

Mint: Java Multi-stage Programming Using Weak Separability

Edwin Westbrook Mathias Ricken Jun Inoue Yilong Yao Tamer Abdelatif¹ Walid Taha

Rice University

{emw4,mgricken,ji2,yy3}@cs.rice.edu, eng.tamerabdo@gmail.com, taha@cs.rice.edu

Abstract

Multi-stage programming (MSP) provides a disciplined approach to run-time code generation. In the purely functional setting, it has been shown how MSP can be used to reduce the overhead of abstractions, allowing clean, maintainable code without paying performance penalties. Unfortunately, MSP is difficult to combine with imperative features, which are prevalent in mainstream languages. The central difficulty is scope extrusion, wherein free variables can inadvertently be moved outside the scopes of their binders. This paper proposes a new approach to combining MSP with imperative features that occupies a “sweet spot” in the design space in terms of how well useful MSP applications can be expressed and how easy it is for programmers to understand. The key insight is that escapes (or “anti-quotes”) must be *weakly separable* from the rest of the code, i.e. the computational effects occurring inside an escape that are visible outside the escape are guaranteed to not contain code. To demonstrate the feasibility of this approach, we formalize a type system based on Lightweight Java which we prove sound, and we also provide an implementation, called Mint, to validate both the expressivity of the type system and the effect of staging on the performance of Java programs.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; D.3.3 [Programming Languages]: Language Constructs and Features

General Terms Languages

Keywords Multi-staged languages, Multi-stage programming, Type systems, Java

1. Introduction

Abstraction mechanisms, such as reflection and design patterns, are useful for writing clean, maintainable code. Often, however, such mechanisms come with a steep performance overhead, making them less useful in real systems. Our goal is to allow software developers to use such abstractions, but have them be executed in an efficient manner. One approach to this problem is multi-stage programming (MSP), a language feature that provides a disciplined form of runtime code generation. Just by inserting staging annotations, a form of quasi-quotation, the developer can change

programs that use expensive abstractions into program *generators* which generate programs without the abstractions. This reduces the runtime cost of abstractions, because overhead is only paid when the generators are executed, not at the time the programs they generate are run.

A key issue for MSP is type safety, which ensures statically that all programs generated at runtime will be well-formed. Unfortunately, although it is known how to ensure type safety for MSP in the purely functional setting [4, 25, 26], it remained an open challenge how to extend this guarantee to mainstream languages such as Java. In particular, standard features of mainstream languages, such as imperative assignment, can lead to *scope extrusion*, in which variables in code fragments may move out of the scopes where they are defined. Several approaches to this problem have been proposed [1, 10, 11, 13] which give the expert MSP user fine-grained control over scoping in code; however, there is still a need for a type system that makes MSP accessible to general programmers. Mint provides a type system that is well-suited for the prevalent culture of programming in languages like Java. Specifically, Mint allows imperative generators and does not require a functional programming style, a style that often is inconvenient in Java. The type system remains simple while being more expressive than any previous type system that statically ensures type safety of generated code.

Contributions

To make MSP accessible to programmers in mainstream languages, we propose a new approach to type-safe MSP that we argue occupies a “sweet spot” in the design space in terms of how well useful MSP applications can be expressed and how easy it is for programmers to understand. Our contributions include:

- After a brief introduction to staging in Java (Section 2), we analyze and explain why scope extrusion can arise with naive approaches to MSP in Java or in similar languages (Section 3).
- We introduce the notion of *weak separability* as a solution to the scope extrusion problem (Section 4). Weak separability ensures that any effects that can be observed outside escaped expressions will not involve code objects. We present a type system that realizes this idea, and show that it is sufficient to prevent scope extrusion.
- We demonstrate the expressivity of the type system using a number of small examples illustrative of important classes of programs for which MSP can be useful (Section 5). The examples show the use of MSP for building interpreters, numerical codes, and reflective programs. The examples also emphasize that the type system allows expressing imperative generators that, for example, throw exceptions or freely access locations outside the scope of a surrounding dynamic binder as long as no code objects are involved.
- We formalize the semantics and prove the type safety of a core calculus illustrating the key features of the proposed type sys-

¹ Ain Shams University.

```

public static
Integer power(Integer x, Integer n){
    if (n == 1)
        return x;
    else
        return x * power(x, n-1);
}

```

(a) The Unstaged Power Function

```

public static
Code<Integer> spower(Code<Integer> x, int n){
    if (n == 1)
        return x;
    else
        return <| 'x * '(spower(x, n-1)) |>;
}

```

(b) The Staged Power Function

Figure 1. Staging the Power Function in Mint

tem. (Section 6). Proving type safety establishes that a well-typed program is guaranteed to be free of any runtime errors, including possible scope extrusion and generation (and execution) of ill-formed code. Lemmas are only stated here; full proofs are given in the companion technical report [27].

- We have implemented the type system in an MSP extension of OpenJDK [16] called Mint [17]. The implementation was used to type check all examples presented in the paper. We also use the implementation to confirm that MSP in Java — a language typically implemented using JIT compilation — can lead to performance speedups similar to those seen by MSP extensions of other languages (Section 7).

2. Multi-Stage Programming in Mint

Mint extends Java 6 with the three standard MSP constructs: brackets, escape, and run (see e.g. [24]). Brackets are written as `<| |>` and delay the enclosed computation by returning it as a code object. For example, `<| 2 + 3 |>` is a value. Brackets can contain a block of statements if it is surrounded by curly braces:

```

<| { C.foo();
    C.bar(); } |> // has type Code<Void>

```

Code objects have type `Code<T>`, where `T` is the type of the expression contained. For example, `<| 2 |>` has type `Code<Integer>`. A bracketed block of statements always has type `Code<Void>`.

Code objects can be escaped or run. Escapes are written as `'` and allow code objects to be spliced into other brackets to create bigger code objects. For example,

```

Code<Integer> x = <| 2 + 3 |>;
Code<Integer> y = <| 1 + 'x |>;

```

stores `<| 1 + (2 + 3) |>` into `y`. Run is provided as a method `run()` that code objects support. For example, executing

```

int z = y.run();

```

after the above example sets `z` to 6.

Mint also allows cross-stage persistence (CSP), wherein a variable bound outside brackets can be used inside the brackets, as in

```

int x = 1;
Code<Integer> c = <| x + 1 |>;

```

Weak separability, the requirement Mint uses to ensure safety, places certain restrictions on CSP; see Section 4.

Basic MSP in Mint can be illustrated using the classic power example. Figure 1(a) displays the unstaged power function in Java. Figure 1(b) displays a staged version. The staged method `spower` takes in an argument `x` that is a piece of code for an integer, along with an integer `n`, and returns code that multiplies `x` by itself `n` times.

3. The Scope Extrusion Problem

One of the most important properties of MSP languages is the guarantee that program generators will always produce well-formed

code. It is known how to achieve this in the purely functional setting [4, 25, 26]. In the presence of imperative features, however, guaranteeing this is more challenging because of the possibility of *scope extrusion*, where a code object containing a variable is used outside the scope of the binder for that variable. If such a code object were allowed to be compiled and run, a runtime error would be emitted, because the result of compiling and running code with free variables is undefined.

Scope extrusion can be caused by the following situations:

1. Assigning a code object to a variable or field that is reachable outside the escape, for example:

```

Code<Integer> x;
<| { Integer y = foo();
    '(x = <| y |>); } |>;

```

2. Throwing an exception that contains a code object, for example:

```

Code<Integer> meth(Code<Integer> c) {
    throw new CodeContainerException(c);
}
<| { Integer y; '(meth(<| y |>)); } |>

```

3. Cross-stage persistence (CSP) of a code object, an example of which is displayed in Figure 2.

The first two cases are straightforward; the first example extrudes `y` from its scope by assigning `<| y |>` to the variable `x` bound outside of the scope of `y`, while the second example throws an exception containing `<| y |>` outside the scope of `y`. The third example, however, is more subtle. This example creates an anonymous inner subclass of `Thunk`, whose `call` method returns a code object containing the variable `y`. This `Thunk` object is then passed to `doCSP` in the escape, yielding the code object

```

<| { Integer y = foo();
    Code<Integer> d = T.call(); } |>

```

where `T` is the anonymous inner subclass of `Thunk`. In a substitution-based semantics which substitutes values directly for variables, no scope extrusion would occur, because running this code object would substitute the return value of `foo()` for `y` in `T`, thus producing a new copy of `T` whose `call` method returns `<| r |>`, where `r` is the return value of `foo()`. Such a semantics, however, would be impractical for a language like Java, because it would involve traversing the (possibly compiled) method definitions of `T`, and it would also be confusing, because the `call` method of `T` itself would not be called when this example was run. In an environment-based semantics, which moves the value to some location in a separate environment and substitutes an index of that location (usually the variable itself) for the variable, running the above code would simply return the value of `T.call()`, which would produce the code object `<| y |>`, with `y` being out of its scope. The cross-stage persistence of the `Thunk` is essentially the same as storing the value of `T` in a global reference cell or hash table, and as soon as that

is done, the `Thunk` has been moved outside the scope of `y`. Our formalization of `Mint` properly models this indirection, although we model the environment as part of the heap rather than as a separate entity (see Section 6).

