Tilburg University

Metacasanova

di Giacomo, Francesco; Abbadi, Mohamed; Cortesi, Agostino; Spronck, Pieter; Maggiore, G.

Published in:

Proceedings of the 10th ACM SIGPLAN International Conference on Software Language Engineering

Publication date:

2017

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

di Giacomo, F., Abbadi, M., Cortesi, A., Spronck, P., & Maggiore, G. (2017). Metacasanova: an optimized meta-compiler for Domain-Specific Languages. In *Proceedings of the 10th ACM SIGPLAN International Conference on Software Language Engineering* http://delivery.acm.org/10.1145/3140000/3136015/sle17-sle17main2.pdf?ip=137.56.129.121&id=3136015&acc=OPENTOC&key=0C390721DC3021FF%2E8E8A7FC83EB1C6A0%2E4D4702B0C3E38B35%2EC1E31BC46E58D5B8&__acm__=1519031195_2a1bb0fc9a6b2bebe68f57b46885317d

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Metacasanova: An Optimized Meta-compiler for Domain-Specific Languages

Francesco Di Giacomo
Università Ca' Foscari
Venice, Italy
francesco.digiacom@unive.it

Mohamed Abbadi
Hogeschool Rotterdam
Rotterdam, Netherlands
abbam@hr.nl

Agostino Cortesi
Università Ca' Foscari
Venice, Italy
cortesi@unive.it

Pieter Spronck
Tilburg University
Tilburg, Netherlands
p.spronck@tilburguniversity.edu

Giuseppe Maggiore
Hogeschool Rotterdam
Rotterdam, Netherlands
giuseppemag@gmail.com

Abstract

Domain-Specific Languages (DSL's) offer language-level abstractions that General-Purpose Languages do not offer, thus speeding up the implementation of the solution of problems within a specific domain. Developers have the choice of developing a DSL by building an interpreter/compiler for it, which is a hard and time-consuming task, or embedding it in a host language, thus speeding up the development process but losing several advantages that having a dedicated compiler might bring. In this work we present a meta-compiler called *Metacasanova*, whose meta-language is based on operational semantics. Then, we propose a language extension with functors and modules that allows to embed the type system of a language definition inside the meta-type system of *Metacasanova* and improves the performance of manipulating data structures at run-time. Our results show that *Metacasanova* dramatically reduces the code lines required to develop a compiler, and that the running time of the Meta-program is improved by embedding the host language type system in the meta-type system with the use of functors in the meta-language.

CCS Concepts • Software and its engineering → Translator writing systems and compiler generators;

Keywords meta-compiler, optimization, operational semantics

ACM Reference Format:

Francesco Di Giacomo, Mohamed Abbadi, Agostino Cortesi, Pieter Spronck, and Giuseppe Maggiore. 2017. Metacasanova: An Optimized Meta-compiler for Domain-Specific Languages. In *Proceedings of 10th ACM SIGPLAN International Conference on Software Language Engineering (SLE'17)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3136014.3136015>

1 Introduction

Domain-Specific Languages (DSL's) are becoming more and more relevant in software engineering thanks to their ability to provide abstractions at language level to target specific problem domains [20, 21]. Notable examples of the use of DSL's are (i) game development (UnrealScript, JASS, Statusquo, NWScript), (ii) Database programming and design (SQL, LINQ), and (iii) numerical analysis and engineering (MATLAB, Octave). A notable amount of work on DSL's is made in the field of game development, as games are complex and they require, for performance reasons, to implement abstractions that exhibit the behaviours of threads. Indeed the overhead of threads is too big to be used in games to update every single entity in it, as their number can approach to the order of thousands (just think about a shooter game where the player can shoot with an automatic rifle at a rate of 30 rounds per second).

Two main alternatives have been proposed for the development of DSL's: the (i) *Embedding* technique, and (ii) the *Interpretation/Compilation* technique [16].

The former approach extends an existing programming language with the additional abstractions of the DSL. This is the case, for instance, of NWScript, a scripting language for the Neverwinter Nights game extending the C language, and LINQ, which offers SQL-like abstractions extending C#. This technique has the advantage that the infrastructure of the host language can be widely re-used, thus reducing the development effort. Moreover, people who are expert on the host language can become proficient with the DSL extension in a short time. The disadvantages are that the syntax is likely to be far from that of the DSL formal syntax, since

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SLE'17, October 23–24, 2017, Vancouver, Canada

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5525-4/17/10...\$15.00

<https://doi.org/10.1145/3136014.3136015>

General-Purpose Languages generally do not offer syntax extensions, and Domain-specific optimizations are difficult to achieve [13, 19].

