

Program Generation and Components

D. Ancona and E. Moggi*

DISI, Univ. of Genova, v. Dodecaneso 35, 16146 Genova, Italy
{davide, moggi}@disi.unige.it

Abstract. The first part of the paper gives a brief overview of meta-programming, in particular program generation, and its use in software development. The second part introduces a basic calculus, related to FreshML, that supports program generation (as described through examples and a translation of MetaML into it) and programming in-the-large (this is demonstrated by a translation of CMS into it).

1 Introduction

This paper explains what is program generation and what are the most promising uses of it, recalls the role of names in software components, and then presents a basic calculus, called MML_ν^N , which in the authors' view is a suitable formalism to describe program generation in terms of more primitive notions, namely name generation, name resolution and linking.

In calculi of module systems as CMS [AZ99, MT00, WV00, AZ02] modules can refer to *deferred* components by means of names. These calculi provide primitive operators for linking modules and resolving external names of deferred components, thus supporting programming in-the-large [Car97]. Analogously, in the MML_ν^N calculus of [AM04] names are used to refer to external components which need to be provided from the outside (by a name resolver).

In module (and record) calculi names are taken from some infinite set. On the contrary, in MML_ν^N at any time during execution there is a finite set of names, which can be extended dynamically by a construct $\nu X.e$ for generating a *fresh* name, borrowed from FreshML of [SPG03].

Fraenkel and Mostowski's set theory (see [GP99]) provides the mathematical underpinning of name generation for FreshML and MML_ν^N , but to understand the operational semantics and type system there is no need to be acquainted with FM-sets. Besides names $X \in \text{Name}$, the calculus has

- terms $e \in \mathbf{E}$, a closed term corresponds to an *executable* program;
- name resolvers, which denote partial functions $\text{Name} \xrightarrow{\text{fin}} \mathbf{E}$ with finite domain.

We write $r.X$ for the term obtained by applying r to *resolve* name X .

Terms include fragments $\mathbf{b}(r)e$, i.e. term e abstracted w.r.t. resolver r , which denote functions $(\text{Name} \xrightarrow{\text{fin}} \mathbf{E}) \rightarrow \mathbf{E}$. We write $e\langle r \rangle$ for the term obtained by *linking* fragment e using resolver r .

* Supported by EU projects DART IST-2001-33477 and APPSEM-II IST-2001-38957.

Remark 1. If resolvers were included in terms, we would get a λ -calculus with extensible records [CM94]; indeed, a record amounts to a partial function mapping names (of components) to their values. More precisely, $\mathbf{b}(r)e$ would become an abstraction $\lambda r.e$ and $e\langle r \rangle$ an application $e r$. Even when resolvers are second class terms, one can express in the calculus the staging constructs of MetaML [Tah99, She01] and the mixin operations of CMS.

The ability to generate a fresh name is essential to prevent accidental overriding. If we know in advance what names need to be resolved within a fragment (we call such a fragment closed), then we can statically choose a name which is fresh (for that fragment). However, generic functions manipulating *open* fragments will have to generate fresh names at run-time. There are several reasons for working with open fragments: reusability is increased, the need for naming conventions (between independent developers) is reduced, and decisions can be delayed.

We present MML_{ν}^N as a monadic metalanguage, i.e. its type system makes explicit which terms have computational effects, and its operational semantics is given by a (semantic preserving) simplification relation on terms and a computation relation on *configurations*. Generation of fresh names is a computational effect, as in FreshML, thus typing $\nu X.e$ requires computational types.

Summary. Section 2 explains program generation within the broader context of meta-programming, and mentions its most promising uses. Section 3 recalls the role of names in software components. Section 4 recalls syntax, type system and operational semantics of MML_{ν}^N . Section 5 gives several programming examples: programming with open fragments, and benchmark examples for comparison with other calculi (in particular MetaML). Section 6 introduces a 2-level version of MetaML, and show how it can be translated into MML_{ν}^N . Section 7 introduces ML_{Σ}^N , a variant of MML_{ν}^N with records and recursive definitions, where one can recover all (mixin) module operations of CMS. Finally, Section 8 discusses related calculi based on names, i.e. FreshML of [SPG03] and ν^{\square} of [NPar].

Notation. In the paper we use the following notations and conventions.

- m ranges over the set \mathcal{N} of natural numbers. Furthermore, $m \in \mathcal{N}$ is identified with the set $\{i \in \mathcal{N} \mid i < m\}$ of its predecessors.
- Term equivalence \equiv is α -conversion. $\text{FV}(e)$ is the set of variables free in e , while $e[x_i : e_i \mid i \in m]$ denotes parallel capture avoiding substitution.
- $f : A \xrightarrow{\text{fin}} B$ means that f is a partial function from A to B with a finite domain, written $\text{dom}(f)$. The image of f is denoted by $\text{img}(f)$. $A \rightarrow B$ denotes the set of total functions from A to B . We use the following operations:
 - $\{a_i : b_i \mid i \in m\}$ is the partial function mapping a_i to b_i (where the a_i must be different, i.e. $a_i = a_j$ implies $i = j$);
 - \emptyset is the everywhere undefined partial function;
 - $f_{\setminus a}$ denotes the partial function f' s.t. $f'(a') = b$ iff $b = f(a')$ and $a' \neq a$;

- $f\{a:b\}$ denotes the (partial) function f' s.t. $f'(a) = b$ and $f'(a') = f(a')$ when $a' \neq a$;
 - f, f' denotes the union of two partial functions with disjoint domains.
- $A\#B$ means that the sets A and B are disjoint.

2 Program Generation

We explain what is program generation by placing it in the broader context of meta-programming. We borrow from Sheard's invited talk at SAIG'01 [She01], which in addition discusses several areas of meta-programming research. Then we mention some of the most promising uses of program generation in the context of software development. We make no attempt to be exhaustive, instead we advise the interested reader to browse through the proceedings of the conference on Generative Programming and Component Engineering (GPCE) [BCT02, PS03, KV04], and a compendium [LBCO04] of contributions presented at a Dagstuhl seminar on "Domain-specific program generation", where a new IFIP WG on "Program Generation" was proposed.

2.1 What Is It?

In general programs manipulate data. Meta-programs are programs that manipulate object-programs, or more precisely data representing other programs. We are not committed to a particular (programming) language, thus by object-program we mean a syntactic element in a formal language. Broadly speaking we can classify meta-programs in three categories:

- Generators, which construct object-programs. For instance, specializers generate a specialized program solving an instance of a general problem.
- Analyzers, which analyze the structure of object-programs. For instance, type-checkers or tools that perform various kinds of static analysis.
- Transformers, which combine the features of analyzers and generators. For instance, optimizers that perform source-to-source transformation, or aspect weavers that insert code to address a cross-cutting concern.

A compiler is a typical example of program where the three kinds of meta-programs co-exist: a static analyzer for the source language, an optimizer for some intermediate language, and a program generator for the target language.

Object-programs can be represented at different level of abstraction (we call *code* the data representing a program):

- White-box abstraction, where code is text (i.e. a string) or an abstract syntax tree (AST). This representation is the most versatile, since it gives a low-level description of the object-program, but could be error prone.
- White-box abstraction modulo α -conversion. This representation of syntax has been considered in the context of logical frameworks to deal with binders in object languages, alongside other representations like higher-order abstract syntax. FreshML [GP99, SPG03] is the leading programming language that supports this abstraction.

- Black-box abstraction, where code can be executed and combined, but not analyzed. This is the most abstract representation.

The black-box abstraction is incompatible with program analyzers and transformers. On the other hand, program generators can work with any of the abstractions, and the black-box abstraction ensures the maximum separation of concerns between the generator and the user of the generated code.

We consider related concepts to clarify how they differ from the concept of code and meta-program. A program *configuration* is a snapshot of a program during execution, a *computation* is a description of the program execution, while code represents (the syntax of) a program independently from execution. A reflective program is a program that manipulates *itself*, thus it is a particular instance of a meta-program. One can identify three forms of reflection:

- Introspection is the ability of a program to analyze itself, namely its code (this introspection is called structural reflection) or its current configuration or execution history (this introspection is called behavioral reflection).
- Self-modification is the ability to modify itself.
- Intercession is the ability to manipulate its semantics, i.e. the interpreter or virtual machine for the reflective program.

A computation is *staged* when it is decomposed into stages along the temporal dimension. The change of stage is usually triggered by the acquisition of new information. In a meta-programming system supporting program generation there are two natural stages: the computation of the meta-program and the computation of the generated object-program. Depending on the nature of the meta-program, its computation is called differently (e.g. generation-time, compile-time, design-time, specialization-time), while the computation of the object-program is usually called run-time or use-time computation. Moreover, if the meta-programming system is *homogeneous*, i.e. the object-language for describing object-programs coincides with the meta-language, then it provides a natural support for *multi-stage* programming. MetaML [TS97] and MetaOCaml [CTHL03, Met01] are among the leading multi-stage programming languages. For several applications *heterogeneous* meta-programming systems are enough. In this case, the main issue is to provide support for a variety of object-languages.

2.2 What Is for?

Generative and component approaches have the potential to revolutionize software development in a similar way as automation and components revolutionized manufacturing. Generative programming (developing programs that synthesize other programs), component engineering (raising the level of modularization and analysis in application design), and domain-specific languages (elevating program specifications to compact domain-specific notations that are easier to write and maintain) are key technologies for automating program development. Before focusing on program generation, we mention some trends in software engineering, in order to provide a broader picture. [GS04] identifies *software factories* as the next methodology for software development. A software factory is

a collection of *reusable* assets (like patterns, models, frameworks, tools) for *rapidly* and *cheaply* producing an *open-ended* set of unique variants of a software *product*.

Clearly, for a software product that has a big market and need to evolve, like an operating system, this is an economically feasible approach. However, to make software factories economically feasible for specific domains, it is essential to empower the domain-experts and end-users.

Domain-specific languages (DSLs) are a way to give programming abilities to domain-experts and end-users. Descriptions given in a DSL can be treated as high-level source code, rather than non-executable requirement specifications. Relational databases are a “classic” example of success story in the use of DSLs, while UML is a counter-example of DSL. In fact, UML is too general (i.e. it is not meant for any specific domain) and too imprecise (e.g. it cannot be used as source code, although some subsets might). Other examples of domains where DSL technology has been used successfully or appear highly promising are: language parsers, reactive real-time programs and the telephony domain. For instance, MetaCase Consulting (<http://www.metacase.com>) gave a demo at GPCE’03 entitled “MetaEdit+ revolutionized the way Nokia develops mobile phone software”, and the choice of MetaEdit+ was motivated as follows

When Nokia was searching for an effective CASE tool, the prime criteria was encapsulation of domain knowledge, flexible method support and code generation. After evaluating a number of off-the-shelf CASE tools, they undertook the development of their own solution using the MetaEdit+ metaCASE tool.

DSLs provide language-based abstraction, which goes well beyond the abstraction provided by libraries (i.e. platform extension). In this context program generators play the role of compiler back-ends for domain-specific languages, and provide the following benefits

- Increase automation by exploiting domain features and knowledge.
- Improve performance via partial evaluation.

Usually program generators are *co-designed* with the DSL they implement, unlike compilers for general-purpose languages. It is worth to adopt this approach, when the effort to implement a program generator is comparable to that of implementing a software library for the specific domain. In the DSL approach there are other things one can do before the program generation stage, namely

- analysis, which exploits domain-knowledge to identify problems prior generation, or to provide static guarantees
- transformation, for instance to perform domain-specific optimization.

