W. Partain, and P. Wadler. The Glasgow Haskell compiler: a technical overview. In *Proceedings UK Joint Framework for Information Technology (JFIT) Technical Conference*, pages 249–257. 1993.
- [PL91] S. Peyton Jones and J. Launchbury. Unboxed values as first class citizens. In *Functional Programming Languages and Computer Architecture*, volume 523 of *Lecture Notes in Computer Science*, pages 636–666. Springer Verlag, 1991.
- [PW92] S. Peyton Jones and P. Wadler. A static semantics for Haskell. Unpublished draft, Department of Computing Science, University of Glasgow. Available at <http://research.microsoft.com/users/simonpj>, 1992.
- [Qia99] Z. Qian. A formal specification of JavaTM virtual machine instructions for objects, methods and subroutines. In J. Alves-Foss, editor, *Formal Syntax and Semantics of Java*, volume 1532 of *Lecture Notes in Computer Science*, pages 271–312. Springer Verlag, 1999.
- [SA98] R. Stata and M. Abadi. A type system for Java bytecode subroutines. In *Proceedings POPL'98*, pages 149–160. ACM Press, 1998.
- [Sha97] Z. Shao. An overview of the FLINT/ML compiler. In *Proc. 1997 ACM SIGPLAN Workshop on Types in Compilation (TIC'97)*, Amsterdam, The Netherlands, June 1997.
- [Sym98] D. Syme. *Declarative Theorem Proving for Operational Semantics*. PhD thesis, University of Cambridge, 1998.

- [Sym99] D. Syme. Proving Java type soundness. In J. Alves-Foss, editor, *Formal Syntax and Semantics of Java*, volume 1532 of *Lecture Notes in Computer Science*, pages 83–119. Springer Verlag, 1999.
- [TMC⁺96] D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. TIL: A type-directed optimizing compiler for ML. In *Proc. PLDI'96*, pages 181–192, 1996.
- [TT97] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [Van96] M. VanInwegen. *The Machine-Assisted Proof of Programming Language Properties*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, May 1996.
- [vN99] D. von Oheimb and T. Nipkow. Machine-checking the Java specification: Proving type-safety. In J. Alves-Foss, editor, *Formal Syntax and Semantics of Java*, volume 1532 of *Lecture Notes in Computer Science*, pages 119–156. Springer Verlag, 1999.
- [Yel99] P.M. Yelland. A compositional account of the JavaTM Virtual Machine. In *26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 57–69. ACM Press, 1999.

A Facts Needed in the Proof of Theorem 1

This appendix encompasses the main lemmas needed in the proof of the main type safety theorem of the paper. Our proofs are for the definitions of (Good fields) and *pointerFree*(–) described in Section 3. The proofs can trivially be adapted for the original definitions in Section 2. Appendix A.1 covers basic lemmas about the subtype and conformance relations. Appendix A.2 presents an alternative characterisation of the pointer conformance judgement $\Sigma \models ptr : A\&$. Finally, Appendix A.3 presents the definitions and typing properties of the store lookup and update functions.

A.1 Basic Lemmas

We begin with two lemmas about the subtype relation. Subtyping is trivial for all types except reference types. Only reference types can be supertypes of other reference types.

Lemma 1 *Assume $B \neq \text{class } c$ for all c . If $A <: B$ or $B <: A$ then $A = B$.*

Lemma 2 *If $\text{class } c <: A$ then there exists c' such that $A = \text{class } c'$ and c inherits c' .*

Although a subsumption rule is not part of the definition of the result conformance relation $\Sigma \models v : A$, it is derivable.

Lemma 3 *If $\Sigma \models v : A$ and $A <: A'$ then $\Sigma \models v : A'$.*

The next three lemmas concern how varying the size of the stack affects conformance.

Lemma 4 states that a pointer-free result well-formed in a store type (H, S) is also well-formed in the store type (H, \emptyset) . This justifies moving pointer-free results from the current frame to the heap.

Lemma 5 states that any result well-formed in a store type (H, S) is also well-formed in the store type $(H, S Fr)$. This justifies passing results from the current frame into the frame of a called method.

Lemma 6 states that a pointer-free result well-formed in a store type $(H, S Fr)$ is also well-formed in the store type (H, S) . This justifies returning pointer-free results from a called frame to the previous frame.

Lemmas 4 and 6 do not apply to pointer results because if the result is a pointer into the top stack frame it is not well-formed in a smaller stack.

Lemma 4 *If $(H, S) \models v : A$ and $\text{pointerFree}(A)$ then $(H, \emptyset) \models v : A$.*

Lemma 5 *If $(H, S) \models v : A$ then $(H, S Fr) \models v : A$.*

Lemma 6 *If $(H, S Fr) \models v : A$ and $\text{pointerFree}(A)$ then $(H, S) \models v : A$.*

Next, we have two lemmas concerned with method call and return.