The latter approach requires to develop an interpreter or compiler for the language. This is the case, for instance, of UnrealScript, JASS, and SQL. This approach has the advantages of providing a syntax close to the formal definition of the Domain-specific language, good error reporting, and domain-specific optimization through code analysis. However, designing and implementing a compiler for a DSL is a hard and time-consuming task, since a compiler is a complex piece of software made of different modules that perform several translation steps [3]. For this reason this option is not always considered feasible.

The translation steps performed by a compiler are not part of the creative aspect of designing the language [4, 8]. Therefore, they can be automated. The most common automated part is the Lexing/Parsing phase with Parser generators such as Yacc. A further effort in fully automating the development of a compiler has been done by employing *Meta-compilers*, that are computer programs that take as input the definition of a language (usually defined in a *meta-language*), a program written in that language and output executable code for the program. Meta-compilers usually automate not only the parsing phase, but also the type checking and the semantics implementation.

In this paper we present *Metacasanova*, a meta-compiler whose meta-language is based on operational semantics [10, 12], which was a project whose goal was easing the development and extension of the DSL for games *Casanova* [1, 2, 11], and a possible extension that aims to improve the performance of languages built in Metacasanova. In Section 2 we further discuss how developing a compiler leads to repetitive steps that could be automated and we formulate the problem statement of this paper; in Section 3 we explain how the meta-language of Metacasanova is defined and what its semantics are; in Section 4 we explain how the meta-compilation process is implemented in Metacasanova and how the target code is generated; in Section 5 we propose a further language abstraction for Metacasanova in order to improve the performance of the generated code; in Section 6 we evaluate the performance of the code generated by Metacasanova after re-implementing Casanova [2], a DSL for game development, and a subset of the C language. We also evaluate the performance gain given by the presented code optimization.

2 Repetitive Steps in Compilers Development

In Section 1 we briefly stated that the process of developing a compiler includes several steps that are repetitive, i.e. their behaviour is always the same regardless of the language for which the compiler is built. In this section we show in which

way this process is repetitive and what the common pattern is.

2.1 Type Checking

Type systems are generally expressed in the form of logical rules [7], made of a set of premises, that must be verified in order to assign to the language construct the type defined in the conclusion. For example the following rule defines the typing of an if-then-else statement in a functional programming language:¹

$$\frac{\Gamma \vdash c : \text{bool} \quad \Gamma \vdash t : \tau \quad \Gamma \vdash e : \tau}{\Gamma \vdash \text{if } c \text{ then } t \text{ else } e : \tau}$$

In this rule Γ is the environment. The type rule first evaluates the premises, which means that if the condition of the if-then-else has type `bool` and both then and else blocks have the same type, then the whole if-then-else has the type of either blocks.

Typing a construct of the language requires to evaluate its corresponding typing rule. In order to do so, the behaviour of each typing rule must be implemented in the host language in which the compiler is defined. Independently of the chosen language, the behaviour will always be the following : (i) evaluate a premise, (ii) if the evaluation of the premise fails then the construct fails the type check and an error is returned, (iii) repeat step 1 and 2 until all the premises have been evaluated, and (iv) assign the type to the construct that is defined in the rule conclusion.

2.2 Semantics

Semantics define how the language abstractions behave and can be expressed in different ways, for example with a term-rewriting system [15] or with the operational semantics [10]. For the scope of this work, we choose to rely on the operational semantics. The definition of the operational semantics of a language abstraction is, again, in the form of a logical rule where the conclusion (which is the final behaviour of the construct) is achieved if the evaluation of the premises lead to the desired results. For instance, the operational semantics of a while loop could be the following:

$$\frac{\langle c \rangle \Rightarrow \text{true}}{\langle \text{while } c \text{ do } L ; k \rangle \Rightarrow \langle L ; \text{while } c \text{ do } L ; k \rangle}$$

$$\frac{\langle c \rangle \Rightarrow \text{false}}{\langle \text{while } c \text{ do } L ; k \rangle \Rightarrow \langle k \rangle}$$

Again, the behaviour of the semantics rule must be encoded in the host language in which the compiler is being developed, but the pattern it follows is always the same. This

¹Note that the type rule of if-then-else in an imperative programming language is different.

step, depending on the implementation choice, might also require to translate this behaviour into an *intermediate language* representation that is more suitable for the subsequent code generation phase.