It is important that analysis and transformation take place before program generation, so that they can be understood and managed by the user of the DSL, who can be expected to be knowledgeable of the domain but not of the target language for the program generator.

3 Names and Software Components

It has been argued that the notion of software component is so general that cannot be defined in a precise and comprehensive way [CE00]. For instance, [Szy02] provides three different definitions, that adopt different levels of abstraction. Nevertheless, names play an important role, independently of the particular definition adopted for software components. For clarity, we identify the notion of software component with that of *mixin module*¹, as done in CMS of [AZ02].

A basic module could be described as follows

```
import  $X_1$  as  $x_1, \dots, X_m$  as  $x_m$ 
export  $Y_1 = E_1, \dots, Y_n = E_n$ 
local  $z_1 = E'_1, \dots, z_p = E'_p$ 
```

A basic module make use of *names* and *variables*. The former are the names of the components the module either *imports* from the outside (*input* components X_1, \dots, X_m) or *exports* to the outside (*output* components Y_1, \dots, Y_n). The latter are the variables used in definitions inside the module (i.e. in the expressions $E_1, \dots, E_n, E'_1, \dots, E'_p$). These variables can be either *deferred* (x_1, \dots, x_n), i.e. associated with some input component, or locally defined (z_1, \dots, z_p).

The distinction between names and variables is crucial: names correspond to external references, while variables correspond to internal references. Technically speaking the main differences between variables and names are: expressions include variables but not names; variables declared in a basic module are local and can be α -converted; while component names belong to a global name space and allow modules to talk to each other.

A useful operator which can be easily encoded in CMS is the *link* operator $link(M_1, M_2)$, used for merging two modules and resolving input names. This operator may be regarded as either an operation provided by a module language in order to define structured module expressions or an extra-linguistic mechanism to combine object files provided by a tool for modular software development. $link(M_1, M_2)$ is well-defined if the sets of the output components of M_1 and M_2 are disjoint. In this case, $link(M_1, M_2)$ corresponds to a module where some input component of one module has been bound to the definition of the corresponding output component of the other module, and conversely.

For instance, let the modules **BOOL** and **INT** define the evaluation of some boolean and integer expressions in a mutually recursive way:

```
module BOOL is
  import IntEv as ext_ev
  export BoolEv = ev
  local
    fun ev EQ(ie1,ie2) = ext_ev(ie1)==ext_ev(ie2)
    | ...
end BOOL;
```

¹ In the sequel we will interchangeably use the terms “module” and “mixin” as abbreviations of “mixin module”.

```

module INT is
  import BoolEv as ext_ev
  export IntEv = ev
  local
    fun ev IF(be,ie1,ie2) = if ext_ev(be) then ev(ie1) else ev(ie2)
    | ...
end INT;

```

The result of *link*(BOOL,INT) corresponds to the module

```

module BOOL_INT is
  export IntEv = iev
  export BoolEv = bev
  local
    fun bev EQ(ie1,ie2) = iev(ie1)==iev(ie2)
    | ...
    fun iev IF(be,ie1,ie2) = ifbev(be) then iev(ie1) else iev(ie2)
    | ...
end BOOL_INT;

```

The separation between component names and variables allows one to use internally the same name *ev* for the evaluation function in the two modules; in the compound module, indeed, *ev* of BOOL and *ev* of INT are α -renamed to *bev* and *iev*, respectively.

The *link* operation described above can be decomposed in two steps. First, put together the declarations of the two arguments in one module, yielding

```

module
  import IntEv as ext_iev
  import BoolEv as ext_bev
  export IntEv = iev
  export BoolEv = bev
  local
    fun bev EQ(ie1,ie2) = ext_iev(ie1)==ext_iev(ie2)
    | ...
    fun iev IF(be,ie1,ie2) = if ext_bev(be) then iev(ie1) else iev(ie2)
    | ...
end;

```

Then, bind import components with export components with the same name, yielding BOOL_INT. Formally, this corresponds to the fact that *link* is a derived operator which can be expressed by the *sum* and *freeze* basic operators of CMS.

CMS provides also a primitive operation for deleting module components, which allows redefinition of components when used in conjunction with the *link* operator. This is an important feature for enabling reuse of software components and amortizing the investment over multiple applications [Szy02].

4 A Core Calculus with Names: MML_ν^N

This section recalls the monadic metalanguage MML_ν^N of [AM04]. For simplicity, we focus on the key feature, i.e. *names*, and exclude imperative computations

and functional types. Moreover, we restrict the formal treatment to a simply typed language (Section 4.1), and recall only the main statements concerning type safety (Section 4.4). We refer to [AM04] for details and a polymorphic extension of the type system, which is essential for typing the examples on open fragments generators (see Example 1 in Section 5). The operational semantics is given according to the general pattern proposed in [MF03], namely by a confluent *simplification* relation \longrightarrow defined as the *compatible closure* of a set of rewrite rules (see Section 4.2), and a *computation* relation \mapsto describing how *configurations* may evolve (see Section 4.3).

Names X are syntactically pervasive, i.e. they occur both in types and in terms. The term $\nu X.e$ allows to generate a *fresh* name for private use within e . Following FreshML of [SPG03], we consider generation of a fresh name a computational effect, therefore for typing $\nu X.e$ we need computational types.

We parameterize typing judgments w.r.t. a finite set of names, namely those that can occur (free) in the judgment. The mathematical underpinning for names is provided by [GP99]. In particular, properties are invariant w.r.t. name permutation (equivariance), but not w.r.t. name substitution.

The syntax of MML_{ν}^N is abstracted over symbolic names $X \in \text{Name}$, basic types b , term variables $x \in X$ and resolver variables $r \in R$. The syntactic category of types and signatures (i.e. the types of resolvers) is parameterized w.r.t. a finite set $\mathcal{X} \subseteq_{fin} \text{Name}$ of names that can occur in the types and signatures.

- $\tau \in \mathbb{T}_{\mathcal{X}} ::= b \mid [\Sigma|\tau] \mid M\tau$ \mathcal{X} -types, where
 $\Sigma \in \Sigma_{\mathcal{X}} \triangleq \mathcal{X} \xrightarrow{fin} \mathbb{T}_{\mathcal{X}}$ is a \mathcal{X} -signature $\{X_i: \tau_i \mid i \in m\}$
- $e \in \mathbb{E} ::= x \mid \theta.X \mid e\langle\theta\rangle \mid \mathbf{b}(r)e \mid \mathbf{ret} \ e \mid \mathbf{do} \ x \leftarrow e_1; e_2 \mid \nu X.e$ terms, where
 $\theta \in \mathbb{ER} ::= r \mid ? \mid \theta\{X:e\}$ is a name resolver term.

We give an informal semantics of the language (see Section 5 for examples).

- The type $[\Sigma|\tau]$ classifies fragments which produce a term of type τ when linked with a resolver for Σ . The terms $\theta.X$ and $e\langle\theta\rangle$ use θ to *resolve* name X and to *link* fragment e . The term $\mathbf{b}(r)e$ *represents* the fragment obtained by abstracting e w.r.t. r .
- The resolver $?$ cannot resolve any name, while $\theta\{X:e\}$ resolves X with e and *delegates* the resolution of other names to θ .
- The monadic type $M\tau$ classifies programs computing values of type τ . The terms $\mathbf{ret} \ e$ and $\mathbf{do} \ x \leftarrow e_1; e_2$ are used to terminate and sequence computations, $\nu X.e$ generates a *fresh* name for use within the computation e .

As a simple example, let us consider the fragment $\mathbf{b}(r)(r.X*r.X)$ which can be correctly linked with resolvers mapping X to integer expressions and whose type is $[X: \text{int} \mid \text{int}]$. Then we can link the fragment with the resolver $?\{X:2\}$, as in $\mathbf{b}(r)(r.X*r.X)\langle?\{X:2\}\rangle$, and obtain $2*2$ of type int . Note that $\mathbf{b}(r)(r.X*r.X)$ is not equivalent to $\mathbf{b}(r)(r.Y*r.Y)$, whose type is $[Y: \text{int} \mid \text{int}]$. This is in clear contrast with what happens with variables and λ -abstractions: $\lambda x \rightarrow x*x$ and $\lambda y \rightarrow y*y$ are equivalent and have the same type. The sequel of this

section is devoted to the formal definition of MML_{ν}^N . More interesting examples (with informal explanatory text) can be found in Section 5.

One can define (by induction on τ , e and θ) the following syntactic functions:

- the set $\text{FV}(_)\subseteq_{fin} \mathbf{Name}\uplus\mathbf{X}\uplus\mathbf{R}$ of free names and variables in $_$, in particular $\text{FV}(\{X_i:\tau_i|i\in m\}) = (\cup_{i\in m}\text{FV}(\tau_i)) \cup \{X_i|i\in m\}$
- the capture-avoiding substitution $_ [x_0:e_0]$ for term variable x_0 .
- the capture-avoiding substitution $_ [r_0:\theta_0]$ for resolver variable r_0 .
- the action $_ [\pi]$ of a name permutation π on $_$.

4.1 Type System

The typing judgments are $\mathcal{X};\Pi;\Gamma\vdash e:\tau$ (i.e. e has type τ) and $\mathcal{X};\Pi;\Gamma\vdash\theta:\Sigma$ (i.e. θ resolves the names in the domain of Σ , and only them, with terms of the assigned type), where

- τ is a \mathcal{X} -type and Σ is a \mathcal{X} -signature
- $\Pi:\mathbf{R}\xrightarrow{fin}\Sigma_{\mathcal{X}}$ is a signature assignment $\{r_i:\Sigma_i|i\in m\}$ for resolver variables
- $\Gamma:\mathbf{X}\xrightarrow{fin}\mathbf{T}_{\mathcal{X}}$ is a type assignment $\{x_i:\tau_i|i\in m\}$ for term variables

The typing rules are given in Table 1. All the rules, except that for $\nu X.e$, use the same finite set \mathcal{X} of names in the premises and the conclusion. The typing rule for $e\langle\theta\rangle$ supports a limited form of *width* subtyping, namely it allows linking of a fragment $e:[\Sigma|\tau]$ with a resolver θ whose signature Σ' includes Σ . All the other rules are standard.

4.2 Simplification

We define a confluent relation on terms, called *simplification*. There is no need to define a deterministic simplification strategy, since computational effects are *insensitive* to further simplification. Simplification \longrightarrow is the compatible closure of the following rules

- (resolve) $(\theta\{X:e\}).X\longrightarrow e$
- (delegate) $(\theta\{X:e\}).X'\longrightarrow\theta.X'$ if $X'\neq X$
- (link) $(b(r)e)\langle\theta\rangle\longrightarrow e[r:\theta]$

Simplification enjoys the following properties.

Theorem 1 (Church-Rosser). *The simplification relation \longrightarrow is confluent.*

Theorem 2 (Subject Reduction).

- If $\mathcal{X};\Pi;\Gamma\vdash e:\tau$ and $e\longrightarrow e'$, then $\mathcal{X};\Pi;\Gamma\vdash e':\tau$.
- If $\mathcal{X};\Pi;\Gamma\vdash\theta:\Sigma$ and $\theta\longrightarrow\theta'$, then $\mathcal{X};\Pi;\Gamma\vdash\theta':\Sigma$.