Lemma 7 says that a frame is well-formed in the store $(H, S Fr)$ if it is well-formed in the store (H, S) . This justifies passing an argument frame to a called method.

Lemma 8 says that a store (h, s) conforms to the store type (H, S) if the store $(h, s fr)$ conforms to a store type $(H, S Fr)$. This justifies returning from a method.

The proof of Lemma 8 depends on showing that no pointer in the final store (h, s) refers to the frame *fr*.

Lemma 7 *If $(H, S) \models fr : Fr$ then $(H, S Fr) \models fr : Fr$.*

Lemma 8 *If $(H, S Fr) \models (h, s fr)$ then $(H, S) \models (h, s)$.*

Recall that we state Theorem 1 in terms of a relation $H \leq H'$ defined to mean that $\text{dom}(H) \subseteq \text{dom}(H')$ and $H(p) = H'(p)$ for all $p \in \text{dom}(H)$. We may call this the *heap extension* relation. Heap extension is a partial order.

Lemma 9 *The relation $H \leq H'$ is reflexive and transitive (that is, for all H, H' , and H'' , $H \leq H$, and, if $H \leq H'$ and $H' \leq H''$ then $H \leq H''$).*

The next three lemmas state that heap extension preserves the conformance relations for results, objects, and frames.

Lemma 10 *If $(H, S) \models v : A$ and $H \leq H'$ then $(H', S) \models v : A$.*

Lemma 11 *If $H \models o : c$ and $H \leq H'$ then $H' \models o : c$.*

Lemma 12 *If $(H, S) \models fr : Fr$ and $H \leq H'$ then $(H', S) \models fr : Fr$.*

The final lemma of this section justifies boxing of results. If the heap h and the object o both conform to the heap type H , and p is a fresh reference, then the extended heap obtained by allocating o at p is well-formed.

Lemma 13 *If $H \models h$ and $p \notin \text{dom}(h)$ and $H \models o : c$ then $H, p \mapsto c \models h, p \mapsto o$.*

A.2 Another Formulation of Pointer Conformance

In the next section we present the recursive definitions of the *lookup* and *update* functions on pointers. To show properties of these functions, it is convenient to present in this section a reformulation of the pointer conformance relation $\Sigma \models ptr : A\&$. Essentially, we show that every well-formed pointer takes the form of either (1) a pointer to an argument in a frame, followed by a possibly empty path of field selections, or (2) a reference to a boxed object of a value class, followed by a possibly empty path of field selections, or (3) a reference to a boxed object (not necessarily of a value class) followed by a non-empty path of field selections.

This reformulation begins with a notion of a path, a possibly empty sequence of field names.

Path Within an Object:

$$\vec{f} ::= f_1 \cdots f_n \quad \text{sequence of fields (written } \epsilon \text{ if } n = 0)$$

Next, we define a relation $A \xrightarrow{\vec{f}} B$ to mean that either the sequence \vec{f} is empty and $A = B$, or that A is a value class, and selecting the fields in the series \vec{f} in order yields the type B . This is defined in terms of $A \xrightarrow{f} B$, an auxiliary single step relation.

Actions of Fields on Types: $A \xrightarrow{f} B$ and $A \xrightarrow{\vec{f}} B$

$$\begin{aligned} A \xrightarrow{f} B & \text{ if and only if } A = \text{value class } vc \text{ and} \\ & \text{fields}(vc) = f_i \mapsto A_i \text{ }^{i \in 1..n} \text{ and } f = f_j \text{ and } B = A_j. \\ A \xrightarrow{\epsilon} B & \text{ if and only if } A = B. \\ A \xrightarrow{f_1 \cdots f_n} B & \text{ if and only if } A \xrightarrow{f_1} \cdots \xrightarrow{f_n} B \text{ (where } n > 0). \end{aligned}$$

Given these notations, we reformulate pointer conformance as follows.

Lemma 14 *The judgment $\Sigma \models ptr : A\&$ holds if and only if either:*

- (1) *there exist (i, j) , \vec{f} , and B such that $ptr = (i, j).\vec{f}$ and $\Sigma \models (i, j) : B\&$ and $B \xrightarrow{\vec{f}} A$, or*
- (2) *there exist p , \vec{f} , and vc such that $ptr = p.\vec{f}$ and $\Sigma \models p : \text{value class } vc\&$ and $\text{value class } vc \xrightarrow{\vec{f}} A$, or*
- (3) *there exist p , f_j , \vec{f} , and c such that $ptr = p.f_j.\vec{f}$ and $\Sigma \models p : \text{class } c$ and $A_j \xrightarrow{\vec{f}} A$, where $\text{fields}(c) = f_i \mapsto A_i \text{ }^{i \in 1..n}$ and $j \in 1..n$.*

We use this lemma to prove the typing properties of store lookup and update functions stated in the next section.