2.3 Discussion

The examples above show how the behaviour of the type checking and semantics rules must be hard-coded in the language chosen for the compiler implementation, regardless of the fact that their pattern is constantly repeated in every rule. This pattern can be captured in a meta-language that is able to process the type system and operational semantics definition of the language and produce the code to execute the behaviour of the rules. In this work we describe the meta-language for *Metacasanova*, a meta-compiler that is able to read a program written in terms of type system/operational semantics rules defining a programming language, a program written in that language, and output executable code that mimics the behaviour of the semantics. Such a language relieves the programmer from writing boiler-plate code when implementing a compiler for a (Domain-Specific) language. For this reason we formulate the following research question:

Research question 1: *To what extent does Metacasanova ease the development speed of a compiler for a Domain-Specific Language, in terms of code length compared to the hard-coded implementation, and how much does the abstraction layer of the Metacompiler affect the performance of the generated code?*

Another problem that arises when using meta-compilers is the performance decay given by the introduction of their additional abstraction layer. One of the reasons for this performance decay (see Section 4.3) is that the meta-language (and thus the meta-type system) is unaware of the type system and the memory model of the language implemented in the meta-compiler. For this reason, checking the types and accessing the memory requires to dynamically look up a symbol table defined with the abstractions provided by the meta-language. The need for performance is for *Metacasanova* important because it is being used to extend the DSL for games *Casanova* [1, 2]. Thus, we formulate a second research question:

Research question 2: *In what way can we embed the type system of the implemented language in Metacasanova in order to get rid of the dynamic lookups at runtime and what is the performance gain of this optimization?*

We try to answer these two research questions by using a two-step methodology: (i) we present an architecture for *Metacasanova* aimed to automate the process of code generation, and then (ii) we propose a language extension to embed the implemented language type system in the meta-type system of *Metacasanova*.

2.4 Related Work

RML [17] is a meta-compiler based on operational semantics that is similar to *Metacasanova*. Its syntax is very close to that of ML and it generates C code. A notable effort was done to optimize the tail calls in the generated code for the rules, but the problem arisen by Research Question 2 is not addressed.

Stratego [5] is a meta-compiler based on a transformation system. A transformation language consists of a series of constructor calls to construct the terms of the grammar and functions that specify how to evaluate the terms. *Stratego* is not a typed language, so it does not ensure that the terms and transformation functions are used consistently.

A language extension for Haskell involving *template meta-programming* exists [18]. Although a valuable and elegant approach, using Haskell language extensions is not suitable for domain-specific languages for games due to the wide use of monads and monad transformers, which greatly affect the performance [14], and thunks, which affects the memory usage. In Section 4 we underline how this project was born to ease the extension of a domain-specific language for game development, thus this was not a suitable choice for our initial goals.

Syntax Macro meta-programming [6] is an approach that operates during the parsing phase. Macros are used to produce an abstract syntax tree that is replaced when the macro is invoked. One notable example of this kind of meta-programming can be found in the Lisp language family. Macros guarantee syntactic safety [22], but it is not possible to define the meta-types of the newly introduced syntactic elements.

3 Metacasanova Syntax and Semantics

In the previous section we showed that the process of evaluating typing and semantics rules is always the same, regardless of the specific language implementation. We have also discussed how this evaluation must be re-implemented every time in a hard-coded compiler by using the abstractions provided by the host language, which leads to verbose code and the loss of the clarity and simplicity originally encoded in the type rules and semantics. In this section we define the requirements of *Metacasanova*, we informally present, through an example, how a meta-program works, and we finally propose the syntax and semantics of its meta-language.

3.1 Requirements of Metacasanova

In order to relieve programmers of manually defining the behaviour described in Section 2 in the back-end of the compiler, we propose the following features for *Metacasanova*:

- It must be possible to define custom operators (or functions) and data containers. This is needed to define the syntactic structures of the language we are defining.