4. Weak Separability

The three situations mentioned in the previous section necessarily all involve code objects; effects that do not involve code cannot cause scope extrusion. We therefore first define the term code-free:

Definition 1. A *type* is **code-free** if it is not a subtype of `Code<T>`, the types of all of its fields are code-free, all of its methods’ return types are code-free, and its class is *final*. A *value* is **code-free** if its type is code-free.

If a value `x` is code-free, then no code object is reachable from `x`, and scope extrusion cannot occur due to effects involving `x`.

The requirement that a class is *final* means that the class is not allowed to be subclassed. This restriction ensures that a subclass with an additional field of type `Code<T>` cannot be substituted at runtime. Making Java classes *final* is unusual, but only because there are almost no benefits to declaring a class *final*. In our experience, most Java classes written by application developer can be *final*.

Commonly used code-free types include number types such as `Integer` and `Double`, the `String` class, arrays of code-free types, and all of Java’s reflection classes such as `Class` and `Field`. It does not include `Object`, for example, as this type is not *final*, and an `Object` could be a code object at runtime.

Scope extrusion can be prevented by requiring that escapes are *weakly separable*, which we define informally:

Informal Definition 1. A term is **weakly separable** if both side effects observable outside the term and cross-stage persistence involve only code-free values.

Intuitively, requiring escapes to be weakly separable will prevent scope extrusion, because no code object can be moved outside of an escape. A similar but stronger restriction was used in the systems of Kameyama et al. [10, 11] to ensure that no effects occurring inside escapes could be visible outside the escapes.¹ We call the condition introduced by Kameyama et al. **separability**.

We leave the notion of weak separability only informally defined here both because formalizing it would require complex semantic definitions and because it is undecidable in general. Instead, we provide a conservative approximation of weak separability that is used by `Mint`. This approximation is decidable, and we show below that it still leaves an expressive language. Unless otherwise specified, the phrase “weakly separable” in the remainder of this document refers to this approximation:

Definition 2. A `Mint` term `e` is **weakly separable** iff:

1. Assignment is made only to variables bound within `e`, or to fields or variables of code-free types;
2. Exceptions are only thrown by a `throw new C(e1, ..., en)`; construct where the `ei` are code-free, or the exception is caught by an enclosing `try-catch` construct before it can leave `e`;
3. Cross-stage persistence occurs only for `final` variables of code-free types;
4. Only weakly separable methods and constructors are called.

We are using the word *term* to describe an expression, statement, or method body. The notion of *pseudo-expression* in Section 6.2 is a direct formalization of this kind of term.

¹Technically, Kameyama et al. placed this restriction on future-stage binders.

```
interface Thunk { Code<Integer> call(); }

Code<Code<Integer>> doCSP(Thunk t) {
    return <| t.call() |>;
}

<| {
    Integer y = foo();
    Code<Integer> d = '(doCSP(new Thunk() {
        Code<Integer> call() {
            return <| y |>;
        }}});
} |>.run();
```

Figure 2. Cross-stage Persistence of Code Objects

The first three of the clauses in the definition above directly preclude the three cases of scope extrusion in the previous section. The restriction on throwing exceptions is syntactic and allows for a static check of code-freedom. Note that the *final* restriction on CSP variables exists so that the value of the variable does not change over the lifetime of the code object; Java has a similar restriction for variables referenced inside anonymous inner classes. The last clause ensures that all methods called from the body of a weakly separable term also satisfy weak separability. To check this condition, methods that are going to be called from the body of an escape are explicitly annotated in `Mint` with the keyword *separable*.

The following code shows a number of examples where weak separability prevents leakage of code values. Each of the statements in the method `foo` represents coding patterns that potentially cause scope extrusion, and the type checker rejects them. The assignment is rejected because the field being assigned to does not have code-free type. The exception is not allowed to be thrown because the exception object receives a non-code-free argument in the constructor.

```
class C {
    Code<C> c;
}
...
// binding for a future-stage variable y here
separable Code<C> foo(final C x) {
    // x.c is not code-free, Def 2(1)
    x.c = <|y|>;
    // CodeContainerException takes a
    // non-code-free argument, Def 2(2)
    throw new CodeContainerException(<|y|>);
    // CSP of x, but C is not code-free
    // not separable by Def 2(3)
    return <|x|>;
}
```

5. Expressivity

Weak separability is a highly expressive notion that excludes only a few coding patterns, such as generators that store open code in a global data structure. Such situations can be addressed by introducing a local data structure and ensuring that it is not used outside the dynamic binder. The separability restriction is only problematic when the coding pattern *requires* scope extrusion.

In an imperative language like Java, being able to write code generators with side effects is a desirable property: An imperative programming style fits better into the language’s culture, while functional programming in Java can often be inconvenient and verbose.

Weak separability does not severely restrict expressiveness because there is a tendency for the computational effects used in code generators to be separable: It is rare that generators are required to export code objects through side effects; all other side effects not involving code are weakly separable and therefore allowed. Furthermore, the `run()` method is only called outside of any brackets in almost all applications of MSP, and cross-stage persistence is mostly used for primitive types. Generated code is not restricted by weak separability at all.

To illustrate these points, the remainder of this section describes the implications of weak separability and examines a number of MSP examples in Mint, namely: staging an interpreter, a classic MSP example; staging array views [21] to remove abstraction overhead; and loop unrolling. The staged interpreter shows that throwing a code-free exception in a code generator is allowed. Both staging array views and loop unrolling demonstrate generators for imperative code. Section 5.3 gives another example, a staged serializer that uses Mint’s reflection capabilities. The performance of all these examples is evaluated in Section 7.

5.1 Staged Interpreter

Staged interpreters are a classic application of MSP. To demonstrate that staged interpreters can be written in Mint, we have implemented an interpreter for a small programming language called `lint` [24], which supports integer arithmetic, conditionals, and recursive function definitions of one variable.

The unstaged interpreter represents expressions with the `Exp` interface, and instantiates this interface with one class for each kind of AST node in the language. This interface specifies the single method `eval` for evaluating the given expression, which takes two environments, one for looking up variables and the other for looking up defined functions. The environments are modeled as functions, implemented using anonymous inner classes. The empty environments unconditionally throw an exception.

Applications of defined functions, for example, are implemented as follows:

```
class App implements Exp {
    private String _s;
    private Exp _a; // argument
    public App(String s, Exp a) {
        _s = s; _a = a;
    }
    public int eval(Env e, FEnv f) {
        return f.get(_s).apply(_a.eval(e, f));
    }
}
```

where `f.get(_s)` looks up the defined function named by the string `_s` as an object of the class `Fun`. The `Fun` class has an `apply` method for applying a function to an argument, and this is used here to apply the function to the result of evaluating the argument `_a`. If `_s` does not name a valid, defined function in `f`, then `get` throws an exception.

The staged interpreter redefines the `eval` method to return `Code<Integer>`, so that evaluating an expression yields code to compute its value. This method is marked as `separable` so that it can be called from inside an escape. For example, staging the `App` class above yields the following:

```
class App implements Exp { /* ... */
    separable Code<Integer> eval(Env e, FEnv f) {
        return <| (f.get(_s)).apply(
            (_a.eval(e,f)) |>;
    }
}
```

The `get` method is again used to look up the function named by `_s`. The return type of `get` is now `Code<Fun>`, meaning that `get` returns code for the defined function. This code is spliced

into the returned code, and its result is applied to the evaluation of the argument using the `apply` method. The argument for `apply` is obtained by splicing in the code object returned by `_a.eval`. If `_s` does not name a valid, defined function, then the `get` method throws an exception.

```
interface Env {
    public separable Code<Integer> get(String y);
}
static separable Env ext(final Env env,
    final String x, final Code<Integer> v) {
    return new Env() {
        public separable
        Code<Integer> get(String y) {
            if (x==y) return v; // lookup succeeded
            else return env.get(y); // recurse
        };
    };
}
Env env0 = new Env() { // empty environment
    public separable Code<Integer>
    get(String s) {
        throw new Yikes(); // unconditional throw
    }
};
```

Note that this is the only computational effect in the whole staged interpreter that happens inside a code generator. It is weakly separable, however, because the thrown exception need only contain the string argument `_s` that was not found in the environment; the exception therefore is code-free.

5.2 Array Views

As discussed, weak separability does not restrict the computational effects in generated code; it only restricts effects in the code generators themselves. As an example of this, we use staging to remove overhead in array views, which are useful for parallel programming.

It can be challenging for a compiler to parallelize Java code that uses multi-dimensional arrays, implemented in Java as nested one-dimensional arrays. This arrangement prevents the compiler from assuming that `A[i][j]` and `A[i+1][j]` refer to different locations. In one-dimensional arrays, each index refers to a different location, and this knowledge allows for simpler parallelization.

To address the problems with multi-dimensional arrays in Java, the `Habanero` project [7] provides *array views* that map subsections of multi-dimensional arrays to one-dimensional arrays without requiring the programmer to perform the arithmetic manually [21].