Table 1. Type System for MML_{ν}^N

$x \frac{}{\mathcal{X}; \Pi; \Gamma \vdash x: \tau} \Gamma(x) = \tau$	resolve $\frac{\mathcal{X}; \Pi; \Gamma \vdash \theta: \Sigma}{\mathcal{X}; \Pi; \Gamma \vdash \theta.X: \tau} \tau = \Sigma(X)$
link $\frac{\mathcal{X}; \Pi; \Gamma \vdash e: [\Sigma]\tau}{\mathcal{X}; \Pi; \Gamma \vdash e(\theta): \tau} \Sigma \subseteq \Sigma'$	box $\frac{\mathcal{X}; \Pi, r: \Sigma; \Gamma \vdash e: \tau}{\mathcal{X}; \Pi; \Gamma \vdash \mathbf{b}(r)e: [\Sigma]\tau}$
$r \frac{\Pi(r) = \Sigma}{\mathcal{X}; \Pi; \Gamma \vdash r: \Sigma}$? $\frac{}{\mathcal{X}; \Pi; \Gamma \vdash ? : \emptyset}$
ret $\frac{\mathcal{X}; \Pi; \Gamma \vdash e: \tau}{\mathcal{X}; \Pi; \Gamma \vdash \mathbf{ret} e: M\tau}$	extr $\frac{\mathcal{X}; \Pi; \Gamma \vdash \theta: \Sigma}{\mathcal{X}; \Pi; \Gamma \vdash \theta\{X: e\}: \Sigma\{X: \tau\}}$
	do $\frac{\mathcal{X}; \Pi; \Gamma \vdash e_1: M\tau_1}{\mathcal{X}; \Pi; \Gamma \vdash \mathbf{do} x \leftarrow e_1; e_2: M\tau_2}$
	$\nu \frac{\mathcal{X}, X; \Pi; \Gamma \vdash e: M\tau}{\mathcal{X}; \Pi; \Gamma \vdash \nu X.e: M\tau} X \notin \text{FV}(\Pi, \Gamma, \tau)$

4.3 Computation

The computation relation $Id \mapsto Id' \mid \text{done}$ is defined using evaluation contexts and configurations $Id \in \text{Conf}$. A configuration records the current name space as a finite set \mathcal{X} of names. The computation rules (see Table 2) consist of those given in [MF03] for the monadic metalanguage MML (these rules do not change the name space) plus a rule for generation of a fresh name (this is the only rule that extends the name space).

- $\boxed{E \in \text{EC} ::= \square \mid E[\mathbf{do} x \leftarrow \square; e]}$ evaluation contexts
- $(\mathcal{X} \mid e, E) \in \text{Conf} \triangleq \mathcal{P}_{\text{fin}}(\text{Name}) \times \mathbf{E} \times \text{EC}$ configurations consist of the current name space \mathcal{X} (which grows as computation progresses), the program fragment e under consideration, and its evaluation context E
- $\boxed{rc \in \text{RC} ::= \mathbf{ret} e \mid \mathbf{do} x \leftarrow e_1; e_2 \mid \nu X.e}$ computational redexes.

Simplification \longrightarrow is extended in the obvious way to a confluent relation on configurations (and related notions). The bisimulation property, i.e. computation is insensitive to further simplification, is like that stated in [MF03] for MML.

Theorem 3 (Bisimulation). *If $Id \equiv (\mathcal{X} \mid e, E)$ with $e \in \text{RC}$ and $Id \xrightarrow{*} Id'$, then*

1. $Id \mapsto D$ implies $\exists D'$ s.t. $Id' \mapsto D'$ and $D \xrightarrow{*} D'$
2. $Id' \mapsto D'$ implies $\exists D$ s.t. $Id \mapsto D$ and $D \xrightarrow{*} D'$

where D and D' range over $\text{Conf} \cup \{\text{done}\}$.

Table 2. Computation Relation

Administrative steps	
(A.0)	$(\mathcal{X} \text{ret } e, \square) \mapsto \text{done}$
(A.1)	$(\mathcal{X} \text{do } x \leftarrow e_1; e_2, E) \mapsto (\mathcal{X} e_1, E[\text{do } x \leftarrow \square; e_2])$
(A.2)	$(\mathcal{X} \text{ret } e_1, E[\text{do } x \leftarrow \square; e_2]) \mapsto (\mathcal{X} e_2[x: e_1], E)$
Name generation step	
(ν)	$(\mathcal{X} \nu X.e, E) \mapsto (\mathcal{X}, X e, E)$ with X renamed to avoid clashes, i.e. $X \notin \mathcal{X}$

Table 3. Well-formed Evaluation Contexts

\square	$\frac{\mathcal{X}; \square: M\tau_2 \vdash E: M\tau' \quad \mathcal{X}; \emptyset; x: \tau_1 \vdash e: M\tau_2}{\mathcal{X}; \square: M\tau_1 \vdash E[\text{do } x \leftarrow \square; e]: M\tau'}$
-----------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4.4 Type Safety

Following Felleisen, type safety can be decomposed in two properties: subject reduction and progress. We refer to [MF03] for a formulation of these properties in the context of a monadic metalanguage.

Definition 1 (Well-formed configuration). $\vdash (\mathcal{X}|e, E): \tau' \stackrel{\Delta}{\iff} \tau' \in \mathbb{T}_\emptyset$ and $\exists \tau \in \mathbb{T}_\mathcal{X}$ s.t. $\mathcal{X}; \emptyset; \emptyset \vdash e: M\tau$ and $\mathcal{X}; \square: M\tau \vdash E: M\tau'$ (see Table 3).

Theorem 4 (Subject Reduction).

- If $\vdash Id_1: \tau'$ and $Id_1 \longrightarrow Id_2$, then $\vdash Id_2: \tau'$.
- If $\vdash Id_1: \tau'$ and $Id_1 \mapsto Id_2$, then $\vdash Id_2: \tau'$.

Theorem 5 (Progress). If $\vdash (\mathcal{X}|e, E): \tau'$, then

1. either $e \notin \text{RC}$ and $e \longrightarrow$
2. or $e \in \text{RC}$ and $(\mathcal{X}|e, E) \mapsto$

5 Programming Examples

We demonstrate the use and expressivity of MML_ν^N with a few examples:

- the first exemplifies programming with *open* fragments;
- the second recasts the multi-stage programming method of [TS97] by making a simplified use of *open* fragments;

- the third uses *closed* fragments, and allows a further comparison with other calculi for run-time code generation and staging.

To improve readability we use ML-like notation for functions (β -reduction is a sound simplification in monadic metalanguages) and operations on references, and Haskell’s do-notation `do {x1 <- e1; ...; xn <- en; e}`. In the sequence of commands of a do-expression we allow computations `ei` whose value is not bound to a variable (because it is not used by other commands) and non-recursive let-bindings like `xi = ei` (which amounts to replace `xi` with `ei` in the commands following the let-binding).

Example 1. We consider an example of generative programming, which motivates the need for fresh name generation. In our calculus, a component is identified with a fragment of type $[\Sigma|\tau]$, where Σ specifies what information needs to be provided for deployment. Generative programming supports dynamic manufacturing of customized components from elementary (highly reusable) components. The most appropriate building block for generative programming are polymorphic functions $G: \forall p. [p, \Sigma_i | \tau_i] \rightarrow M[p, \Sigma | \tau]$ (we refer to [AM04] for a polymorphic extension of the type system). The result type of G is computational, because generation may require computational activities, while the *signature variable* p classifies the information passed to the parameters of G , but not directly used or provided in the implementation of G itself. Applications of G may instantiate p with different signatures, thus we say that G manipulates *open* fragments. An over-simplified example of open fragment generator is

```
Ac: [p|a->a] -> M[p|{add: a -> M unit, update: M unit}]
```

`Ac` creates a data structure to maintain an (initially empty) set of accounts. Since we don’t really need to know the structure of an account, we use a type variable `a`. The generator makes available two functionalities for operating on a set (of accounts): `add` inserts a new account in the set, and `update` modifies all the accounts in the set by applying a function of type `a->a`, which depends on certain parameters (e.g. the interest rate) represented by the signature variable `p`. These parameters are decided by the bank after the data structure has been created, and they change over time.

In many countries bank accounts are taxed, according to local criteria. So we need a more refined generator, with an extra parameter for computing the new balance based on the state of the account after the bank’s update

```
TaxedAc: [p'|a->a] -> [p|a->a] ->
          M[p'|[p|{add: a -> M unit, update: M unit}]]
```

The signature variable `p'` classifies the information needed to compute local taxes. In general `p'` and `p` are unrelated, and identifying them means that banks and local authorities rely on the same information. `TaxedAc` is defined as follows

```
fun TaxedAc tax upd = nu Tax.
  do {m <- Ac(b(r2) fn x => r2.Tax (upd<r2> x));
      ret (b(r') b(r1) m<r1{Tax:tax<r'>>})};
```

It is essential that the name `Tax` is fresh and private to `TaxedAc`, otherwise we may override some information in `r1`, which is needed by `upd`. In fact, `TaxedAc` is an open fragment generator that does not know in advance how the signature variable `p` could be instantiated. On the other hand, with *closed* fragment generators $G: [\Sigma_i | \tau_i] \rightarrow M[\Sigma | \tau]$ the problem does not arise, but reusability is impaired. For instance, it is not reasonable to expect that all banks will use the same parameters to update the accounts of their customers.

Example 2. We recast the *multi-stage programming method* of [TS97] (see also [CMS03]) using the power function, which is a classical example for staged programming.

1. The method starts from a “conventional” program `exp` with two parameters. In the specific example, `exp` takes an exponent n , a base x , and then computes x^n by making recursive calls

```
fun exp n x = if n=0 then ret(1.0)
              else do {y <- (exp (n-1) x); ret(x*y)};
> exp = ... : int -> real -> M real
```

The result type of `exp` is computational, because we consider recursion a computational effect.

2. Then one obtains a “staged” version `exp_a`, which replaces the second parameter (the base x) with an open fragment u , and builds an open fragment representing the desired result (i.e. x^n)

```
fun exp_a n u = if n=0 then ret(b(r) 1.0)
                 else do {v <- exp_a (n-1) u; ret(b(r) u<r>*v<r>)};
> exp_a = ... : int -> [p|real] -> M[p|real]
```

The staged version is polymorphic in the signature variable `p`.

3. By exploiting the polymorphism of `exp_a`, one defines a *code generator* `exp_cg`. Given the base n , the generator calls `exp_a` with a “dummy” parameter `b(r) r.X`, then builds an open fragment representing a function

```
fun exp_cg n = nu X. do {v <- exp_a n (b(r) r.X);
                       ret (b(r) fn x => v<r>{X:x}>)};
> exp_cg = ... : int -> M[p|real -> real]
```

The type of `exp_cg` says that recursion is unfolded at “specialization” time, when the exponent n is known.