- It must be typed: each syntactic structure can be associated to a specific type in order to be able to detect meaningless terms (such as adding a string to an integer) and notify the error to the user.
- It must be possible to have polymorphic syntactical structures. This is useful to define equivalent “roles” in the language for the same syntactical structure; for instance we can say that an integer literal is both a *Value* and an *Arithmetic expression*.
- It must natively support the evaluation of semantics rules, as those shown above.

We can see that these specifications are compatible with the definition of meta-compiler, as the software takes as input a language definition written in the meta-language, a program for that language, and outputs runnable code that mimics the code that a hard-coded compiler would output.

3.2 General Overview

A Metacasanova program is made of a set of Data and Function definitions, and a sequence of rules. A data definition specifies the constructor name of the data type (used to construct the data type), its field types, and the type name of the data. Optionally it is possible to specify a priority for the constructor of the data type. For instance this is the definition of the sum of two arithmetic expressions

```
Data Expr -> "+" -> Expr : Expr
```

Note that Metacasanova allows you to specify any kind of notation for data types in the language syntax, depending on the order of definition of the argument types and the constructor name. In the previous example we used an infix notation. The equivalent prefix and postfix notations would be:

```
Data "+" -> Expr -> Expr : Expr
Data Expr -> Expr -> "+" : Expr
```

A function definition is similar to a data definition but it also has a return type. For instance the following is the evaluation function definition for the arithmetic expression above:

```
Func "eval" -> Expr : Value
```

In Metacasanova it is also possible to define polymorphic data in the following way:

```
Value is Expr
```

In this way we are saying that an atomic value is also an expression and we can pass both a composite expression and an atomic value to the evaluation function defined above.

Metacasanova also allows to embed C# code² into the language by using double angular brackets. This code can be used to embed .NET types when defining data or functions, or to run C# code in the rules. For example in the following snippets we define a floating point data which encapsulates

a floating point number of .NET to be used for arithmetic computations:

```
Data "$f" -> <<float>> : Value
```

A rule in Metacasanova, as explained above, may contain a sequence of function calls and clauses. In the following snippet we have the rule to evaluate the sum of two floating point numbers:

```
eval a => $f c
eval b => $f d
<<c + d>> => res
-----
eval (a + b) => $f res
```

Note that if one of the two expressions does not return a floating point value, then the entire rule evaluation fails. Also note that we can embed C# code to perform the actual arithmetic operation. Metacasanova selects a rule by means of pattern matching (in order of declaration of rules) on the function arguments. This means that both of the following rules will be valid candidates to evaluate the sum of two expressions:

```
...
-----
eval expr => res

...
-----
eval (a + b) => res
```

Finally the language supports expression bindings with the following syntax:

```
x := $f 5
```

3.3 Formalization

In what follows we assume that the pattern matching of the function arguments in a rule succeeds, otherwise a rule will fail to return a result. The informal semantics of the rule evaluation in Metacasanova is the following:

- R1 A rule with no clauses or function calls always returns a result.
- R2 A rule returns a result if all the clauses evaluate to true and all the function calls in the premise return a result.
- R3 A rule fails if at least one clause evaluates to false or one of the function calls fails (returning no results).

We will express the semantics, as usual, in the form of logical rules, where the conclusion is obtained when all the premises are true. In what follows we consider a set of rules defined in the Metacasanova language R . Each rule has a set of function calls F and a set of clauses (boolean expressions) C . We use the notation f^r to express the application of the function f through the rule r . We will define the semantics by using the notation $\langle expr \rangle$ to mark the evaluation of an expression, for

²See Section 4 for the motivation.

example $\langle f^r \rangle$ means evaluating the application of f through r . The following is the formal semantics of the rule evaluation in Metacasanova, based on the informal behaviour defined above:

$$\begin{array}{l}
 \text{R1: } \frac{C = \emptyset \quad F = \emptyset}{\langle f^r \rangle \Rightarrow \{x\}} \\
 \text{R2: } \frac{\forall c_i \in C, \langle c_i \rangle \Rightarrow \text{true} \quad \forall f_j \in F, \exists r_k \in R \mid \langle f_j^{r_k} \rangle \Rightarrow \{x_{r,k}\}}{\langle f^r \rangle \Rightarrow \{x_r\}} \\
 \text{R3(A): } \frac{\exists c_i \in C \mid \langle c_i \rangle \Rightarrow \text{false}}{\langle f^r \rangle \Rightarrow \emptyset} \\
 \text{R3(B): } \frac{\forall r_k \in R, \exists f_j \in F \mid \langle f_j^{r_k} \rangle \Rightarrow \emptyset}{\langle f^r \rangle \Rightarrow \emptyset}
 \end{array}$$

R1 says that, when both C and F are empty (we do not have any clauses or function calls), the rule in Metacasanova returns a result. R2 says that, if all the clauses in C evaluates to true and, for all the function calls in F we can find a rule that returns a result (all the function applications return a result for at least one rule of the program), then the current rule returns a result. R3(a) and R3(b) specify when a rule fails to return a result: this happens when at least one of the clauses in C evaluates to false, or when one of the function applications does not return a result for any of the rules defined in the program.