Much simplified, an array view stores a reference to a one-dimensional base array and provides `get` and `set` methods for several numbers of dimensions. The example below shows `get` and `set` for a two-dimensional double array:

```
class DoubleArrayView {
    double[] base;
    public double get(int i, int j) {
        return base[offset + (j-j0) + jSize*(i-i0)];
    }
    public void set(double v, int i, int j) {
        base[offset + (j-j0) + jSize*(i-i0)] = v;
    }
}
```

The methods calculate the the index in the base array for the coordinates (here: `i`, `j`) using the minimum values for those dimensions (here: `i0`, `j0`), the size of the the dimensions (here: `iSize`, `jSize`) and the start index of the view (`offset`). This calculation is performed for every array access.

The overhead of this calculation can be removed by staging the array view. The base array itself and the parameters `i0`, `j0`, `iSize`, `jSize`, and `offset` describing the array view are replaced by code values that can be spliced together as needed. The `get` method returns `Code<Double>`, code to retrieve the value in the

array. The `set` method returns `Code<Void>` and is a generator for code performing the array assignment.

```
class SDoubleArrayView {
    Code<double[]> base;
    public separable
    Code<Double> get(final int i, final int j) {
        return <| '(base)['(offset) + (j-'(j0)) +
                    '(jSize)*(i-'(i0))] |>; }

    public separable
    Code<Void> set(final Code<Double> v,
                  final int i, final int j) {
        return <| {
            '(base)['(offset) + (j-'(j0)) +
                '(jSize)*(i-'(i0))] = '(v); } |>; }
}
```

Staging allows us to work with the array view at a high level without paying for the overhead at runtime. For example, a matrix transpose can be written as

```
public Code<Void> stranspose(int m, int n,
    final SDoubleArrayView input,
    final SDoubleArrayView output) {
    Code<Void> stats = <| { } |>;
    for(int i = 0; i < m; i++)
    for(int j = 0; j < m; j++)
        stats = <| {
            'stats;
            '(output.set(input.get(i,j),j,i)); } |>;
    return stats;
}
Code<Void> c = stranspose(4, 4, a, b);
```

This method generates code consisting of direct array accesses:

```
b[0+(0-0)+4*(0-0)] = a[0+(0-0)+4*(0-0)];
b[0+(0-0)+4*(1-0)] = a[0+(1-0)+4*(0-0)]; // ...
```

An optimizing compiler will replace the computations with constants, completely removing the cost of the abstraction in the generated code.

Note that the size of the dimensions does not have to be known statically, but at generation time; only the number of dimensions of the array has to be known statically. Therefore, the size of the array can vary, and this optimization is applicable in many situations.

This example performs loop unrolling as well. The code objects that perform the array assignments, returned by the `set` method, are accumulated into another code object `stats`. At the end of the two nested loops, `stats` will contain the sequence of statements for the entire matrix transpose operation. The code generator is written in an imperative style consistent with the prevalent Java culture. The body of the method is weakly separable because `stats` is bound inside the method; the code inside `stats` returned by the `stranspose` method is not weakly separable, but that is not required of generated code.

Staged array views and loop unrolling serve as compelling examples how MSP can reduce abstraction overhead in an imperative setting.

5.3 Staged Reflection Primitives

Neverov observed that staging and reflection in languages like C# and Java can be highly synergistic [14]. He also noticed that fully exploiting this synergy requires providing a special library of staged reflection primitives. Mint provides such a library. The primitives are based on the standard reflection primitives in the Java library, including the `Class<A>` and `Field` classes.²

²The Mint reflection library does not support all reflection primitives. For example, `Method` and `Constructor` have multiple parameters. This would require adding indexed types to Java, and is therefore outside the scope of this work.

To represent these in Mint, the library adds two corresponding types, `ClassCode<A>` and `FieldCode<A,B>`. The `ClassCode<A>` type is indexed by the class itself, just like the type `Class<A>` it is modelled after. For example, the corresponding class for `Integer` objects has type `ClassCode<Integer>`. Any `ClassCode<A>` object provides methods for manipulating the class, corresponding to the methods of `Class<A>`. For example, the `cast` method of `ClassCode<A>` takes any code object of type `Code<Object>` and inserts a cast in the code object, yielding a code object of type `Code<A>`. Because the cast is inserted into the code, any exceptions raised by the cast will not happen until the code is run with the `run()` method. The class also provides methods for looking up a class by name and for retrieving the fields of a class.

The type `FieldCode<A,B>` represents a field in class `A` that has type `B`. It provides a `get` method which takes a `Code<A>` value and returns a value of type `Code`. This method constructs field selection (intuitively, a `<| ('a).f |>` code fragment) on that object. The type also provides a `getFieldClassCode` method to return a `ClassCode` object for the type `B`.

The fields of a class are returned using the `getFields()` method in `ClassCode<A>`. The return type of `getFields()` is `FieldCode<A,?>[]`, where the `?` represents an existential type in Java's generics. The method therefore returns an array of fields contained in class `A`, all of which have “some” type.

The following example illustrates the use of these classes. The code defines a serializer, which recursively converts an object and all of its fields to a string representation. Serializers are often slow, however, because they must use Java's reflection primitives to determine the fields of an object at runtime. Here we show how to write a staged serializer, which generates a serializer for a given static type. This approach performs the necessary reflection when the serializer is generated, and produces code to serialize all of a given object's fields without reflection:

```
public static separable
<A> Code<Void> sserialize(ClassCode<A> t,
    final Code<A> o) {
    if (t.getCodeClass()==Byte.class)
        return <| {
            writeByte('((Code<Byte>)o)); } |>;
    else if (t.getCodeClass()==Integer.class)
        return <| {
            writeInt('((Code<Integer>)o)); } |>;
    Code<Void> result = <| { } |>;
    for(FieldCode<A,?> fc: t.getFields()) {
        result = <| {
            'result;
            '(sserializeField(fc, o)); } |>;
    }
    return result;
}
```

The code to write primitive fields is generated directly. Non-primitive fields are visited recursively. The code is then spliced together and returned. This example was inspired by a similar example given by Neverov and Roe [14].

6. Type Safety

We now turn to formalizing a subset of Mint, called Lightweight Mint (LM), and to proving type safety. Type safety implies that scope extrusion is not possible in Mint.

LM is based on Lightweight Java [23] (LJ), a subset of Java that includes imperative features. LM includes staging constructs (brackets, escapes, and run), assignments, and anonymous inner classes (AICs). These features—especially the staging constructs and AICs—make the operational semantics and type system large; staging constructs alone double the number of rules in the operational semantics, while AICs increase the complexity of the type

system. All of these features, however, are necessary to capture the safety issues that arise in Mint. Specifically, assignments are required to cause many forms of scope extrusion, and AICs are required to create the scopes (i.e., the additional variable bindings) that can be extruded. AICs also lead to more complex possibilities for scope extrusion as shown in Section 3, which we wish to show are prevented by our system.

A noteworthy feature of the type system is the use of a *stack* of store typings instead of a single store typing. The stack of store typing is represented as a sequence. Each time we enter the scope of a variable binding, say of x , we push a new Σ to the right of the sequence that may type heap locations holding code values containing x free. During execution, type assignments to non-code-free heap locations l allocated within this scope are always added to this Σ . The Σ is popped from the right upon leaving the scope of x , so references to l are only typable in the scope of x .

When we pop a Σ from the stack of Σ 's, we always salvage code-free locations and tack them on to the Σ 's that remain, so that those locations can be used elsewhere. The Smashing Lemma guarantees that this contraction, or smashing, of the store typing stack is safe.

For this strategy to work, the Σ 's must obey a condition somewhat similar to the Barendregt variable convention, which governs $\bar{\Gamma}$. Whenever Σ is appended to the store typing stack, the locations in Σ must be fresh with respect to other Σ 's in the same typing derivation, i.e. the domains of all distinct Σ 's appearing in a single derivation tree must be disjoint. Otherwise, a location l that types $\langle x \rangle$ under one Σ can be used to type an incompatible object in a disjoint subtree of the derivation tree. A more rigorous definition of this constraint can be found in the companion technical report [27].

To simplify the formalism, LM disallows assignments to local variables; all assignments must be to object fields. This restriction by itself would completely rule out assignments in escapes, however. To rectify this problem, we add a restricted form of `let`, written as

```
let x <= new C (...) in ...
```

which always allocates a new instance of a class C that is not an AIC. We then relax the restrictions on escapes to allow field assignments if the object containing the field was allocated by a `let` inside the escape. Local variable assignment can then be modeled by replacing any local variable binding x of type C for which there is an assignment by a `let`-binding of a new variable x_cell of type $CCell$, defined as follows:

```
public class CCell { public C x; }
```

Uses of x , including assignments to x , can then be replaced by uses of $x_cell.x$. Thus, we model the environment as part of the heap, which simplifies the formalism without causing it to degenerate to a substitution-based semantics that we mentioned is unsatisfactory in Section 3.