4. By instantiating `p` with the empty signature, one gets an *optimized* program

```
fun exp_o n = do {v <- exp_cg n; ret(v<?>)};
> exp_o = ... : int -> M(real -> real)
```

The type of `exp_o` differs from the type of the conventional program `exp` to reflect the different timing in unfolding recursion. Namely, `exp_o` unfolds the recursion when the parameter n is known. For instance, when $n = 2$

```
do sq_o <- exp_o 2;
> sq_o = (fn x => x*(x*1.0)) : real -> real
```

The multi-stage programming method makes use of open fragments of type $[p|\tau]$ (these types are similar to the code types annotated with environment classifiers $\langle\tau\rangle^\alpha$ used by [TN03, CMT04]). One can easily recast the multi-stage programming method also in the presence of more complex computational effects (while in MetaML there are typing problems). For instance, when the conventional program is an imperative variant $p:\text{int}\rightarrow\text{real}\rightarrow(\text{real ref})\rightarrow M \text{ unit}$ of exp

```
fun p n x y = if n=0 then y:=1.0
              else do {p (n-1) x y; y' <- !y; y:=x*y'};
> p = ... : int -> real -> (R real) -> M unit
```

p takes an exponent n , a base x and a reference y , then it initializes y with 1.0 and repeatedly multiplies the content of y with x until it becomes x^n . The “staged” version, p_a , is defined in the obvious way, and its type says that some computations are postponed to the second stage

```
fun p_a n u v= if n=0 then ret(b(r) v<r>:=1.0)
               else do {
                   w <- p_a (n-1) u v;
                   ret(b(r) do {w<r>; y' <- !v<r>; v<r>:=u<r>*y'})};
> p_a = ... : int -> [p|real] -> [p|Ref real] -> M[p|M unit]
```

In comparison to MetaML, we don’t face the problems due to execution of *potentially open* code or *scope extrusion*, which motivated the introduction of closed types in [CMS03]. The reason is that in MML_ν^N one has a better control of the name space and name resolution.

Example 3. We reconsider the power function $\text{exp}:\text{int}\rightarrow\text{real}\rightarrow M \text{ real}$, and give an alternative way to define $\text{exp}_o:\text{int}\rightarrow M(\text{real}\rightarrow \text{real})$, which does not involve fresh name generation.

```
(* conventional program *)
fun exp n x = if n=0 then ret(1.0)
              else do {y <- (exp (n-1) x); ret(x*y)};
> exp = ... : int -> real -> M real
(* staged program *)
fun exp_a n u = if n=0 then ret(b(r) 1.0)
                else do {v <- exp_a (n-1) u; ret(b(r) u<r>*v<r>)};
> exp_a = ... : int -> [p|real] -> M[p|real]
(* exp_c generates a fragment with hook X for base *)
fun exp_c n = do {v <- exp_a n (b(r) r.X);
                  ret (b(r) fn x => v<r>)};
> exp_c = ... : int -> M[X:real|real]
(* optimized program *)
fun exp_o n = do {v <- exp_c n; ret(fn x => u<?{X:x}>)};
> exp_o = ... : int -> M(real -> real)
```

The definition of `exp_c` relies on a pre-existing name X , while `exp_cg` uses a freshly generated name. `MetaML` does not allow to mention names explicitly, thus it has no analogue of `exp_c` nor of the type $[X:\text{real}|\text{real}]$. On the other hand, ν^\square has an analogue of `exp_c` (see `exp'` in [NPar, Example 2]), but the name X has to be declared of type `real` globally.

6 Relating MML_ν^N to `MetaML`

In this section we define a monadic CBV translation of a 2-level version of `MetaML` into MML_ν^N (extended with functional types), and show that the translation preserves the operational semantics. We make no formal claim about preservation of typing, since we have not been able to extend the translation to types. We have not defined a monadic CBV translation of the whole `MetaML`, since key ideas would get confused with orthogonal issues involved in the translation of a multi-level language. Restricting to a 2-level language allows to bring these ideas in the foreground.

6.1 `MetaML`₂

We give the formal definition (syntax, a simplified type system and big-step CBV operational semantics) of `MetaML`₂, a 2-level version of `MetaML`. As customary for 2-level languages, the syntax (type system and operational semantics) is stratified in two levels: the meta-level 0, and the object-level 1.

- $\begin{array}{l} \tau^0 \in \mathbb{T}^0 ::= b \mid \tau_1^0 \rightarrow \tau_2^0 \mid \langle \tau^1 \rangle \\ \tau^1 \in \mathbb{T}^1 ::= b \mid \tau_1^1 \rightarrow \tau_2^1 \end{array}$ types at level 0 and 1, note that $\mathbb{T}^1 \subset \mathbb{T}^0$
- $\begin{array}{l} e^0 \in \mathbb{E}^0 ::= \underline{v}^0 \mid e_1^0 e_2^0 \mid \langle e^1 \rangle \mid \text{run } e^0 \\ v^0 \in \mathbb{V}^0 ::= x^0 \mid \lambda x^0. e^0 \mid \langle v^1 \rangle \\ e^1 \in \mathbb{E}^1 ::= \underline{v}^1 \mid \lambda x^1. e^1 \mid e_1^1 e_2^1 \mid \sim e^0 \\ v^1 \in \mathbb{V}^1 ::= x^1 \mid \lambda x^1. v^1 \mid v_1^1 v_2^1 \end{array}$ terms and values at level 0 and 1

We give an informal semantics of the language.

- A value v^1 corresponds to an object-level program, while a term e^1 may require some meta-level computation to get an object-level program.
- The type $\langle \tau^1 \rangle$ classifies code, i.e. meta-level values of the form $\langle v^1 \rangle$ representing an object-level program v^1 . Brackets $\langle e^1 \rangle$ and escape $\sim e^0$ allow to move between meta-level code and the corresponding object-level program, in particular $\sim \langle e^1 \rangle$ and e^1 evaluate to the same object-level program.
- The construct `run` e^0 first evaluates e^0 to code $\langle v^1 \rangle$, then evaluates the object-level program v^1 (provided it is a *complete* program, i.e. $\text{FV}(v^1) = \emptyset$).

Besides the stratification in two levels, there are the following syntactic differences between `MetaML`₂ and `MetaML` of [CMT04]:

- MetaML₂ values are explicitly marked in terms. This avoids re-evaluation and a general pitfall of monadic CBV translations, namely preservation of the operational semantics (as stated in Theorem 6) would fail.
- Cross-stage persistence, i.e. the ability to include meta-level values ($v^0: \tau^1$) into object-level programs, is excluded from MetaML₂. This simplifies some definitions and technical lemmas.
- In MetaML₂ code types are not annotated with environment classifiers.

A type system (without the environment classifiers of [TN03, CMT04]) is given by the following rules, where n ranges over levels (i.e. is either 0 or 1):

$$\begin{array}{c}
 \frac{\Gamma(x^n) = \tau^n}{\Gamma \vdash_n x^n: \tau^n} \quad \text{val} \quad \frac{\Gamma \vdash_n v^n: \tau^n}{\Gamma \vdash_n \underline{v}^n: \tau^n} \quad \text{brk}_v \quad \frac{\Gamma \vdash_1 v^1: \tau^1}{\Gamma \vdash_0 \langle v^1 \rangle: \langle \tau^1 \rangle} \\
 \lambda^n \quad \frac{\Gamma, x^n: \tau_1^n \vdash_n e^n: \tau_2^n}{\Gamma \vdash_n \lambda x^n. e^n: \tau_1^n \rightarrow \tau_2^n} \quad @^n \quad \frac{\Gamma \vdash_n e_1^n: \tau_1^n \rightarrow \tau_2^n \quad \Gamma \vdash_n e_2^n: \tau_1^n}{\Gamma \vdash_n e_1^n e_2^n: \tau_2^n} \\
 \lambda_v \quad \frac{\Gamma, x^1: \tau_1^1 \vdash_1 v^1: \tau_2^1}{\Gamma \vdash_1 \lambda x^1. v^1: \tau_1^1 \rightarrow \tau_2^1} \quad @_v \quad \frac{\Gamma \vdash_1 v_1^1: \tau_1^1 \rightarrow \tau_2^1 \quad \Gamma \vdash_1 v_2^1: \tau_1^1}{\Gamma \vdash_1 v_1^1 v_2^1: \tau_2^1} \\
 \text{run} \quad \frac{\Gamma \vdash_0 e^0: \langle \tau^1 \rangle}{\Gamma \vdash_0 \text{run } e^0: \tau^1} \quad \text{brk} \quad \frac{\Gamma \vdash_1 e^1: \tau^1}{\Gamma \vdash_0 \langle e^1 \rangle: \langle \tau^1 \rangle} \quad \text{esc} \quad \frac{\Gamma \vdash_0 e^0: \langle \tau^1 \rangle}{\Gamma \vdash_1 \sim e^0: \tau^1}
 \end{array}$$

The operational semantics of Table 4 consists of two relations $e^n \xrightarrow{n} v^n$, that evaluate terms to values. The operational semantics uses only the substitution $\llbracket x^0: v^0 \rrbracket$ and enjoys the following properties.

Proposition 1 (Operational properties).

- demote $\frac{\{x_i^1: \tau_i^1 \mid i \in m\} \vdash_1 v^1: \tau^1}{\{x_i^0: \tau_i^1 \mid i \in m\} \vdash_0 v^1 \downarrow: \tau^1}$
- SR $\frac{\Gamma \vdash_n e^n: \tau^n \quad e^n \xrightarrow{n} v^n}{\Gamma \vdash_n \underline{v}^n: \tau^n}$

6.2 Translation of MetaML₂ into MML_v^N

Table 5 defines (by induction on the syntax of MetaML₂) a translation $\llbracket e^n \rrbracket^\rho$, $\llbracket v^0 \rrbracket^\rho$ and $\llbracket v^1 \rrbracket_\theta^\rho$, where θ is a resolver and ρ is a (partial) mapping from variables of MetaML₂ to terms of MML_v^N such that

- a meta-level variable x^0 is mapped to a term e
- an object-level variable x^1 is mapped to a fragment $\mathbf{b}(r)e$

The parameters ρ and θ are convenient to state some properties (see Lemma 1), but for the definition of the translation it suffices to take $\theta = r$ and $\rho(x^0) = x$.

Some clauses in the definition of the translation deserve to be commented:

- terms of the form \underline{v}^n are always translated into terms of the form $\text{ret } e$, since values v^n do not require meta-level computation

Table 4. Big-Step Operational Semantics for MetaML₂

$\frac{v^n \xrightarrow{c^n} v^n}{e^1 \xrightarrow{c^1} v^1}$	$\frac{e_1^0 \xrightarrow{c^0} \lambda x^0.e^0 \quad e_2^0 \xrightarrow{c^0} v^0 \quad e^0[x^0:v^0] \xrightarrow{c^0} v_1^0}{e_1^0 e_2^0 \xrightarrow{c^0} v_1^0}$
$\frac{e^1 \xrightarrow{c^1} v^1}{\langle e^1 \rangle \xrightarrow{c^0} \langle v^1 \rangle}$	$\frac{e^0 \xrightarrow{c^0} \langle v^1 \rangle \quad v^1 \downarrow \xrightarrow{c^0} v^0}{\text{run } e^0 \xrightarrow{c^0} v^0} \quad \text{FV}(v^1) = \emptyset$
$\frac{e^1 \xrightarrow{c^1} v^1}{\lambda x^1.e^1 \xrightarrow{c^1} \lambda x^1.v^1}$	$\frac{e_1^1 \xrightarrow{c^1} v_1^1 \quad e_2^1 \xrightarrow{c^1} v_2^1}{e_1^1 e_2^1 \xrightarrow{c^1} v_1^1 v_2^1} \quad \frac{e^0 \xrightarrow{c^0} \langle v^1 \rangle}{\tilde{e}^0 \xrightarrow{c^1} v^1}$

where demotion $v^1 \downarrow$ is defined by induction on v^1

$$x^1 \downarrow = \underline{x^0} \quad (\lambda x^1.v^1) \downarrow = \underline{\lambda x^0.v^1 \downarrow} \quad (v_1^1 v_2^1) \downarrow = v_1^1 \downarrow v_2^1 \downarrow$$

- the translations of $\langle e^1 \rangle$ and e^1 are the same (and similarly for \tilde{e}^0 and e^0), because the bijection between meta-level code and object-level programs is collapsed to an equality
- the translation of $\langle v^1 \rangle$ is a fragment, which results into an object-level program after linking, thus the translation of v^1 depends on a resolver θ
- the translations of $e_1^1 e_2^1$ and $\lambda x^1.e^1$ are meta-level computations to generate code representing an application and abstraction in the object language
- the translation of values at level 1 (i.e. object-level programs) is like the monadic CBV translation of the λ -calculus, as the object language is CBV.