In the following section we describe how the code generation process works, namely how the Data types of Metacasanova are mapped in the target language, and how the rule evaluation is implemented.

4 Code Generation

In Section 3 we defined the syntax and semantics of Metacasanova. In this section we explain how the abstractions of the language are compiled into the generated code. We chose C# as target language because the development of Metacasanova started with the idea of expanding the DSL for game development Casanova with further functionalities. Casanova hard-coded compiler generates C# code as well because it is compatible with game engines such as Unity3D and Monogame. At the same time, C# grants decent performance without having to manually manage the memory such as for lower-level languages like C/C++. Code generation in different target languages is possible but still an ongoing project (see Section 7).

4.1 Data Structures Code Generation

The type of each data structure is generated as an interface in C#. Each data structure defined in Metacasanova is mapped to a class in C# that implements such interface. The class contains as many fields as the number of arguments the data structure contains. Each field is given an automatic name argC where C is the index of the argument in the data structure definition. The data structure symbols used in the definition might be pre-processed and replaced in order to avoid illegal characters in the C# class definition. The class contains an additional field that stores the original name of the data structure before the replacement is performed, used for its “pretty print”. For example the data structure

Data "\$i" -> int : Value

will be generated as

```

public interface Value { }

public class __opDollari : Value
{
    public string __name = "$i";
    public int __arg0;

    public override string ToString()
    {
        return "(" + __name + " " + __arg0 + ")";
    }
}

```

4.2 Code Generation for Rules

Each rule contains a set of premises that in general call different functions to produce a result, and a conclusion that contains the function evaluated by the current rule and the result it produces. The code generation for the rules follows the steps below:

1. Generate a data structure for each function defined in the meta-program.
2. For each function f extract all the rules whose conclusion contains f .
3. Create a switch statement with a case for each rule that is able to execute the function (the function is in its conclusion).
4. In the case block of each rule, define the local variables defined in the rule.
5. Apply pattern matching to the arguments of the function contained in the conclusion of the rule. If it fails, jump immediately to the next case (rule).
6. Store the values passed to the function call into the appropriate local variables.
7. Run each premise by instantiating the class for the function used by it and copying the values into the input arguments.
8. Check if the premise outputs a result and, in the case of an explicit data structure argument, check the pattern

matching. If the premise result is empty or the pattern matching fails for all the possible executions of the premise then jump to the next case.

9. Generate the result for the current rule execution.

In what follows, we use as an example the code generation for the following rule (which computes the sum of two integer expressions in a programming language):

```
eval a -> $i c
eval b -> $i d
<< c + d >> -> e
-----
eval (a + b) -> $i e
```

From now on we will refer to an argument as *explicit data argument* when its structure appears explicitly in the conclusion or in one of the premises, as in the case of `a + b` in the example above.

4.2.1 Data Structure for the Function

As first step the meta-compiler generates a class for each function defined in the meta-program. This class contains one field for each argument the function accepts. It also contains a field to store the possible result of its evaluation. This field is a struct generated by the meta-compiler defined as follows:

```
public struct __MetaCnvResult<T> { public T Value
; public bool HasValue; }
```

The result contains a boolean to mark if the rule actually returned a result or failed, and a value which contains the result in case of success.

For example, the function

```
Func eval -> Expr : Value
```

will be generated as

```
public class eval
{
    public Expr __arg0;
    public __MetaCnvResult<Value> __res;
    ...
}
```

4.2.2 Rule Execution

The class defines a method `Run` that performs the actual code execution. The meta-compiler retrieves all the rules whose conclusion contains a call to the current function, which define all the possible ways the function can be evaluated with. It then creates a switch structure where each case represents each rule that might execute that function. The result of the rule is also initialized here (the struct will contain a default value and the boolean flag will be set to false). Each case defines a set of local variables, that are the variables used within the scope of that rule.