6.1 Syntax

In this section, we formalize the syntax of LM. We use the following sequence notation:

Notation. We write $\langle a, b, c, \dots \rangle$ for sequences, with the shorthand $\langle A_i \rangle_{i=I}^J$ for $\langle A_I, A_{I+1}, \dots, A_J \rangle$. I may be omitted, and it defaults to 1. J may also be omitted when clear from context. The empty sequence is written $\langle \rangle$. Concatenation of sequences s_1 and s_2 is written $s_1 \circ s_2$, with the shorthand $\langle A_i \rangle_i, A$ for $\langle A_i \rangle_i \circ \langle A \rangle$. We also use $\langle e_i \rangle_{i=I}^J [i_0 \rightarrow x]$ to denote $\langle e_i \rangle_i$ with e_{i_0} replaced by x .

The syntax of LM is given in Figure 3. Expressions are stratified into levels. An expression is at level n if, for every point in the

extensible class names	D
final class names	F
variables	x
field names	f
method names	m
heap locations	l
classes	$C ::= D \mid F$
separability marker	$S ::= \text{sep} \mid \text{insep}$
types	$\tau ::= C \mid \text{Code}(S, \tau)$
class declarations	$CL ::= \text{class } C \text{ extends } D$ $\quad \{ \langle \tau_i f_i \rangle_i^I; \langle M_j^0 \rangle_j^J \}$
method declarations	$M^n ::= S \tau m (\langle \tau_i x_i \rangle_i) \{ e^n \}$
class hierarchy	$P ::= \langle CL_i \rangle_i$
programs	$p ::= P, e^0$
expressions	$e^n ::= x \mid l \mid e^n.f \mid (e^n.f := e^n)$ $\quad \mid e^n.m(\langle e_i^n \rangle_i)$ $\quad \mid \text{let } x \leftarrow \text{new } C(\langle e_i^n \rangle_i) \text{ in } e^n$ $\quad \mid \text{new } D(\langle e_i^n \rangle_i) \{ \langle M_j^n \rangle_j^J \}$ $\quad \mid \langle e^{n+1} \rangle \mid \langle e^{n-1} \rangle [n > 0]$ $\quad \mid e^n.\text{run}()$
values	$v^n ::= l \mid e^{n-1} [n > 0]$

NB: Production rules marked $[n > 0]$ can be used only if $n > 0$.

Figure 3. Lightweight Mint syntax.

expression, the nesting of escapes is at most n levels deeper than brackets. Clearly, a level- n expression is also a level- $(n + 1)$ expression. This stratification induces a similar structure on method declarations. A complete program must not have any unmatched escapes, so the bodies of methods declared in the class hierarchy are required to be at level 0. Likewise, the initial expression in a program is required to be at level 0. Values are also stratified: a value at level 0 is just a heap location, and a value at any level $n > 0$ is any expression at level $< n$.

Remark. The n in e^n is a constraint on the shapes of terms that this metavariable ranges over and is not something that forms a part of any concrete term. Similar precautions apply to other superscripted metavariables in this formalism, like v^n and M^n .

We categorize classes as final (F) or extensible (D) depending upon their names. In the implementation, they are rather categorized according to the manner in which they are declared, but using disjoint sets of names gives a simpler system. $\text{Code}(S, \tau)$ falls under neither classification. We do not allow an AIC to have fields or methods that its parent does not, although we allow method overrides. Additional fields or methods can be emulated by declaring (statically) a new subclass with those fields and creating anonymous subclasses of those.

We do not include the syntax (`new C(...)`) for instantiating ordinary (i.e., non-AIC) classes because one can write (`let x ← new C(...)` in x) instead. Sequencing ($e_1; e_2$) is also omitted because this code can be written $\text{seq.call}(e_1, e_2)$, where seq.call is a method that ignores its first argument and returns its second.

The code type is indexed by a separability marker which indicates whether a code object is itself separable. Specifically, $\text{Code}(\text{sep}, \tau)$ is the type of code objects containing separable code, which is a subtype of the standard code type, written $\text{Code}(\text{insep}, \tau)$. This distinction is necessary in the case of a separable expression which itself contains a nested escape ' e ', since we must know for type preservation that ' e is guaranteed to reduce only to separable code. In this case, e must have type $\text{Code}(\text{sep}, \tau)$.

All judgments and functions in the following discussions implicitly take a class hierarchy P as a parameter. We avoid writing

it out explicitly because it is fixed for each program and there is no fear of confusion.

6.2 Operational Semantics

Figure 4 shows the small-step semantics for Lightweight Mint. This figure defines the judgment $H_1, \hat{e}_1 \xrightarrow{n} H_2, \hat{e}_2$ which states that heap H_1 and pseudo-expression \hat{e}_1 evaluate in a single step at level n to H_2 and \hat{e}_2 , respectively. The notion of heap is standard; a heap is a finite mapping from locations to heap elements, where a heap element contains a runtime type tag with either the contents of the object or a code value if the tag is `Code`. The notion of a pseudo-expression, defined as either an expression or a method body, is used because execution can occur within bodies of methods defined at level > 0 . Similarly, we define the notion of pseudo-value, which is either a value or a method body with no unresolved escapes at the current level.

Figure 4 contains a number of helper functions. `fields()` extracts the fields of a type. `method()` looks up a method, respecting overriding rules. `mbody()` extracts the specified method’s formal arguments and body. Code types do not have methods (`run()` is formally not a method). `mname` extracts the method name from a method declaration.

The single-step evaluation judgment \xrightarrow{n} is defined as the closure of the primitive one-step relation $\xrightarrow{k}_{\text{prim}}$ under n, k -evaluation contexts $\mathcal{E}^{n,k}$. An n, k -evaluation context takes a level- k expression as argument and yields a level- n pseudo-expression. See the technical report for a BNF-like definition of $\mathcal{E}^{n,k}$.

Most of the primitive reduction steps are straightforward, including rules for class instantiation, method invocation, and assignment. These reductions only occur at level 0, to prevent reductions from occurring inside code objects. As local variables are immutable, we can model method invocation and `let`-execution by substitution. The local binding L found in LJ [23] is therefore unnecessary, and the small-step judgment is made between heap-term pairs rather than bindings-heap-term triples. Recall that, as discussed at the beginning of this section, this is not the same as using a substitution-based semantics.

Each staging construct induces a primitive reduction rule. Escape and run extract expressions from brackets. Escape reduces only at level 1 and run only reduces at level 0. These are standard in multi-stage languages [25], except that the code values are on the heap. Brackets allocate a code object on the heap. CSP, which can be regarded as execution at arbitrarily high levels, is automatically taken care of by substitution and does not give rise to a redex.

6.3 Type System

Figure 5 shows the type system. A variable typing (or type environment) comes in pairs, separated by a $|$. The predicate $\text{iscf}_{\langle F_i \rangle}(\tau)$ asserts that τ is code-free assuming that the final classes $\langle F_i \rangle$ are. Thus $\text{iscf}_{\emptyset}(\tau)$, also written $\text{iscf}(\tau)$, means that τ is code-free. Note, however, that $\text{iscf}()$ does not require method types to be code-free because the formal language does not model class declarations within anonymous inner classes. The auxiliary functions `fypes()`, `ftypei()`, and `mtype()` are similar to `fields()` and `method()`, but they extract type information. The top-level judgment $\vdash p$ asserts that program p is a valid “initial state” of execution: the class hierarchy P contained in p must be well-formed and the expression e contained in p must be well-typed and free of heap locations. A class hierarchy P is well-formed, written $\vdash P$, if P is acyclic, field names and types (including inherited ones) do not clash within each class, and each class is well-formed. We omit a formalization of the first two checks but will use them implicitly by assuming that auxiliary functions like `fields()` and `mtype()` are always unambiguous and that the sequence returned by `fields()` is