There are problems in extending the translation to types (thus we make no formal claim about preservation of typing). More precisely, the problem is to identify a signature to replace of \dots in the following inductive definition

$$\llbracket b \rrbracket = b \quad \llbracket \tau_1^n \rightarrow \tau_2^n \rrbracket = \llbracket \tau_1^n \rrbracket \rightarrow M[\llbracket \tau_2^n \rrbracket] \quad \llbracket \langle \tau^1 \rangle \rrbracket = [\dots | M[\llbracket \tau^1 \rrbracket]]$$

The translation preserves the operational semantics in the following sense:

Theorem 6. $e^0 \xrightarrow{c^0} v^0$ and $\text{FV}(e^0) = \emptyset$ imply $(\emptyset \llbracket e^0 \rrbracket, \square) \xrightarrow{*} (\mathcal{X} | \text{ret } \llbracket v^0 \rrbracket, \square)$ for some \mathcal{X} , where $\xrightarrow{*} \triangleq \longrightarrow \cup \dashv \longrightarrow$.

The result is a consequence of the following lemmas (stated without proof).

Lemma 1 (Properties of Translation).

1. If $e = \llbracket v^0 \rrbracket^\rho$, then $\llbracket _ [x^0:v^0] \rrbracket^\rho = \llbracket _ \rrbracket^{\rho, x^0:e}$ and $\llbracket _ [x^0:v^0] \rrbracket_\theta^\rho = \llbracket _ \rrbracket_\theta^{\rho, x^0:e}$
2. $\llbracket v^1 \rrbracket_\theta^\rho = \llbracket v^1 \downarrow \rrbracket_{\theta'}^{\rho'}$, if $\rho'(x^0) = e[r:\theta]$ when $\rho(x^1) = \mathbf{b}(r)e$ and $x^1 \in \text{FV}(v^1)$
3. $\llbracket v^1 \rrbracket_{\theta_1}^{\rho_1} \xrightarrow{*} \llbracket v^1 \rrbracket_{\theta_2}^{\rho_2}$, if $e_1[r:\theta_1] \xrightarrow{*} e_2[r:\theta_2]$ when $\rho_i(x^1) = \mathbf{b}(r)e_i$ and $x^1 \in \text{FV}(v^1)$.

Table 5. Translation of MetaML₂ terms and values

e^0	$\llbracket e^0 \rrbracket^\rho$
v^0	$\text{ret } \llbracket v^0 \rrbracket^\rho$
$e_1^0 e_2^0$	$\text{do } x_1 \leftarrow \llbracket e_1^0 \rrbracket^\rho; x_2 \leftarrow \llbracket e_2^0 \rrbracket^\rho; x_1 x_2$
$\langle e^1 \rangle$	$\llbracket e^1 \rrbracket^\rho$
$\text{run } e^0$	$\text{do } x \leftarrow \llbracket e^0 \rrbracket^\rho; x(?)$
v^0	$\llbracket v^0 \rrbracket^\rho$
x^0	e where $e = \rho(x^0)$
$\lambda x^0. e^0$	$\lambda x. \llbracket e^0 \rrbracket^{\rho, x^0 : x}$
$\langle v^1 \rangle$	$\text{b}(r) \llbracket v^1 \rrbracket_r^\rho$
e^1	$\llbracket e^1 \rrbracket^\rho$
v^1	$\text{ret } (\text{b}(r) \llbracket v^1 \rrbracket_r^\rho)$
$e_1^1 e_2^1$	$\text{do } x'_1 \leftarrow \llbracket e_1^1 \rrbracket^\rho; x'_2 \leftarrow \llbracket e_2^1 \rrbracket^\rho; \text{ret } (\text{b}(r) \text{do } x_1 \leftarrow x'_1(r); x_2 \leftarrow x'_2(r); x_1 x_2)$
$\lambda x^1. e^1$	$\nu X. \text{do } x' \leftarrow \llbracket e^1 \rrbracket^{\rho, x^1 : \text{b}(r)r.X}; \text{ret } (\text{b}(r) \text{ret } \lambda x. x' \langle r\{X : x\} \rangle)$
\tilde{e}^0	$\llbracket e^0 \rrbracket^\rho$
v^1	$\llbracket v^1 \rrbracket_\theta^\rho$
x^1	$\text{ret } e[r : \theta]$ where $\text{b}(r)e = \rho(x^1)$
$\lambda x^1. v^1$	$\text{ret } \lambda x. \llbracket v^1 \rrbracket_\theta^{\rho, x^1 : \text{b}(r)x}$
$v_1^1 v_2^1$	$\text{do } x_1 \leftarrow \llbracket v_1^1 \rrbracket_\theta^\rho; x_2 \leftarrow \llbracket v_2^1 \rrbracket_\theta^\rho; x_1 x_2$

Lemma 2 (Preservation of Evaluation). For any \mathcal{X} and E

- $e^0 \xrightarrow{c^0} v^0$ implies $(\mathcal{X} \llbracket e^0 \rrbracket^\rho, E) \xrightarrow{*} (\mathcal{X}' \text{ret } \llbracket v^0 \rrbracket^\rho, E)$ for some \mathcal{X}'
- $e^1 \xrightarrow{c^1} v^1$ implies $(\mathcal{X} \llbracket e^1 \rrbracket^\rho, E) \xrightarrow{*} (\mathcal{X}' \text{ret } \text{b}(r) \llbracket v^1 \rrbracket_r^\rho, E)$ for some \mathcal{X}'

provided $x^1 \in \text{FV}(e^n)$ implies $\rho(x^1) = \text{b}(r)r.X$ for some $X \in \mathcal{X}$.

We conclude with three examples of MetaML₂-terms. Each term evaluates to the same MetaML₂-value $v^0 \triangleq \langle \lambda x^1. x^1 \rangle$, i.e. the code representing the object-level identity function. However, the MML_v^N -translation of these terms reflect the different complexity of the evaluation to v^0 .

- The term $e_0^0 \triangleq \underline{v}^0$ evaluates immediately to v^0 .

The translation $\llbracket v^0 \rrbracket^\theta$ is given by $e \triangleq \text{b}(r) \text{ret } \lambda x. \text{ret } x$, where $\text{ret } \lambda x. \text{ret } x$ is the CBV monadic translation of $\lambda x. x$. The translation $\llbracket e_0^0 \rrbracket^\theta$ is simply $\text{ret } e$.

Therefore, Theorem 6 holds trivially for $e_0^0 \xrightarrow{c^0} v_0^0$.

- The term $e_1^0 \triangleq \langle \lambda x^1. \underline{x}^1 \rangle$ evaluates to v^0 , but the evaluation steps are more complex. This complexity is mirrored in the translation $\llbracket e_1^0 \rrbracket^\theta$ given by

$$e_1 \triangleq \nu X. \text{do } x' \leftarrow \text{ret } (\mathbf{b}(r)\text{ret } r.X); \\ \text{ret } (\mathbf{b}(r)\text{ret } \lambda x.x' \langle r\{X:x\} \rangle)$$

Theorem 6 for $e_1^0 \xrightarrow{0} v^0$ yields $(\emptyset | e_1, \square) \xrightarrow{+} (X | \text{ret } e, \square)$.

- The term $e_2^0 \triangleq \langle \lambda x^1. \sim((\lambda x^0.x^0)\langle x^1 \rangle) \rangle$ evaluates to v^0 , and requires evaluation within the body of the λx^1 -binder. The translation $\llbracket e_2^0 \rrbracket^\emptyset$ is given by

$$e_2 \triangleq \nu X. \text{do } x' \leftarrow \text{do } x_1 \leftarrow \text{ret } (\lambda x. \text{ret } x); \\ x_2 \leftarrow \text{ret } (\mathbf{b}(r)\text{ret } r.X); \\ x_1 x_2; \\ \text{ret } (\mathbf{b}(r)\text{ret } \lambda x.x' \langle r\{X:x\} \rangle)$$

Theorem 6 yields $(\emptyset | e_2, \square) \xrightarrow{+} (X | \text{ret } e, \square)$, but the steps needed to reach the final configuration are strictly more than in the previous case.

7 Relating MML_ν^N to CMS

In this section we recall CMS [AZ02], a purely functional calculus of mixin modules, and introduce ML_Σ^N , a variant of MML_ν^N . Then we define a translation of CMS in ML_Σ^N preserving CMS typing and simplification up to Ariola's equational axioms [AB02] for recursion. We summarize the main differences between MML_ν^N and CMS (for those already familiar with CMS).

- CMS has a fixed infinite set of names (but a program uses only finitely many of them) and no fresh name generation facility.
- CMS is a pure calculus, thus we can restrict to the fragment of MML_ν^N without computational types, called ML^N .
- In CMS recursion is bundled in mixin, and removing it results in a very inexpressive calculus. On the contrary, ML^N is an interesting calculus (comparable to the λ -calculus) even without recursion, and one can add recursion following standard approaches.

7.1 CMS

We recall the calculus of mixin modules CMS, and refer to [AZ99, AZ02] for further details. The syntax of CMS is abstracted over symbolic names $X \in \text{Name}$, and term variables $x \in \mathbf{X}$. For simplicity, we avoid to introduce core terms and types (in [AM04] the calculus is parametrized w.r.t. a core calculus).

- $\tau \in \text{CMST} ::= [\Sigma_1; \Sigma_2]$ types, where $\Sigma: \text{Name} \xrightarrow{\text{fin}} \text{CMST}$
- $E \in \text{CMSE} ::= x \mid [l; o; \rho] \mid E_1 + E_2 \mid E \setminus X \mid E!X \mid E.X$ terms, where $l: \mathbf{X} \xrightarrow{\text{fin}} \text{Name}$, $o: \text{Name} \xrightarrow{\text{fin}} \text{CMSE}$, $\rho: \mathbf{X} \xrightarrow{\text{fin}} \text{CMSE}$ and we implicitly require that $\text{dom}(l) \# \text{dom}(\rho)$ for well-formed of $[l; o; \rho]$.

Free variables are defined as follows (omitting trivial cases): $\text{FV}([l; o; \rho]) = (\text{FV}(o) \cup \text{FV}(\rho)) \setminus (\text{dom}(l) \cup \text{dom}(\rho))$. Thus one can freely rename the bound

variables in $\text{dom}(\iota) \cup \text{dom}(\rho)$, as done implicitly in the reduction rule (sum) below. We first give an informal overview of the calculus:

- The type $[\Sigma_1; \Sigma_2]$ specifies the names and types of the *deferred* (Σ_1) and *defined* (Σ_2) components of a mixin. The deferred components can be referred in the mixin, but are not defined, therefore they need to be resolved (see the freeze operation described below). The defined components corresponds to the exported definitions of the mixin.
- Term variables are used for local referencing of components, whereas names are needed for dealing with global access and linking of components. As in MML_ν^N , names are not terms.
- In a basic mixin $[\iota; \rho]$, ι specifies the deferred components. The mapping to names is needed for component resolution (see the freeze operation described below). The defined components (ρ) are associated with names, whereas local components (ρ) are introduced by variables and can be mutually recursive.
- The sum operation ($E_1 + E_2$) performs the union of the deferred components (in the sense that components with the same name are shared), and the disjoint union of the defined and local components of the two mixins.
- The *freeze* operation ($E!X$) binds the deferred component X to the expression of the defined component X in the same mixin; in this way a name can be resolved, and a deferred component becomes local. Cross-module recursion is obtained as a combination of the sum and the freeze operations.
- The *delete* operation ($E \setminus X$) is used for hiding defined components.
- Selection of a defined component ($E.X$) is only allowed for mixin with no deferred components.