4.2.3 Local Variables Definitions and Pattern Matching of the Conclusion

At the beginning of each case, the meta-compiler defines the local variables initialized with their respective default values. It also generates then the code necessary for the pattern-matching of the conclusion arguments. Since variables always pass the pattern-matching, the code is generated only for arguments explicitly defining a data structure (see the examples about arithmetic operators in Section 3) and literals. If the pattern matching fails then the execution jumps to the next case (rule). For instance, the code for the following conclusion

```
...
-----
eval (a + b) -> $i e
```

is generated as follows

```
case 0:
{
    Expr a = default(Expr);
    Expr b = default(Expr);
    int c = default(int);
    int d = default(int);
    int e = default(int);
    if (!(__arg0 is __opPlus)) goto case 1;
    ...
}
```

Note that an explicit data argument, such in the example above, might contain other nested explicit data arguments, so the pattern-matching is recursively performed on the data structure arguments themselves.

4.2.4 Copying the Input Values Into the Local Variables

When each function is called by a premise, the local values are stored into the class fields of the function defined in Section 4.2.1. These values must be copied to the local variables defined in the case block representing the rule. Particular care must be taken when one argument is an explicit data. In that case, we must copy, one by one, the content of the data into the local variables bound in the pattern matching. For example, in the rule above, we must separately copy the content of the first and second parameter of the explicit data argument into the local variables `a` and `b`. The generated code for this step, applied to the example above, will be:

```
__opPlus __tmp0 = (__opPlus)__arg0;
a = __tmp0.__arg0;
b = __tmp0.__arg1;
```

Note that the type conversion from the polymorphic type `Expr` into `opPlus` is now safe because we have already checked during the pattern matching that we actually have `opPlus`.

4.2.5 Generation of Premises

Before evaluating each premise, we must instantiate the class for the function that they are invoking. The input arguments of the function call must be copied into the fields of the instantiated object. If one of the arguments is an explicit data argument, then it must be instantiated and its arguments should be initialized, and then the whole data argument must be assigned to the respective function field. After this step, it is possible to invoke the Run method of the function to start its execution. The first premise of the example above then becomes (the generation of the second is analogous):

```
eval a -> $i c
```

```
eval __tmp1 = new eval();
__tmp1.__arg0 = a;
__tmp1.Run();
```

4.2.6 Checking the Premise Result

After the execution of the function called by a premise, we must check if a rule was able to correctly evaluate it. In order to do so, we must check that the result field of the function object contains a value, and if not the rule fails and we jump to the next case (rule), which is performed in the following way:

```
if (!(__tmp1.__res.HasValue)) goto case 1;
```

If the premise was successfully evaluated by one rule, then we must check the structure of the result, which leads to the following three situations: (i) the result is bound to a variable, (ii) the result is constrained to be a literal, and (iii) the result is an explicit data argument. In the first case, as already explained above, the pattern matching always succeeds, so no check is needed. In the second case, it is enough to check the value of the literal. In the last case, all the arguments of the data argument must be checked to see if they match the expected result. In general this process is recursive, as the arguments could be themselves other explicit data arguments. If the result passes the check, then the result is copied into the local variables, in a fashion similar to the one performed for the function premise. For instance, for the premise

```
eval a -> $i c
```

the meta-compiler generates the following code to check the result

```
if (!(__tmp1.__res.Value is __opDollari)) goto
case 1;
__MetaCnvResult<Value> __tmp2 = __tmp1.__res;
__opDollari __tmp3 = (__opDollari)__tmp2.Value;
c = __tmp3.__arg0;
```

4.2.7 Generation of the Result

When all premises correctly output the expected result, the rule can output the final result. In order to do that, the generated code must copy the right part of the conclusion (the result) into the res variable of the function class. If the right part of the conclusion is, again, an explicit data argument, then the data object must first be instantiated and then copied into the result. For example the result of the rule above is generated as follows:

```
res = c + d;
__opDollari __tmp7 = new __opDollari();
__tmp7.__arg0 = res;
__res.HasValue = true;
__res.Value = __tmp7;
break;
```

After this step, the rule evaluation successfully returns a result.