operational terms	
heaps	$H : l \xrightarrow{\text{fin}} h$
runtime type tags	$T ::= C \mid \text{sub } D \{ \langle M_i^0 \rangle_i \} \mid \text{Code}$
heap elements	$h ::= (C, \langle l_i \rangle_i) \mid (\text{Code}, \langle e^0 \rangle) \mid (\text{sub } D \{ \langle M_i^0 \rangle_i \}, \langle l_j \rangle_j)$
pseudo-expressions	$\hat{e}^n ::= e^n \mid M^n$
pseudo-values	$\hat{v}^n ::= v^n \mid M^{n-1} [n > 0]$
$H, e^n \xrightarrow{n}_{\text{prim}} H, e^n$ and $H, \hat{e}^n \xrightarrow{n} H, \hat{e}^n$	
$l \notin \text{dom } H$	
$H, \text{new } D(\langle l_i \rangle_i) \{ \langle M_j^0 \rangle_j \} \xrightarrow{0}_{\text{prim}} H[l \mapsto (\text{sub } D \{ \langle M_j^0 \rangle_j \}, \langle l_i \rangle_i), l$	
$\frac{H(l) = (T, \langle l_i \rangle_i) \quad \text{fields}(T) = \langle f_i \rangle_i}{H, l.f_{i_0} \xrightarrow{0}_{\text{prim}} H, l_{i_0}}$	
$\frac{H(l) = (T, \langle l_i \rangle_i)}{H, (l.f_{i_0} := l') \xrightarrow{0}_{\text{prim}} H[l \mapsto (T, \langle l_i \rangle_i [i_0 \rightarrow l']), l'}$	
$\frac{H(l) = (T, \dots) \quad \text{mbody}(m, T) = (\langle x_i \rangle_i, e^0)}{H, l.m(\langle l_i \rangle_i) \xrightarrow{0}_{\text{prim}} H, [\langle l_i \rangle_i / \langle x_i \rangle_i] [l/\text{this}] e^0}$	
$l \notin \text{dom } H$	
$H, \text{let } x \leftarrow \text{new } C(\langle l_i \rangle_i) \text{ in } e^0 \xrightarrow{0}_{\text{prim}} H[l \mapsto (C, \langle l_i \rangle_i), [l/x] e^0$	
$\frac{H(l) = (\text{Code}, \langle e^0 \rangle)}{H, \text{' } l \xrightarrow{1}_{\text{prim}} H, e^0} \quad \frac{H(l) = (\text{Code}, \langle e^0 \rangle)}{H, l.\text{run}() \xrightarrow{0}_{\text{prim}} H, e^0}$	
$l \notin \text{dom } H$	
$\frac{H, \langle e^0 \rangle \xrightarrow{0}_{\text{prim}} H[l \mapsto (\text{Code}, \langle e^0 \rangle)], l}{H_1, e_1^k \xrightarrow{k}_{\text{prim}} H_2, e_2^k}$	
$\frac{H_1, \mathcal{E}^{n,k}[e_1^k] \xrightarrow{n} H_2, \mathcal{E}^{n,k}[e_2^k]}{H_1, \mathcal{E}^{n,k}[e_1^k] \xrightarrow{n} H_2, \mathcal{E}^{n,k}[e_2^k]}$	

Figure 4. Small-step semantics for Lightweight Mint.

finite and has no duplicates. Classes are well-formed if they contain no locations, their methods are well-typed, and any overridden methods have the same types as in the superclass.

The bottom half of Figure 5 defines typing for pseudo-expressions with the judgment $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n \hat{e}^n : \hat{\tau} | S$, which states that the pseudo-expression \hat{e}^n has type $\hat{\tau}$ at level n under the stack $\langle \Sigma_i \rangle_i$ of store typings and the pair $\bar{\Gamma}$ of contexts. If $S = \text{sep}$, this judgment further states that the pseudo-expression \hat{e}^n is weakly separable. Variable typing $\bar{\Gamma}$ is partitioned into two parts in order to check weak separability of field assignments. The right part contains the variables that were bound within the current method or enclosing escape, which are precisely the variables whose fields can be assigned to without violating weak separability. We always assume that no variables are repeated in $\bar{\Gamma}$.

Most of the rules for typing pseudo-expressions are straightforward. The first rule generalizes subtypes to supertypes. (Subtyping rules are omitted due to space limitations; they are the same as in

typing terms

variable typing $\Gamma : x \xrightarrow{\text{fin}} \tau^n$
store typing $\Sigma : l \xrightarrow{\text{fin}} \tau$
variable typing pair $\bar{\Gamma} ::= (\Gamma | \Gamma)$
pseudo-types $\hat{\tau} ::= \tau | \langle \tau_i \rangle_i \xrightarrow{S} \tau$

$\text{isfc}_{\langle F_i \rangle}(\tau)$, $\text{isfc}(\tau)$

$$\frac{F \in \langle F_i \rangle}{\text{isfc}_{\langle F_i \rangle}(F)} \quad \frac{}{\text{isfc}(\langle \rangle(\tau))} \quad \frac{\forall i. \text{isfc}_{\langle F_j \rangle, F}(\text{ftype}_i(F))}{\text{isfc}_{\langle F_j \rangle}(F)}$$

$\text{cf}(\Sigma)$, $\text{locs}(\hat{e})$, $\text{ftypes}(\tau)$, $\text{ftype}_i(\tau)$, $\text{ftype}(f, \tau)$

$\text{cf}(\Sigma) = \Sigma|_L$ where $L = \{l \in \text{dom}(\Sigma) : \text{isfc}_{\langle \rangle}(\Sigma(l))\}$.
 $\text{locs}(\hat{e}) = \{l : l \text{ is a subterm of } \hat{e}\}$
 $\text{ftypes}(\tau) = \langle \tau_i \rangle_i$ assuming $\text{fields}(\tau) = \langle \tau_i f_i \rangle_i$
 $\text{ftype}_i(\tau) = \text{ftype}(f_i, \tau) = \tau_i$ assuming $\tau_i f_i \in \text{fields}(\tau)$

$\vdash p, \vdash P, \vdash CL$

$$\frac{\vdash P \quad \langle \rangle; \emptyset | \emptyset \vdash^0 e^0 : \tau | S \quad \text{locs}(e^0) = \emptyset}{\vdash P, e^0}$$

inheritance is acyclic no field names clash $\langle \vdash CL_i \rangle_i$
 $\vdash \langle CL_i \rangle_i$

$$\frac{\langle \text{locs}(M_i^0) = \emptyset \rangle_i \quad \langle \langle \rangle; \emptyset | \text{this} : C^0 \vdash^0 M_i^0 \rangle_i}{\langle \text{mtype}(\text{mname}(M_i), D) = \text{undef or mtype}(M_i^0) \rangle_i} \quad \vdash \text{class } C \text{ extends } D \{ \langle \tau_j f_j \rangle_j; \langle M_i^0 \rangle_i \}$$

$\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n : \tau | S$ Additional constraints on $\langle \Sigma_i \rangle_i$ are discussed in the text.

$$\frac{\tau' \prec \tau \quad \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n : \tau' | S}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n : \tau | S} \quad \frac{\bar{\Gamma}(x) = \tau^n \quad \text{isfc}(\tau) \vee k = 0}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^{n+k} x : \tau | S} \quad \frac{(\cup_i \Sigma_i)(l) = \tau \quad \text{isfc}(\tau) \vee n = 0}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n l : \tau | S}$$

$$\frac{\frac{\langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n e_j^n : \text{ftype}_j(C) | S \rangle_j}{\langle \Sigma_i \rangle_i; \Sigma; \Gamma_1, \Gamma_2 | x : C^n \vdash^n e^n : \tau | S}}{\langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n (\text{let } x \leftarrow \text{new } C \langle (e_j^n)_j \rangle \text{ in } e^n) : \tau | S}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n : \tau | S}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n . f : \text{ftype}(f, \tau) | S}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e_j^n : \tau_j | S \rangle_{j=1}^2 \quad \text{ftype}(f, \tau_1) = \tau_2 \quad S = \text{insep} \vee \text{isfc}(\tau_2)}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n (e_1 . f := e_2) : \tau_2 | S}$$

$$\frac{\bar{\Gamma}(x) = \tau_1^n \quad \langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n e^n : \tau_2 | \text{sep} \quad x \in \text{dom } \Gamma_2 \quad \text{ftype}(f, \tau_1) = \tau_2}{\langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n (x . f := e^n) : \tau_2 | \text{sep}}$$

$$\frac{(\cup_i \Sigma_i)(l) = \tau_1 \quad \text{isfc}(\tau_2) \vee (n = 0 \wedge l \in \text{dom } \Sigma_I) \quad \text{ftype}(f, \tau_1) = \tau_2 \quad \langle \Sigma_i \rangle_i^I; \bar{\Gamma} \vdash^n e^n : \tau_2 | \text{sep}}{\langle \Sigma_i \rangle_i^I; \bar{\Gamma} \vdash^n (l . f := e^n) : \tau_2 | \text{sep}}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n : \tau | S \quad \langle \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e_j^n : \tau_j | S \rangle_j}{\text{mtype}(m, \tau) = \langle \tau_j \rangle_j \xrightarrow{S} \tau} \quad \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n . m \langle (e_j^n)_j \rangle : \tau | S$$

$$\frac{\langle \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e_j^n : \text{ftype}_j(D) | S \rangle_j}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n \text{new } D \langle (e_j^n)_j \rangle \{ \langle M_k^n \rangle_k \} : D | S}$$

$$\frac{\langle \Sigma_i \rangle_i; \Gamma_1, \Gamma_2 | \emptyset \vdash^{n+1} e : \tau | S}{\langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n \langle e \rangle : \text{Code} \langle S, \tau \rangle | S'}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e^n : \text{Code} \langle S, \tau \rangle | \text{sep}}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^{n+1} \langle e \rangle : \tau | S}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e : \text{Code} \langle S, \tau \rangle | S'}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n e . \text{run}() | \text{insep}}$$

$\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n \text{sub } D \{ \langle M_i^n \rangle \}$

$$\frac{\langle \langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2, \text{this} : D^n \vdash^n M_j^n : \tau_j | S_j \rangle_j \quad n > 0 \vee \text{dom}(\cup_i \Sigma_i) \supseteq \text{locs}(\text{sub } D \{ \langle M_j^n \rangle_j \})}{\langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n \text{sub } D \{ \langle M_j^n \rangle_j \}}$$

where $\tau_j = \text{mtype}(\text{mname}(M_j^n), D)$.