Typing Rules. The typing judgment has form $\Gamma \vdash_{\text{CMS}} E: \tau$, where $\Gamma: \mathcal{X} \xrightarrow{\text{fin}} \text{CMST}$. The typing rules are given in Table 6, where two signatures Σ_1 and Σ_2 are *compatible* iff $\Sigma_1(X) = \Sigma_2(X)$ for all $X \in \text{dom}(\Sigma_1) \cap \text{dom}(\Sigma_2)$.

Simplification Rules. We define the relation $\xrightarrow{\text{CMS}}$ as the compatible closure of the simplification rules defined in Table 7.

7.2 ML_Σ^N

The syntax of ML_Σ^N is defined in two steps. First, we remove from MML_ν^N computational types (and consequently monadic operations, like $\nu X.e$). In the resulting calculus, called ML^N , the computation relation disappears (CMS is a pure calculus), and \mathcal{X} could be left implicit in the typing judgments $\mathcal{X}; \Pi; \Gamma \vdash e: \tau$, since the typing judgments of a derivation must use the same \mathcal{X} . Then we add records and mutual recursion:

- $\tau \in \mathbb{T}_{\mathcal{X}} + = \Sigma$ types, where $\Sigma \in \Sigma_{\mathcal{X}} \triangleq \mathcal{X} \xrightarrow{\text{fin}} \mathbb{T}_{\mathcal{X}}$ is a \mathcal{X} -signature
- $e \in \mathbb{E} + = o \mid e.X \mid e_1 + e_2 \mid e \setminus X \mid \text{let } \rho \text{ in } e$ terms, where
 - $o: \text{Name} \xrightarrow{\text{fin}} \mathbb{E}$ is a record $\{X_i: e_i \mid i \in m\}$ and
 - $\rho: \mathcal{X} \xrightarrow{\text{fin}} \mathbb{E}$ is a (recursive) binding $\{x_i: e_i \mid i \in m\}$.

Table 6. Type System for CMS

$$\begin{array}{c}
\text{var} \frac{}{\Gamma \vdash_{\text{CMS}} x: \tau} \quad \Gamma(x) = \tau \quad \text{delete} \frac{\Gamma \vdash_{\text{CMS}} e: [\Sigma_1; \Sigma_2]}{\Gamma \vdash_{\text{CMS}} e \setminus X: [\Sigma_1; \Sigma_2 \setminus X]} \\
\text{mixin} \frac{\begin{array}{l} \{\Gamma, \Sigma_1 \circ \iota, \Gamma' \vdash_{\text{CMS}} o(X): \Sigma_2(X) \mid X \in \text{dom}(o)\} \\ \{\Gamma, \Sigma_1 \circ \iota, \Gamma' \vdash_{\text{CMS}} \rho(x): \Gamma'(x) \mid x \in \text{dom}(\rho)\} \end{array}}{\Gamma \vdash_{\text{CMS}} [\iota; o; \rho]: [\Sigma_1; \Sigma_2]} \quad \begin{array}{l} \text{dom}(\Gamma') = \text{dom}(\rho) \\ \text{dom}(\Sigma_1) = \text{img}(\iota) \\ \text{dom}(\Sigma_2) = \text{dom}(o) \end{array} \\
\text{sum} \frac{\Gamma \vdash_{\text{CMS}} e_1: [\Sigma_1^1; \Sigma_2^1] \quad \Gamma \vdash_{\text{CMS}} e_2: [\Sigma_1^2; \Sigma_2^2]}{\Gamma \vdash_{\text{CMS}} e_1 + e_2: [\Sigma_1^1, \Sigma_1^2; \Sigma_2^1, \Sigma_2^2]} \quad \Sigma_1^1 \text{ compatible with } \Sigma_1^2 \\ \text{dom}(\Sigma_2^1) \# \text{dom}(\Sigma_2^2) \\
\text{freeze} \frac{\Gamma \vdash_{\text{CMS}} e: [\Sigma_1; \Sigma_2]}{\Gamma \vdash_{\text{CMS}} e!X: [\Sigma_1 \setminus X; \Sigma_2]} \quad \tau = \Sigma_1(X) = \Sigma_2(X) \\
\text{select} \frac{\Gamma \vdash_{\text{CMS}} e: [\emptyset; \Sigma]}{\Gamma \vdash_{\text{CMS}} e.X: \tau} \quad \tau = \Sigma(X)
\end{array}$$

Table 7. Simplification rules for CMS

$$\begin{array}{l}
\text{sum}) [\iota_1; o_1; \rho_1] + [\iota_2; o_2; \rho_2] \xrightarrow{\text{CMS}} [\iota_1, \iota_2; o_1, o_2; \rho_1, \rho_2] \quad \text{if } \text{dom}(o_1) \# \text{dom}(o_2) \\
\text{delete}) [\iota; o; \rho] \setminus X \xrightarrow{\text{CMS}} [\iota; o \setminus X; \rho] \\
\text{freeze}) [\iota, \{x: X\}; o, \{X: E\}; \rho]!X \xrightarrow{\text{CMS}} [\iota; o, \{X: E\}; \rho, \{x: E\}] \\
\text{select}) [\iota; o, \{X: E\}; \rho].X \xrightarrow{\text{CMS}} E[x: [\iota; X: \rho(x); \rho]].X \quad x \in \text{dom}(\rho)
\end{array}$$

Free variables are defined as follows (omitting trivial cases): $\text{FV}(\text{let } \rho \text{ in } e) = \text{FV}(e) \setminus \text{dom}(\rho)$.

The type $\Sigma \equiv \{X_i; \tau_i \mid i \in m\}$ classifies records of the form $\{X_i; e_i \mid i \in m\}$, i.e. with a fixed set of components. Notice that records should not be confused with resolvers. In particular, a fragment of type $[\Sigma] \tau$ can be linked with a resolver of any signature $\Sigma' \supseteq \Sigma$. The operations on records correspond to the CMS primitives for mixins: $e.X$ selects the component named X , $e_1 + e_2$ concatenates two records (provided their component names are disjoint), and $e \setminus X$ removes the component named X (if present). The let construct allows mutually recursive declarations, which are used to encode the local components of a CMS module. The order of record components and mutually recursive declarations are immaterial, therefore o and ρ are not sequences but functions (with finite domain).

Table 8 gives the typing rules for the new constructs. The properties of the type system in Section 4 extend in the obvious way to ML_{Σ}^N . We define

Table 8. Additional Typing Rules for ML_Σ^N

o	$\frac{\{\mathcal{X}; \Pi; \Gamma \vdash e_i: \tau_i \mid i \in m\}}{\mathcal{X}; \Pi; \Gamma \vdash \{X_i: e_i \mid i \in m\}: \{X_i: \tau_i \mid i \in m\}}$	select	$\frac{\Sigma(X) = \tau \quad \mathcal{X}; \Pi; \Gamma \vdash e: \Sigma}{\mathcal{X}; \Pi; \Gamma \vdash e.X: \tau}$
plus	$\frac{\mathcal{X}; \Pi; \Gamma \vdash e_1: \Sigma_1 \quad \mathcal{X}; \Pi; \Gamma \vdash e_2: \Sigma_2}{\mathcal{X}; \Pi; \Gamma \vdash e_1 + e_2: \Sigma_1, \Sigma_2}$	dom(Σ_1)#dom(Σ_2)	
	$\text{delete} \frac{\mathcal{X}; \Pi; \Gamma \vdash e: \Sigma}{\mathcal{X}; \Pi; \Gamma \vdash e \setminus X: \Sigma \setminus X}$		
rec	$\frac{\{\mathcal{X}; \Pi; \Gamma, \Gamma' \vdash \rho(x): \Gamma'(x) \mid x \in \text{dom}(\rho)\} \quad \mathcal{X}; \Pi; \Gamma, \Gamma' \vdash e: \tau}{\mathcal{X}; \Pi; \Gamma \vdash \text{let } \rho \text{ in } e: \tau}$	dom(Γ') = dom(ρ)	

simplification \longrightarrow for ML_Σ^N as the compatible closure of the simplification rules for MML_ν^N (see Section 4.2) and the following simplification rules for record operations and mutually recursive declarations:

- select**) $o.X \longrightarrow e$ if $e \equiv o(X)$
plus) $o_1 + o_2 \longrightarrow o_1, o_2$ if $\text{dom}(o_1) \# \text{dom}(o_2)$
delete) $o \setminus X \longrightarrow o_{\setminus X}$
unfolding) $\text{let } \rho \text{ in } e \longrightarrow e[x: \text{let } \rho \text{ in } \rho(x) \mid x \in \text{dom}(\rho)]$

Simplification for ML_Σ^N enjoys confluence (Theorem 1) and subject reduction (Theorem 2).

7.3 Translation of CMS into ML_Σ^N

The key idea of the translation consists in translating a mixin type $[\Sigma_1; \Sigma_2]$ in $[\Sigma'_1 \mid \Sigma'_2]$, in this way we obtain a compositional translation of CMS terms. In contrast, a translation based on functional types, where $[\Sigma_1; \Sigma_2]$ is translated in $\Sigma'_1 \rightarrow \Sigma'_2$, is not compositional (the problem is in the translation of $e_1 + e_2$, which must be driven by the type of e_1 and e_2).

Table 9 gives the translation of CMS in ML_Σ^N . The translation can be easily extended to core terms (see [AM04]).

In the translation of a basic mixin $[\iota; o; \rho]$ the deferred variables x ($x \in \text{dom}(\iota)$) are replaced with the resolution $r.X$ of the corresponding name $X = \iota(x)$, whereas the local variables x ($x \in \text{dom}(\rho)$) are bound by the let construct for mutually recursive declarations. (A similar translation would not work in ν^\square , because of the limitations in typing discussed in Section 8).

The translation of selection $E.X$ uses the empty resolver $?$, since in CMS selection is allowed only for mixins without deferred components.

The freeze operator $E!X$ resolves a deferred component X with the corresponding output component. This resolution may introduce a recursive definition, since the output component X could be defined in terms of the corre-

Table 9. Translation of CMS in ML_{Σ}^N

CMS typing	ML_{Σ}^N typing
$\Gamma \vdash_{\text{CMS}} E: \tau$	$\mathcal{X}; \emptyset; \Gamma' \vdash E': \tau'$
CMS type	ML_{Σ}^N type
$[\Sigma_1; \Sigma_2]$	$[\Sigma'_1 \Sigma'_2]$
CMS term	ML_{Σ}^N term
x	x
$[i; o; \rho]$	$\mathbf{b}(r)(\text{let } \rho' \text{ in } o')[x: r.X \mid i(x) = X]$
$E_1 + E_2$	$\mathbf{b}(r)E'_1\langle r \rangle + E'_2\langle r \rangle$
$E \setminus X$	$\mathbf{b}(r)E'\langle r \rangle \setminus X$
$E.X$	$E'\langle ? \rangle.X$
$E!X$	$\mathbf{b}(r)\text{let } \{x_1: x_2.X, x_2: E'\langle r \{X: x_1\}\rangle\} \text{ in } x_2$

the translations of Γ , Σ , o and ρ are defined pointwise.

sponding deferred component. Therefore, the translation defines the record x_2 by resolving the name X with the X component of the record x_2 itself.