A.3 Facts about Lookup and Update

We omitted the definitions of functions for store lookup $lookup(\sigma, ptr)$ and store update $update(\sigma, ptr, v')$ from the main body of the paper.

The store lookup function is defined in terms of an auxiliary function, result lookup $lookup(v, f_1 \cdots f_n)$, that given the result v , returns the outcome of applying each of the field selections f_1, \dots, f_n in turn. Here is the definition of this auxiliary function, followed by a typing lemma.

Result Lookup: $lookup(v, f_1 \cdots f_n)$

$$\begin{aligned} lookup(v, \epsilon) & \triangleq v \\ lookup(f_i \mapsto u_i \text{ }^{i \in 1..n}, f_j \vec{f}) & \triangleq lookup(u_j, \vec{f}) \quad \text{where } j \in 1..n \end{aligned}$$

Lemma 15 *If $\Sigma \models v : A$ and $A \xrightarrow{\vec{f}} B$ then $\Sigma \models lookup(v, \vec{f}) : B$.*

Next, we present the definition of store lookup, followed by a typing lemma.

Store Lookup via Pointer: $lookup(\sigma, ptr)$

$$\begin{aligned} lookup((h, s), p.\vec{f}) & \triangleq lookup(f_i \mapsto u_i \text{ }^{i \in 1..n}, \vec{f}) \\ & \text{where } h(p) = c[f_i \mapsto u_i \text{ }^{i \in 1..n}] \\ lookup((h, s), (i, j).\vec{f}) & \triangleq lookup(v_j, \vec{f}) \\ & \text{where } s = fr_1 \cdots fr_i \cdots fr_m \text{ with } i \in 1..m, \\ & \text{and } fr_i = \text{.args}(v_0, \dots, v_n) \text{ with } j \in 0..n \end{aligned}$$

Lemma 16 *If $\Sigma \models \sigma$ and $\Sigma \models ptr : A\&$ then $\Sigma \models lookup(\sigma, ptr) : A$.*

The store update function is defined in terms of an auxiliary function, result update $update(v, f_1 \cdots f_n, v')$, that given the result v , returns the outcome of updating the field indicated by the field selections f_1, \dots, f_n with the result v' . Here is the definition, together with a typing lemma.

Result Update: $update(v, f_1 \cdots f_n, v')$

$$\begin{aligned} update(v, \epsilon, v') & \triangleq v' \\ update(f_i \mapsto u_i \text{ }^{i \in 1..n}, f_j \vec{f}, v') & \triangleq \\ & (f_j \mapsto update(u_j, \vec{f}, v'), f_i \mapsto u_i \text{ }^{i \in (1..n) - \{j\}}) \quad \text{for } j \in 1..n \end{aligned}$$

Lemma 17 *If $\Sigma \models u : A$ and $A \xrightarrow{\vec{f}} B$ and $\Sigma \models v : B$ then $\Sigma \models update(u, \vec{f}, v) : A$.*

Given the previous auxiliary function, here is the definition of store update.

Store Update via Pointer: $update(\sigma, ptr, v')$

$$\begin{aligned} update((h, s), p.\vec{f}, v') & \triangleq \\ & (((h - p), p \mapsto c[update(f_i \mapsto u_i \text{ }^{i \in 1..n}, \vec{f}, v')]), s) \\ & \text{where } h(p) = c[f_i \mapsto u_i \text{ }^{i \in 1..n}] \\ update((h, s), (i, j).\vec{f}, v') & \triangleq \\ & (h, fr_1 \cdots \text{.args}(v_0, \dots, update(v_j, \vec{f}, v'), \dots, v_n) \cdots fr_m) \\ & \text{where } s = fr_1 \cdots fr_i \cdots fr_m \text{ with } i \in 1..m, \\ & \text{and } fr_i = \text{.args}(v_0, \dots, v_n) \text{ with } j \in 0..n \end{aligned}$$

Finally, we state two typing lemmas for store update. They are essential facts in the proof of type safety for BIL: the proof of Theorem 1 uses Lemma 18 and Lemma 19 to show that evaluations of *stind* and *starg*, respectively, are type safe.

Lemma 18 *If $\Sigma \models \sigma$ and $\Sigma \models ptr : A\&$ and $\Sigma \models v : A$ and $pointerFree(A)$ then $\Sigma \models update(\sigma, ptr, v)$.*

Lemma 19 *If $\Sigma \models \sigma$ and $\Sigma \models (i, j) : A\&$ and $\Sigma \models v : A$ and $\sigma = (h, fr_1 \cdots fr_i)$ then $\Sigma \models update(\sigma, (i, j), v)$.*