$\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n M^n : \langle \tau_i \rangle_i \xrightarrow{S} \tau | S$

$$\frac{\langle \Sigma_i \rangle_i, \Sigma; \Gamma_1, \Gamma_2, \langle x_i : \tau_i^n \rangle_i | \emptyset \vdash^n e^n : \tau | S}{\langle \Sigma_i \rangle_i; \Gamma_1 | \Gamma_2 \vdash^n S \tau m \langle \langle \tau_i x_i \rangle_i \rangle \{ e^n \} : \langle \tau_i \rangle_i \xrightarrow{S} \tau | S'}$$

$\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash H$

$$\frac{\forall l \in \text{dom}(\cup_i \Sigma_i). \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash H(l) : (\cup_i \Sigma_i)(l)}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash H}$$

$\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash h : \tau$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash h : \tau \quad \tau \prec \tau' \quad \langle (\cup_i \Sigma_i)(l_j) \prec \text{ftype}_j(C) \rangle_j}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash h : \tau' \quad \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash (C, \langle l_j \rangle) : C}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma}^{\geq 1} \vdash^0 \langle e^0 \rangle : \text{Code} \langle S, \tau \rangle | S'}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash (\text{Code}, \langle e^0 \rangle) : \text{Code} \langle S, \tau \rangle}$$

$$\frac{\langle \Sigma_i \rangle_i; \bar{\Gamma}^{\geq 1} \vdash^0 \text{sub } D \{ \langle M_j^0 \rangle_j \} \quad \langle (\cup_i \Sigma_i)(l_k) \prec \text{ftype}_k(D) \rangle_k}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash (\text{sub } D \{ \langle M_j^0 \rangle_j \}, \langle l_k \rangle_k) : D}$$

where $\bar{\Gamma}^{\geq 1}(x) = \tau^n \iff \bar{\Gamma}(x) = \tau^n \wedge n \geq 1$.

Figure 5. Type system for Lightweight Mint.

plain Java.) The next two rules look up the types for variables and locations in the variable and store typings, respectively. CSP is only allowed (by $k > 0$ or $n > 0$, respectively) if the associated type is code-free. The next rule is typing `let`-expressions by extending the current context with the `let`-bound variable, while the rule following is typing field lookups by typing the object and then looking up the relevant field type. In typing the body of a future-stage `let`, a new frame Σ is added to the current stack $\langle \Sigma_i \rangle_i$ to allow for the possibility of heap locations containing code objects with the variable x free.

The next three rules type field assignments ($e_1.f := e_2$) by checking that e_1 has some τ_1 and that the type τ_2 of e_2 matches the appropriate field type of τ_1 . The first of these rules applies to arbitrary e_1 and types the assignment as weakly separable only if the field has code-free type. The second and third rule handle assignments of non-code-free type in a weakly separable context. The second rule requires e_1 to be a variable x and allows the assignment to be typed as separable only if the variable is in Γ_2 (recalling that Γ_2 gives precisely the variables whose fields can be modified without violating weak separability). The third rule essentially captures how a judgment using the first or the second rule transforms under substitution of a location l for the x in the second rule: it requires that either the location l is in the topmost store typing or the type of the field is code-free.

The rule following is typing method calls by looking up the type of the given method. The rule on the immediate right checks well-formedness of AIC definitions. Finally, the last three rules type brackets, escape, and run, where typing $\langle e \rangle$ requires typing e at the next level and adds the code type; typing $'e$ requires typing e at a code type on the previous level and removes the code type; and typing $e.run()$ types e at a code type on the same level and removes the code type. Brackets can always be weakly separable, run is never weakly separable, and escapes $'e$ are only weakly separable if e has type $\text{Code}(\text{sep}, \tau)$.

The remainder of Figure 5 defines the following judgments. $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n \text{sub } D \{ \langle M_i^n \rangle \}$ states that an anonymous subclass of D with method definitions $\langle M_i^n \rangle$ is well-formed. This requires the methods $\langle M_i^n \rangle$ to have the appropriate types. It also requires, if $n = 0$, that all the locations in the AIC are contained in $\text{dom}(\cup_i \Sigma_i)$, ensuring that no new frames can be added to the stack of store typings. The judgment $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n M^n : \langle \tau_i \rangle_i \xrightarrow{S} \tau | S$ states that method M has input types $\langle \tau_i \rangle_i$, output type τ , and further is weakly separable if $S = \text{sep}$. Note that this rule is allowed to push a new frame onto the stack of store typings when the level $n > 0$. This is because there may be some locations in the store that contain code that include the free variables bound inside M . Note also that passing inside a method resets the vertical bar $|$ in $\bar{\Gamma}$ to the end, indicating that weakly separable expressions in the method cannot freely access variables bound at or before the method M .

The judgment $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash H$ states that the store H is well-formed under the given stack of store typings. This judgment includes the typing context $\bar{\Gamma}$ because the store may contain code with free variables. This judgment requires that, for all locations l in the stack of store typings, the heap for $H(l)$ is well-typed. Note that there may be more locations in H than in the domain of $\langle \Sigma_i \rangle_i$, allowing the possibility that other frames could be pushed onto this stack. The judgment $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash h : \tau$ is then used to state that heap form h has type τ . The rules for this judgment require that the expressions contained in the heap form h are well-typed. The typing context used to type these expressions is the restriction of $\bar{\Gamma}$ to the variables of level greater than 0. This is because heap forms are allowed to have code objects with free variables in them, but these free variables must be bound in other code objects, meaning they must have been bound at level greater than 0. Note that, as a side effect of these definitions, if $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash H$ holds then H re-

stricted to $\text{dom}(\cup_i \Sigma_i)$ is closed under reachability, meaning that no location in this domain can reference a location outside of it.

6.4 Soundness

We now outline the key parts of our type soundness proof. Complete proofs can be found in the companion technical report [27]. Type soundness is proved by the usual Preservation and Progress lemmas. Progress follows directly from Unique Decomposition, which states that any well-typed expression is either a value or contains a unique redex, which can be contracted by the operational rules. Uniqueness also ensures that our semantics is deterministic.

Lemma 1 (Unique Decomposition). *If $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n \hat{e}^n : \tau | S$ and \hat{e}^n is not a pseudo-value then \hat{e}^n is uniquely decomposed as $\hat{e}^n = \mathcal{E}^{n,m}[r^m]$, where $=$ denotes syntactic equality modulo α conversion.*

The proof of Preservation is more complicated. One technical difficulty is that the extra Σ added in the rules for binding constructs are unrestricted and therefore may include locations that are not in the current heap. To avoid this problem, we introduce typing for **configurations**—pairs of heaps and pseudo-expressions. The judgment $\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n (H, \hat{e}^n) : \tau | S$ then specifies that the configuration (H, \hat{e}^n) is well-typed. Configuration typing rules are identical to pseudo-expression typing rules except that each rule also requires the heap H to be well-formed under the current context. For example, the rule for `let` becomes:

$$\frac{\langle \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n (H, e_j^n) : \text{ftype}_j(C) | S \rangle_j \quad \langle \Sigma_i \rangle_i; \Sigma; \bar{\Gamma}, x : C^n \vdash^n (H, e^n) : \tau | S \quad \langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash H}{\langle \Sigma_i \rangle_i; \bar{\Gamma} \vdash^n (H, \text{let } x \leftarrow \text{new } C(\langle e_j^n \rangle_j) \text{ in } e^n) : \tau | S}$$

A second technical difficulty is that a reduction step inside a `let` form or AIC that pushes a new frame Σ onto $\langle \Sigma_i \rangle_i$ might modify a code-free location in $\text{dom}(\cup_i \Sigma_i)$ to reference a location in the new frame Σ . Then the resulting heap is ill-formed under $\langle \Sigma_i \rangle_i$, because this portion of the heap is not closed under reachability. The Smashing Lemma solves this problem by smashing the top two Σ 's of $\langle \Sigma_i \rangle_i$ into one, giving a shorter store typing stack. Informally, it states that the side-effects of a small-step happening inside a dynamic binder are invisible from outside except for allocations of code-free locations.

Lemma 2 (Smashing). *If*

1. $\Gamma'_1 \cup \Gamma'_2 \supseteq \Gamma_1 \cup \Gamma_2$
2. $\langle \Sigma_i \rangle_i^I; \Gamma_1 | \Gamma_2 \vdash H_1$
3. $\langle \Sigma_i \rangle_i^I; \Sigma; \Gamma'_1 | \Gamma'_2 \vdash H_2$
4. $H_1|_L = H_2|_L$ where $L = \text{dom}(\cup_{i=1}^I \Sigma_i) - \text{dom}(\text{cf}(\cup_{i=1}^I \Sigma_i))$

then $\langle \Sigma_i \rangle_i^{I-1}, (\Sigma_I \cup \text{cf}(\Sigma)); \Gamma_1 | \Gamma_2 \vdash H_2$.