The typing preservation property of the translation can be proved

Theorem 7 (Typing preservation). *If $\Gamma \vdash_{\text{CMS}} E: \tau$, then $\mathcal{X}; \emptyset; \Gamma' \vdash E': \tau'$, where \mathcal{X} includes all names occurring in the derivation of $\Gamma \vdash_{\text{CMS}} E: \tau$.*

The translation preserves also the semantics of CMS, but this can be proved only up to some equational axioms for mutually recursive declarations

$$\begin{aligned}
\mathbf{C}[\text{let } \rho \text{ in } e] &= \text{let } \rho \text{ in } \mathbf{C}[e] && \text{(lift)} \\
\text{let } \rho_1 \text{ in let } \rho_2 \text{ in } e &= \text{let } \rho_1, \rho_2 \text{ in } e && \text{(ext-merge)} \\
\text{let } \rho_1, x: (\text{let } \rho_2 \text{ in } e_2) \text{ in } e_1 &= \text{let } \rho_1, \rho_2, x: e_2 \text{ in } e_1 && \text{(int-merge)} \\
\text{let } \rho, x: e_1 \text{ in } e_2 &= \text{let } \rho[x: e_1] \text{ in } e_2[x: e_1] \text{ if } x \notin \text{FV}(e_1) && \text{(sub)}
\end{aligned}$$

The (lift) axiom corresponds to Ariola's lift axioms, in principle it can be instantiated with any ML_{Σ}^N context $\mathbf{C}[\]$, but for proving Theorem 8 it suffices to consider the following contexts: $\boxed{\mathbf{C}[\] := \square + e \mid e + \square \mid \square \setminus X \mid \square.X}$.

The (ext-merge) and (int-merge) axioms are Ariola's merge axioms, whereas (sub) is derivable from Ariola's axioms.

Let R denotes the set of the three axioms above, and S denotes the set of equational axioms corresponding to the simplification rules for ML_{Σ}^N ; then the translation is proved to preserve the CMS simplification up to $=_{S \cup R}$ (i.e. the congruence induced by the axioms in $S \cup R$).

Theorem 8 (Semantics preservation). *If $E_1 \xrightarrow{\text{CMS}} E_2$, then $E'_1 =_{S \cup R} E'_2$.*

The translation of the non-recursive subset of CMS (i.e. no local declarations ρ and no freeze $E!X$) is a lot simpler, moreover its simplifications are mapped to plain ML_{Σ}^N simplifications.

We conclude this section with some examples of CMS reductions and their translations in ML_{Σ}^N .

Example 4. Consider the **plus** reduction $E_1 \xrightarrow{\text{CMS}} E_2$, where

$$E_1 \equiv [\{x:A\}; \{R:x\}; \emptyset] + [\emptyset; \{A:y\}; \emptyset],$$

$$E_2 \equiv [\{x:A\}; \{A:y, R:x\}; \emptyset].$$

The translations of E_1 and E_2 are given by

$$E'_1 \equiv \mathbf{b}(r)(\mathbf{b}(r_1)(\text{let } \emptyset \text{ in } \{R:r_1.A\})\langle r \rangle + (\mathbf{b}(r_2)\text{let } \emptyset \text{ in } \{A:y\})\langle r \rangle)$$

$$E'_2 \equiv \mathbf{b}(r)\text{let } \emptyset \text{ in } \{A:y, R:r.A\}.$$

By repeatedly applying simplifications and equational axioms in R we get

$$E'_1 \xrightarrow{*} (\text{by link})$$

$$\mathbf{b}(r)(\text{let } \emptyset \text{ in } \{R:r.A\}) + \text{let } \emptyset \text{ in } \{A:y\} =_R (\text{by lift})$$

$$\mathbf{b}(r)\text{let } \emptyset \text{ in let } \emptyset \text{ in } \{R:r.A\} + \{A:y\} =_R (\text{by ext-merge})$$

$$\mathbf{b}(r)\text{let } \emptyset \text{ in } \{A:y, R:r.A\} \equiv E'_2.$$

Example 5. Consider the **freeze** reduction $E_3 \xrightarrow{\text{CMS}} E_4$, where

$$E_3 \equiv E_2!A \text{ with } E_2 \text{ defined as in the previous example,}$$

$$E_4 \equiv [\emptyset; \{A:y, R:x\}; \{x:y\}].$$

The translations of E_3 and E_4 are given by

$$E'_3 \equiv \mathbf{b}(r)\text{let } \{x_1:x_2.A, x_2:E'_2\langle r\{A:x_1\}\rangle \text{ in } x_2$$

$$E'_4 \equiv \mathbf{b}(r)\text{let } \{x:y\} \text{ in } \{A:y, R:x\}.$$

By repeatedly applying simplifications and equational axioms in R we get

$$E'_3 \longrightarrow (\text{by link})$$

$$\mathbf{b}(r)\text{let } \{x_1:x_2.A, x_2:\text{let } \emptyset \text{ in } \{A:y, R:(r\{A:x_1\}).A\}\} \text{ in } x_2 \longrightarrow (\text{by resolve})$$

$$\mathbf{b}(r)\text{let } \{x_1:x_2.A, x_2:\text{let } \emptyset \text{ in } \{A:y, R:x_1\}\} \text{ in } x_2 =_R (\text{by sub})$$

$$\mathbf{b}(r)\text{let } \{x_1:(\text{let } \emptyset \text{ in } \{A:y, R:x_1\}).A\} \text{ in let } \emptyset \text{ in } \{A:y, R:x_1\} =_R (\text{by ext-merge})$$

$$\mathbf{b}(r)\text{let } \{x_1:(\text{let } \emptyset \text{ in } \{A:y, R:x_1\}).A\} \text{ in } \{A:y, R:x_1\} =_R (\text{by lift})$$

$$\mathbf{b}(r)\text{let } \{x_1:\text{let } \emptyset \text{ in } \{A:y, R:x_1\}.A\} \text{ in } \{A:y, R:x_1\} =_R (\text{by int-merge})$$

$$\mathbf{b}(r)\text{let } \{x_1:\{A:y, R:x_1\}.A\} \text{ in } \{A:y, R:x_1\} \longrightarrow (\text{by resolve})$$

$$\mathbf{b}(r)\text{let } \{x_1:y\} \text{ in } \{A:y, R:x_1\} \text{ which is } \alpha\text{-equivalent to } E'_4.$$

Example 6. Consider the **select** reduction $E_5 \xrightarrow{\text{CMS}} E_6$, where

$$E_5 \equiv E_4.R \text{ with } E_4 \text{ defined as in the previous example,}$$

$$E_6 \equiv [\emptyset; \{X:y\}; \{x:y\}].X.$$

The translations of E_5 and E_6 are given by

$$E'_5 \equiv (\mathbf{b}(r)\text{let } \{x:y\} \text{ in } \{A:y, R:x\})\langle ? \rangle.R,$$

$$E'_6 \equiv (\mathbf{b}(r)\text{let } \{x:y\} \text{ in } \{X:y\})\langle ? \rangle.X.$$

By repeatedly applying simplifications we get

$$E'_5 \longrightarrow (\text{by link})$$

$$(\text{let } \{x:y\} \text{ in } \{A:y, R:x\}).R \longrightarrow (\text{by unfolding})$$

$$\{A:y, R:\text{let } \{x:y\} \text{ in } y\}.R \longrightarrow (\text{by resolve})$$

$$\text{let } \{x:y\} \text{ in } y \longrightarrow (\text{by unfolding})$$

$$y.$$

Analogously, $E'_6 \xrightarrow{*} y$, therefore $E'_5 =_S E'_6$. Unlike the other examples, there is no way to get from E'_5 to E'_6 by applying simplifications and equations axioms in R . This explains the use of $=_{S \cup R}$ in stating Theorem 8.

8 Conclusions and Related Work

This section compares MML_ν^N with the CBV calculi FreshML and ν^\square . First, we recall briefly the main features of these two calculi, then we make a critical assessment based on a comparison with MML_ν^N .

- FreshML of [SPG03]² is an extension of ML, based on a solid mathematical theory [GP99], that provides a convenient support for meta-programming. Namely, in FreshML abstract representations (i.e. modulo α -conversion) of object-level syntax co-exist with pattern matching facilities (similar to those usable on concrete parse trees) to analyse these representations.
- ν^\square of [Nan02, NPar] is a refinement of λ^\square [DP01], which provides better support for symbolic manipulation by exploiting some features of FreshML (the calculus presented in [NPar] does not support program analysis). The stated aim is to combine safely the best features of λ^\square (the ability to execute closed code) and λ° [Dav96] (the ability to manipulate open code). MetaML has similar aims, but adopts the opposite strategy, i.e. it starts from λ° , nor does it build on top of FreshML (names are not part of the MetaML syntax).

If we ignore the different styles for describing the operational semantics of FreshML (a CBV evaluation relation), ν^\square (a CBV small-step reduction relation) and MML_ν^N (a simplification and a computation relation), the key differences are:

- FreshML supports program transformation, in particular the analysis of object-level programs represented as values of an inductive datatype involving *abstraction* types $\langle \text{name} \rangle \tau$.
- ν^\square of [NPar] supports only program generation, (object-level) programs e of type τ with unresolved names in \mathcal{X} are represented by values of type $\square_{\mathcal{X}}\tau$. The typing rules for $\square_{\mathcal{X}}\tau$ are fairly restrictive, because \mathcal{X} has to include **all** unresolved names in e (and e should not contain free variables x).
- MML_ν^N supports only program generation, (object-level) programs e of type τ abstracted w.r.t. a name resolver r of signature Σ are represented by terms of type $[\Sigma|\tau]$. The typing rules for $[\Sigma|\tau]$ are similar to those for a functional type $\Sigma \rightarrow \tau$, in particular e may use other name resolvers besides r .

MML_ν^N *versus* FreshML. Typing judgments of FreshML have the standard format $\Gamma \vdash e : \tau$, because names do not occur in types (and in particular there are no resolvers and no signatures for typing resolvers).

In FreshML names are terms (and there is a type *name* of names), so generation of a fresh name is denoted by $\nu x.e$, where x is a term variable which gets bound to the fresh name, and e is the term where the fresh name can be used. In MML_ν^N names occur both in types and in terms, and using x in place of a name X would entail a type system with dependent types (which would be problematic), thus we must use a different binder $\nu X.e$ for names. FreshML, unlike MML_ν^N , supports the manipulation of object-level syntax modulo α -conversion. This is possible because FreshML has:

² There is another version of FreshML [PG00] with a more elaborate type system, which is able to mask the computational effects due to generation of fresh names.

- an equality type **name** of names
- abstraction types $\langle \text{name} \rangle \tau$ classifying equivalence classes of pairs (X, e) modulo renaming of X with names fresh for e (of type τ), terms $\langle e_1 \rangle e_2$ to form name abstractions, and patterns $\langle x \rangle p$ to deconstructed them.
- a name swapping operation, which is crucial (in combination with name generation) to define the operational semantics of name abstraction matching.

It should be possible to extend MML_ν^N with these features of FreshML. However, one should keep a clear distinction between the types $\langle \text{name} \rangle \tau$ and $[\Sigma] \tau$. The first type classifies representations of object-level syntax modulo α -conversion, while the second classifies fragments modulo simplification, thus it cannot support program analysis.