The different store typing stacks capture different views of the heap. The stack $\langle \Sigma_i \rangle_i^I$ captures the locations that can legally be referenced outside the binder (the outer view). The stack $\langle \Sigma_i \rangle_i^I, \Sigma$ captures what can be referenced from inside the binder (the inner view). Similarly, the variable typing pair $\Gamma_1 | \Gamma_2$ is used outside the binder while $\Gamma'_1 | \Gamma'_2$ is used inside the binder. The inner variable typing pair contains typings for the new variables in addition to those in the outer variable typing pair (condition 1). Code inside the binder may refer to locations that were illegal outside, namely those with code values that (may) contain free variables in $(\Gamma'_1 \cup \Gamma'_2) \setminus (\Gamma_1 \cup \Gamma_2)$. The extra Σ available in the inner view provides this extension to set of visible locations in the outer view.

Suppose a small-step in the inner context takes the heap from H_1 to H_2 . Our concern is whether H_2 is still well-formed, especially in the outer view. The heap had better be well-formed at the beginning in both the inner and outer views; the outer one matters

here (condition 2). After the small-step, we can easily show in the proof of Preservation that the new heap H_2 is well-formed in the inner view, by a suitable inductive hypothesis (condition 3). Recall that any new allocations during the small-step are attributed to the topmost store typing Σ , so the prefix $\langle \Sigma_i \rangle_i^f$ remains the same as in the check on H_1 . All the code-free restrictions on side effects exist to ensure that the small-step could not have touched any non-code-free locations that are visible in the outer view (condition 4). The lemma then asserts under these conditions that H_2 is indeed well-formed in the outer view, provided that we salvage any new code-free allocations from Σ into $\langle \Sigma_i \rangle_i^f$. We add these new allocations to the topmost store typing, Σ_T , in compliance with the convention of only modifying the topmost store typing.

The Smashing Lemma implies the absence of scope extrusion, as it states that any code locations that could potentially cause scope extrusion are not reachable outside their respective scopes.

Lemma 3 (Preservation). *If $\langle \Sigma_i \rangle_i, \Sigma_R; \Gamma_1 | \Gamma_2 \vdash^n (H_1, \widehat{c}_1^n) : \tau | S$ and $(H_1, \widehat{c}_1^n) \rightsquigarrow (H_2, \widehat{c}_2^n)$, then $\exists \Sigma'_R$ such that*

1. $\Sigma'_R \supseteq \Sigma_R$
2. $\langle \Sigma_i \rangle_i, \Sigma_R; \Gamma_1 | \Gamma_2 \vdash^n (H_2, \widehat{c}_2^n) : \tau | S$
3. $H_1|_L = H_2|_L$ where $L = \text{dom}(\cup_i \Sigma_i) - \text{dom}(\text{cf}(\cup_i \Sigma_i))$

This statement is an abridged version. As mentioned at the beginning of this section, we need freshness assumptions about heap locations in Σ , and this lemma must be expanded to incorporate them. The companion technical report contains a complete statement and a proof of this lemma.

7. Performance

To validate both the expressivity of our type system as well as the potential usefulness of MSP for Java, we extended the Java OpenJDK compiler [16] with our proposed type system and with runtime support for staging. All examples presented earlier in the paper were type-checked using this implementation [17]. In this section, we present timing results that confirm that MSP can affect the performance of Java programs in a way similar that observed in other languages.

7.1 Benchmarks

In order to measure the performance impact of MSP in Mint, we have benchmarked the following Mint examples:

- `power` is the power example from Section 2, called with base 2 and exponent 17.
- `fib` recursively computes the 17th element of the generalized Fibonacci function starting from 2 and 3.
- `mmult` performs sparse matrix multiplication, in which every 1 in the left matrix omits the floating-point multiplication at runtime and every 0 omits the multiplication and the addition. The benchmark is multiplies an 11-by-11 unsymmetric sparse matrix [5] with itself.
- `eval-fact` calculates the factorial of 10 using the `lint` interpreter discussed in Section 5.1.
- `eval-fib` calculates the 10th number in the standard Fibonacci sequence using the `lint` interpreter.
- `av-mmult` performs the same sparse matrix multiplication as `mmult`, but accesses the matrix using the array views described in Section 5.2.
- `av-mtrans` performs a matrix transpose using array views.

Benchmark	speedup	unstaged μs	staged μs
<code>power</code>	9.2	0.060	0.0065
<code>fib</code>	8.8	0.058	0.0065
<code>mmult</code>	4.7	13	2.7
<code>eval-fact</code>	20	0.83	0.042
<code>eval-fib</code>	24	18	0.73
<code>av-mmult</code>	65	20	0.30
<code>av-mtrans</code>	14	1.0	0.071
<code>serialize</code>	26	1.5	0.057

Figure 6. Benchmark results.

- `serialize` uses the serializer generator discussed in Section 5.3 to write the primitive fields contained in an object hierarchy two levels deep to an output stream.

For each operation in the benchmarking process (unstaged, staged), we first determine the number of repetitions that are required for the operation to run for 1-2 s. This calibration phase also allows the JIT compiler to finish optimizing the program both for the unstaged and the staged code. We then run as many repetitions of the operation as determined in the previous step and record the total time. The average runtime of a single repetition is calculated for each operation and used for the benchmark.

Timings were recorded on an Apple MacBook with a 2.0 GHz Intel Core Duo processor, 2 MB of L2 cache, and 2 GB main memory, running Mac OS 10.4.11 Tiger and the SoyLatte 1.0.3 JVM [?].

7.2 Results

The results are given in Figure 6. Performance improved in all cases. The speedups achieved range from 4.7 to 65, with speedup defined as unstaged time divided by staged time. The staged versions of `power` and `fib` executed approximately nine times faster than the unstaged code due to the removal of recursion. The `mmult` benchmark involved mostly tight `for` loops and could only be sped up by a factor of 4.7. Staging the `lint` interpreter removed call overhead and improved the performance of the `eval-fact` and `eval-fib` benchmarks by factors of 20 and 24, respectively. In the `av-mmult` and `av-mtrans` benchmarks, loops were unrolled and the layer of indirection in the form of array views was replaced by direct array accesses, resulting in speedups of 65 and 14, respectively. Finally, the `serializer` benchmark also benefited from staging through the removal of reflection, and execution time was reduced by a factor of 26.

The improvements achieved make it clear that the presence of JIT technology in Java does not subsume the need for staging techniques, and that the performance benefits reported in previous work [24?] on languages without JIT technology apply to Java as well.

8. Related Work

A distinguishing feature of Mint is a strong, expressive, and safe type system that permits both manipulation of open terms and imperative programming. Few multi-level imperative languages have such a type system, and fewer yet come with rigorous type safety proofs. ‘C [18] and Jumbo [12] do not guarantee well-formedness of generated code. Cyclone [22] statically guarantees type safety (including well-formedness) of generated code, but does not treat code as first-class values. This design helps Cyclone’s runtime code generator to be very fast and still produce high-quality code, but limits programmers in the way they can write generators. Other works on staging extensions to Java by Sestoft [20], Schultz et al. [19], and Zook et al. [28] focus on exploring novel uses of staging

and/or quantifying performance benefits. As such, these authors do not attempt to establish safety properties of their extensions.

Some multi-stage systems based on Java offer safety properties, but formalizations and proofs are often absent or incomplete. Metaphore [14] comes with a core typed, Java-like calculus but its type soundness is left unproved. The calculus also leaves out side effects. SafeGen [9] is claimed to guarantee well-typedness of generated code, but the authors do not prove such a result or formalize their system. MorphJ [8] focuses on reflection and does not allow manipulation of arbitrary code values (in particular open terms). The paper proves soundness, but the system does not model side effects. Fährndrich et al. [6] propose a system similar to MorphJ that allows the user to perform limited manipulations of code values, using reflection, in a type-safe manner. DynJava [15] has static type checking for dynamically generated code based on annotations about the types of free variables in code fragments. The authors claim type safety, but do not appear to offer a rigorous proof. The type annotations also make their code generation system unhygienic (i.e. α equivalence fails for dynamic code).

Much of the work on type safety proofs in the literature are for functional languages, where imperative extensions similarly cause scope extrusion [2, 3, 10, 11, 13]. Mint either compares favorably or is competitive with all of these systems. Calcagno et al. [3] allow imperative operations on code but do not support imperative operations on open terms. Aktemur [1] and Kim et al. [13] support unrestricted imperative operations on open terms but give up α -equivalence for future-stage code. Kim et al. delegate hygiene to a specialized binder λ^* , whose semantics can be explained only in terms of a “gensym.” Type terms also tend to be very large in their systems which may limit the programmer’s ability to write down or interpret types and correct type errors; this problem is offset to some extent, however, by type inference. Aktemur’s and Kim et al.’s systems allow storing open terms in global variables that can (potentially) outlive the scope of the variables they contain, which cannot be done in Mint. However, we are not aware of any use for this technique, apart from examples by Kim et al. that deliberately violate hygiene [13]. Unhygienic generation brings up the problem of inadvertent variable capture, and preventing it is a desirable feature of MSP.

Ancona and Moggi [2] and Kameyama et al. [10, 11] are the only works we know of that combine imperative operations on open terms and α -equivalence. Kameyama et al.’s approach [11] is closest to the weak separability approach presented in this paper, in that they limit the effects allowed within escapes. What distinguishes Mint from Ancona and Moggi’s and Kameyama et al.’s systems is that Mint allows effects occurring in escapes to be visible outside the escapes as long as they do not involve code objects, whereas Ancona and Moggi and Kameyama et al. unconditionally prohibit such effects. This difference makes our system more expressive, which is demonstrated in Section 5.1, where we throw an exception in a code generator, and in Section 5.2, where we accumulate code in a for-loop. Neither Ancona and Moggi’s nor Kameyama et al.’s system can directly express either of these examples. Even if we consider extending these systems, it is unclear how to extend them to handle Section 5.1’s example. To allow the example in Section 5.2, their calculi must incorporate arrays, which is a nontrivial theoretical exercise. Our type system is thus better suited to Java programming, which in general makes heavy use of effects.

Furthermore, Kameyama et al. take delimited control operators as the primitives for effects, which are not found in mainstream languages like Java. In order to track the use of delimited control operators, they must use an effect type system, which complicates types and typing rules. It is true that delimited control allows Kameyama et al. to express computations that are not easy to transcribe to assignment-based systems like Mint; however, the

advantage of their system is in the underlying imperative primitive and is not in an essential limitation of the weak separability idea itself.

Ancona and Moggi’s system executes effects within dynamic binders not at code generation time but at the time the generated code is run. In other words if dynamic binders are involved, Ancona and Moggi’s system cannot express effectful code generators although it can express generators that generate effectful code. Their system is similar to Aktemur’s [1] and Kim et al.’s [13] in that they use fresh names. All such calculi currently require explicitly listing the free variables in the code type, and explicitly managing hygiene and free variables explicitly in terms. Most importantly, the extra constraints in the type require polymorphism and structural subtyping to be introduced in the language. Mainstream object-oriented languages such as Java and C[#] support nominal subtyping, not structural subtyping.

9. Conclusion

This paper has proposed a practical approach to adding MSP to mainstream languages in a type-safe manner that prevents scope extrusion. The approach is simpler than prior proposals, and we expect that it will be easily and intuitively understood by programmers. The key insight is that safety can be ensured with weak separability, which places straightforward restrictions on the forms and types of computational effects that occur inside escape expressions, so that these effects cannot cause code to leak outside of escapes. The proposal has been validated both by proving that weak separability is enough to ensure safety and by demonstrating by example that many useful MSP applications can still be written that adhere to these restrictions.

A future direction for this work is to try to simplify the idea of weak separability to more closely match the intuition behind the concept. We believe there is some system similar to environment classifiers, in which quantifying on type variables can be used to implicitly capture the property that we wish to express. Instead of quantifying a type variable at the occurrence of `run()` as in environment classifiers, however, we believe that weak separability can be expressed by quantifying a type variable at the occurrence of an escape. This would simplify the type system and possibly add more expressive power to the language.

Acknowledgments

We thank Yannis Smaragdakis, Julia Lawall and Samuel Kamin for their helpful comments. We thank the anonymous reviewers for their valuable feedback.

References

- [1] Baris Aktemur. Type Checking Program Generators Using the Record Calculus, 2009. <http://loome.cs.uiuc.edu/pubs/transformationForTyping.pdf>.
- [2] Davide Ancona and Eugenio Moggi. A fresh calculus for name management. In *GPCE '04: Proceedings of the 3rd International Conference on Generative Programming and Component Engineering*, volume 3286, pages 206–224, 2004.
- [3] Cristiano Calcagno, Eugenio Moggi, and Walid Taha. Closed Types as a Simple Approach to Safe Imperative Multi-stage Programming. In *ICALP '00: Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, pages 25–36, 2000.
- [4] Cristiano Calcagno, Eugenio Moggi, and Walid Taha. ML-like inference for classifiers. In *ESOP '04: Proceedings of the 13th European Symposium on Programming*, pages 79–93, 2004.
- [5] Tim Davis and Yifan Hu. The University of Florida Sparse Matrix Collection. http://www.cise.ufl.edu/research/sparse/matrices/Pajek/Tina_DisCal.html.

- [6] Manuel Fähndrich, Michael Carbin, and James R. Larus. Reflective program generation with patterns. In *GPCE '06: Proceedings of the 5th International Conference on Generative Programming and Component Engineering*, pages 275–284, 2006.
- [7] Habanero Multicore Software Research Project. <http://habanero.rice.edu>.
- [8] Shan Shan Huang and Yannis Smaragdakis. Expressive and safe static reflection with MorphJ. In *PLDI '08: Proceedings of the 2008 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 79–89, 2008.
- [9] Shan Shan Huang, David Zook, and Yannis Smaragdakis. Statically safe program generation with safegen. In *GPCE '04: Proceedings of the 3rd International Conference on Generative Programming and Component Engineering*, pages 309–326, 2005.
- [10] Yuki Yoshi Kameyama, Oleg Kiselyov, and Chung chieh Shan. Closing the stage: from staged code to typed closures. In *PEPM '08: Proceedings of the 2008 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation*, pages 147–157, 2008.
- [11] Yuki Yoshi Kameyama, Oleg Kiselyov, and Chung chieh Shan. Shifting the stage: Staging with delimited control. In *PEPM '09: Proceedings of the 2009 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation*, pages 111–120, 2009.
- [12] Sam Kamin, Lars Clausen, and Ava Jarvis. Jumbo: Run-time code generation for Java and its applications. In *CGO '03: Proceedings of the International Symposium on Code Generation and Optimization*, pages 48–56, 2003.
- [13] Ik-Soon Kim, Kwangkeun Yi, and Cristiano Calcagno. A polymorphic modal type system for lisp-like multi-staged languages. In *POPL '06: Conference record of the 33rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 257–268, New York, NY, USA, 2006. ACM.
- [14] Gregory Neverov and Paul Roe. Metaphor: A Multi-stage, Object-Oriented Programming Language. In *GPCE '04: Proceedings of the 3rd International Conference on Generative Programming and Component Engineering*, pages 168–185, 2004.
- [15] Yutaka Oiwa, Hidehiko Masuhara, and Akinori Yonezawa. DynJava: Type safe dynamic code generation in Java. In *PPL '01: Proceedings of the 3rd JSSST Workshop on Programming and Programming Languages*, March 2001.
- [16] OpenJDK Project. <http://openjdk.java.net>.
- [17] Rice PLT. Mint Multi-stage Java Compiler. Available at <http://www.javamint.org>.
- [18] Massimiliano Poletto, Wilson C. Hsieh, Dawson R. Engler, and M. Frans Kaashoek. C and tcc: a language and compiler for dynamic code generation. *ACM Transactions on Programming Languages and Systems*, 21(2):324–369, 1999.
- [19] U.P. Schultz and J.L. Lawall C. Consel. Automatic Program Specialization for Java. *ACM Transactions on Programming Languages and Systems*, 25(4):452–499, 2003.
- [20] Peter Sestoft. Runtime code generation with JVM and CLR. Available at <http://www.dina.dk/sestoft/publications.html>, 2002.
- [21] Jun Shirako, Hironori Kasahara, and Vivek Sarkar. Language extensions in support of compiler parallelization. In *Languages and Compilers for Parallel Computing: 20th International Workshop, LCPC 2007, Urbana, IL, USA, October 11-13, 2007, Revised Selected Papers*, pages 78–94, Berlin, Heidelberg, 2008. Springer-Verlag.
- [22] Frederick Smith, Dan Grossman, J. Gregory Morrisett, Luke Hornof, and Trevor Jim. Compiling for template-based run-time code generation. *Journal of Functional Programming*, 13(3):677–708, 2003.
- [23] Rok Strniša, Peter Sewell, and Matthew Parkinson. The Java module system: Core design and semantic definition. In *OOPSLA '07: Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object Oriented Programming Systems and Applications*, pages 499–514, 2007.
- [24] Walid Taha. A gentle introduction to multi-stage programming. In *DSPG '03: Proceedings of the International Seminar on Domain-Specific Program Generation*, 2003.
- [25] Walid Taha, Zine el-abidine Benaïssa, and Tim Sheard. Multi-Stage Programming: Axiomatization and Type Safety (Extended Abstract). In *ICALP'98: 25th International Colloquium on Automata, Languages, and Programming*, pages 918–929, 1998.
- [26] Walid Taha and Michael Florentin Nielsen. Environment classifiers. In *POPL '03: Proceedings of the 30th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 26–37, New York, NY, USA, 2003. ACM.
- [27] Edwin Westbrook, Mathias Ricken, Jun Inoue, Yilong Yao, Tamer Abdelatif, and Walid Taha. Multi-stage Programming for Mainstream Languages. Technical Report TR09-02, Rice University, 2009.
- [28] David Zook, Shan Shan Huang, and Yannis Smaragdakis. Generating AspectJ Programs with Meta-AspectJ. In *GPCE '04: Proceedings of the 3rd International Conference on Generative Programming and Component Engineering*, pages 1–18, 2004.