MML_ν^N versus ν^\square . Typing judgments of ν^\square take the form $\Sigma; \Delta; \Gamma \vdash e: \tau[\mathcal{X}]$, where $\mathcal{X} \subseteq \text{dom}(\Sigma)$ includes the names occurring *free* in e , and Δ has declarations of the form $u_i: \tau_i[\mathcal{X}_i]$ with $\mathcal{X}_i \subseteq \text{dom}(\Sigma)$.

In ν^\square the type of a name X is fixed at name generation time. This is a bad name space management policy, which goes against common practice in programming language design (e.g. of modules systems). MML_ν^N follows the approach of mainstream module languages, where different modules can assign to the same name different types (and values). Therefore, programming in ν^\square forces an overuse of name generation, because the language restricts name reuse.

In ν^\square terms includes names, so our $\theta.X$ is replaced by X , in other words there is a *default resolver* which is left implicit. Linking $u(\theta)$ uses a function $\Theta \equiv \langle X_i \rightarrow e_i \mid i \in m \rangle$ to modify the default resolver. The typing judgments for explicit substitutions Θ take the form $\Sigma; \Delta; \Gamma \vdash \Theta: \mathcal{X}[\mathcal{X}']$, where \mathcal{X}' includes the names *used* by the modified resolver to resolve the names in \mathcal{X} , e.g. $\mathcal{X} \subseteq \mathcal{X}'$ when Θ is empty. The following explicit substitution principle is admissible

$$\frac{\Sigma; \Delta; \Gamma \vdash \Theta: \mathcal{X}[\mathcal{X}'] \quad \Sigma; \Delta; \Gamma \vdash e: \tau[\mathcal{X}]}{\Sigma; \Delta; \Gamma \vdash e[\Theta]: \tau[\mathcal{X}']}$$

Our type $[\Sigma] \tau$ corresponds to $\square_{\mathcal{X}} \tau$ with $\mathcal{X} = \text{dom}(\Sigma)$. Typing rules for $\square_{\mathcal{X}} \tau$ are related to those for necessity of S4 modal logic, e.g. $\square_{\mathcal{X}} \tau$ introduction is

$$\frac{\Sigma; \Delta; \emptyset \vdash e: \tau[\mathcal{X}]}{\Sigma; \Delta; \Gamma \vdash \text{box } e: \square_{\mathcal{X}} \tau[\mathcal{X}']}$$

This rule is very restrictive: it forbids having free term variables x in e , and acts like an *implicit binder* for the free names X of e (i.e. it binds the default resolver for e). Without these restrictions substitution would be unsound in the type system of ν^\square . Such restrictions have no reason to exist in MML_ν^N , because we allow multiple name resolvers, and fragments $\mathbf{b}(r)e$ are formed by abstracting over one name resolver. Furthermore, making name resolvers explicit, avoid the need to introduce *non-standard* forms of substitution.

The observations above are formalized by a CBV translation \ulcorner of ν^\square -terms³ into MML_ν^N , where the resolver variable r corresponds to the default resolver, which is implicit in ν^\square .

$e \in \nu^\square$	$e' \in \text{MML}_\nu^N$	$e \in \nu^\square$	$e' \in \text{MML}_\nu^N$
x	$\text{ret } x$	X	$r.X$
$\lambda x: \tau. e$	$\text{ret } (\lambda x. e')$	$u\langle X_i \rightarrow e_i \rangle$	$u\langle r\{X_i: e'_i\} \rangle$
$e_1 \ e_2$	$\text{do } x_1 \leftarrow e'_1; x_2 \leftarrow e'_2; x_1 x_2$	$\text{box } e$	$\text{ret } (\text{b}(r)e')$
$\nu X: \tau. e$	$\nu X. e'$	$\text{letbox } u = e_1 \text{ in } e_2$	$\text{do } u \leftarrow e'_1; e'_2$

We do not define the translation on types and assignments, since in ν^\square the definition of well-formed signatures $\Sigma \vdash$ and types $\Sigma \vdash \tau$ is quite complex.

In conclusion, the key novelty of MML_ν^N is to make name resolvers explicit and to allow a multiplicity of them, as a consequence we gain in simplicity and expressivity. Moreover, by building on top of a fairly simple form of extensible records, we are better placed to exploit existing programming language implementations (like O'Caml).

References

- [AB02] Z. M. Ariola and S. Blom. Skew confluence and the lambda calculus with letrec. *Annals of pure and applied logic*, 117(1-3):95–178, 2002.
- [AM04] D. Ancona and E. Moggi. A fresh calculus for names management. In Karsai and Visser [KV04].
- [AZ99] Davide Ancona and Elena Zucca. A primitive calculus for module systems. In *Proc. Int'l Conf. Principles & Practice Declarative Programming*, volume 1702 of *Lecture Notes in Computer Science*, pages 62–79. Springer, 1999.
- [AZ02] D. Ancona and E. Zucca. A calculus of module systems. *J. Funct. Programming*, 12(2):91–132, March 2002. Extended version of [AZ99].
- [BCT02] D. Batory, C. Consel, and W. Taha, editors. *Generative Programming and Component Engineering*, volume 2487 of *Lecture Notes in Computer Science*. Springer, October 2002.
- [Car97] Luca Cardelli. Program fragments, linking, and modularization. In *Conf. Rec. POPL '97: 24th ACM Symp. Princ. of Prog. Langs.*, pages 266–277, 1997.
- [CE00] Krzysztof Czarnecki and Ulrich Eisenecker. *Generative Programming: Methods, Tools, and Applications*. Addison-Wesley, 2000.
- [CM94] L. Cardelli and J. C. Mitchell. Operations on records. In C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, pages 295–350. The MIT Press, Cambridge, MA, 1994.
- [CMS03] C. Calcagno, E. Moggi, and T. Sheard. Closed types for a safe imperative MetaML. *J. Funct. Programming*, 13(3):545–571, 2003.

³ In [NPar] the operational semantics (and the typing) of $\nu X.e$ differs from that adopted by (us and) FreshML. To avoid unnecessary complications, we work as if ν^\square is FreshML compliant.

- [CMT04] C. Calcagno, E. Moggi, and W. Taha. ML-like inference for classifiers. In *Programming Languages & Systems, 13th European Symp. Programming*, volume 2986 of *Lecture Notes in Computer Science*. Springer, 2004.
- [CTHL03] Cristiano Calcagno, Walid Taha, Liwen Huang, and Xavier Leroy. Implementing multi-stage languages using ASTs, gensym, and reflection. In Krzysztof Czarnecki, Frank Pfenning, and Yannis Smaragdakis, editors, *Generative Programming and Component Engineering (GPCE)*, Lecture Notes in Computer Science. Springer-Verlag, 2003.
- [Dav96] R. Davies. A temporal-logic approach to binding-time analysis. In *the Symposium on Logic in Computer Science (LICS '96)*, pages 184–195, New Brunswick, 1996. IEEE Computer Society Press.
- [DP01] Rowan Davies and Frank Pfenning. A modal analysis of staged computation. *Journal of the ACM*, 48(3):555–604, 2001.
- [ESOP00] *Programming Languages & Systems, 9th European Symp. Programming*, volume 1782 of *Lecture Notes in Computer Science*. Springer, 2000.
- [GP99] Murdoch J. Gabbay and Andrew M. Pitts. A new approach to abstract syntax involving binders. In *Proc. 14th Ann. IEEE Symp. Logic in Comput. Sci.*, pages 214–224, July 1999.
- [GS04] J. Greenfield and K. Short. *Software Factories: Assembling Applications with Patterns, Models, Frameworks and Tools*. Wiley Publishing Inc, 2004.
- [HW03] Christian Haack and J. B. Wells. Type error slicing in implicitly typed, higher-order languages. In *Programming Languages & Systems, 12th European Symp. Programming*, volume 2618 of *Lecture Notes in Computer Science*, pages 284–301. Springer, 2003. Superseded by [HW04].
- [HW04] Christian Haack and J. B. Wells. Type error slicing in implicitly typed, higher-order languages. *Sci. Comput. Programming*, 50:189–224, 2004. Supersedes [HW03].
- [KV04] G. Karsai and E. Visser, editors. *Generative Programming and Component Engineering*, volume 3286 of *Lecture Notes in Computer Science*. Springer, October 2004.
- [LBCO04] Christian Lengauer, Don Batory, Charles Consel, and Martin Odersky, editors. *Domain-Specific Program Generation*. Number 3016 in Lecture Notes in Computer Science. Springer-Verlag, 2004.
- [Met01] MetaOCaml: A compiled, type-safe multi-stage programming language. Available online from <http://www.cs.rice.edu/~{taha}/MetaOCaml/>, 2001.
- [MF03] E. Moggi and S. Fagorzi. A monadic multi-stage metalanguage. In *Proc. FoSSaCS '03*, volume 2620 of *Lecture Notes in Computer Science*. Springer, 2003.
- [MT00] Elena Machkasova and Franklyn A. Turbak. A calculus for link-time compilation. In ESOP '00 [ESOP00], pages 260–274.
- [Nan02] Aleksandar Nanevski. Meta-programming with names and necessity. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP-02)*, ACM SIGPLAN notices, New York, October 2002. ACM Press.
- [NPar] A. Nanevski and F. Pfenning. Staged computations with names and necessity. *J. Funct. Programming*, to appear.

- [PG00] Andrew M. Pitts and Murdoch J. Gabbay. A metalanguage for programming with bound names modulo renaming. In R. Backhouse and J. N. Oliveira, editors, *Proc. Mathematics of Program Construction, 5th Int'l Conf. (MPC 2000)*, volume 1837 of *Lecture Notes in Computer Science*, pages 230–255, Ponte de Lima, Portugal, July 2000. Springer.
- [PS03] F. Pfenning and Y. Smaragdakis, editors. *Generative Programming and Component Engineering*, volume 2830 of *Lecture Notes in Computer Science*. Springer, September 2003.
- [She01] T. Sheard. Accomplishments and research challenges in meta-programming. In W. Taha, editor, *Proc. of the Int. Work. on Semantics, Applications, and Implementations of Program Generation (SAIG)*, volume 2196 of *LNCS*, pages 2–46. Springer-Verlag, 2001.
- [SPG03] Mark R. Shinwell, Andrew M. Pitts, and Murdoch J. Gabbay. FreshML: Programming with binders made simple. In *Proc. 8th Int'l Conf. Functional Programming*. ACM Press, 2003.
- [Szy02] Clemens Szyperski. *Component Software: Beyond Object-Oriented Programming, 2nd Edition*. Addison Wesley, 2002.
- [Tah99] W. Taha. *Multi-Stage Programming: Its Theory and Applications*. PhD thesis, Oregon Graduate Inst. of Science and Technology, 1999. Available from <ftp://cse.ogi.edu/pub/tech-reports/README.html>.
- [TN03] Walid Taha and Michael Florentin Nielsen. Environment classifiers. In *The Symposium on Principles of Programming Languages (POPL '03)*, New Orleans, 2003.
- [TS97] W. Taha and T. Sheard. Multi-stage programming with explicit annotations. In *Proceedings of the Symposium on Partial Evaluation and Semantic-Based Program Manipulation (PEPM)*, pages 203–217, Amsterdam, 1997. ACM Press.
- [WV99] J. B. Wells and René Vestergaard. Confluent equational reasoning for linking with first-class primitive modules (long version). A short version is [WV00]. Full paper, 3 appendices of proofs, August 1999.
- [WV00] J. B. Wells and René Vestergaard. Equational reasoning for linking with first-class primitive modules. In *ESOP '00 [ESOP00]*, pages 412–428. A long version is [WV99].