This implementation choice is due to the fact that we plan to support partial function applications, thus, when a function is partially applied, there is the need to store the values of the arguments that were partially given. This could still be implemented with static methods and lambdas in C#, but not all programming languages natively support lambda abstractions, so we chose to have a set-up that allows us to change the target language without dramatically altering the logic of code generation.

4.3 Discussion

Metacasanova has been evaluated in [9] by re-building the DSL for game development Casanova [1, 2]. Even though the size of the code required to implement the language has been drastically reduced (almost 1/5 shorter), performance dropped dramatically. We identified a main problem causing the performance decay that, if solved, will improve the performance of the generated code.

In order to encode a symbol table in the meta-compiler in the current implementation (used for example to store the variables defined in the local scope of a control structure or to model a class/record data structure), we are left with two options: (i) define a custom data structure made of a list of pairs, containing the field/variable name as a string and its value, in the following way

```
Data "table" -> List[Tuple[string, Value]] :
SymbolTable
```

or (ii) use a dictionary data structure coming from .NET, such as ImmutableDictionary, which was the implementation choice for Casanova. In both cases, the behaviour of the language implemented in Metacasanova will be that of a dynamic language, because whenever the value of a variable or class field must be read, the evaluation rule must look up the symbol table at run time to retrieve the value, whose complexity will be $O(n)$ with the list implementation and $O(\log n)$ with the dictionary implementation. This issue is

caused by the fact that, in the current state of Metacasanova, the meta-type system is unaware of the type system of the language that is being implemented in the meta-compiler. This is not a problem limited to Metacasanova but to all meta-compilers having a meta-type system that does not allow embedding of the host language type system. In the next section we propose an extension to Metacasanova to overcome this problem by embedding the type system of the implemented language in the meta-type system of Metacasanova and inlining the code to access the appropriate variable at compile time.

5 Compile-time Inlining with Functors

In Section 3 and Section 4 we presented the semantics of Metacasanova and we showed how the meta-compiler generates the code necessary to represent the elements of the language and the evaluation of the rules expressed in terms of operational semantics. In Section 4.3 we highlighted the problem of performance degradation, due to the additional abstraction layer of the meta-compiler, and identified a possible cause in how the language manages the memory representation. For now, the memory can only be expressed with a dynamic symbol table that must be looked up at run-time in order to retrieve the value of a variable or of a class/record field. In this section we propose an extension to Metacasanova with parametric *Modules* and *Functors* that will allow to inline the access to record fields at compile time and to embed an arbitrary type system into the meta-type system of Metacasanova. Note that in this scope, we use the term functor with the same meaning used in the scope of the language CamL, i.e. a function that takes some types as input and returns a type. In order to provide additional clarity to the explanation, we introduce, in the next section, an example that we use as reference across the whole section. Moreover, note that we introduce the symbol $\Rightarrow$ that denotes that an evaluation happens at compile-time rather than at runtime.

5.1 Case Study

Assume that we want to represent a physical body with a *Position* and a *Velocity* in a 2D space. This can be defined as a data structure containing two fields for its physical properties (the example below is written in F#).

```
type PhysicalBody = {
    Position      : Vector2
    Velocity      : Vector2
}
```

In the current state of the Metacompiler, a language that wants to support such a data structure, as stated in Section 4, should define it either with a list of pairs (*field, value*) or with a dictionary from .NET.

```
Data "Record" -> List[Tuple[string, Value]] :
    Record
```

Accessing the values of the fields requires to iterate through this list (or dictionary) and find the field we want to read, with two evaluation rules such as

```
field = name
-----
getField ((field,value) :: fields) name
  -> value

field <> name
getField fields name -> v
-----
getField ((field,value) :: fields) name -> v
```

This could be done immediately by inlining the *getter* (or *setter*) for that field directly in the program.

In what follows we add a system of modules and functors to Metacasanova, we explain how the meta-compiler generates the code for them, and we show how to use them to improve the performance of the example above.

5.2 Using Modules and Functors in Metacasanova

A module definition in Metacasanova is parametric with respect to types, in the sense that a module definition might contain some type parameters, and can be instantiated by passing the specific types to use. A module can contain the definition of data structures, functions, or functors.

```
Module "Record" : Record {
    Functor "RecordType" : * }
```

The symbol $*$ reads *kind* and means that the functor might return any type. Indeed the type of a record (or class) in a programming language can be “customized” and depends on its specific definition, thus it is not possible to know it beforehand.

We then define two modules for the *getter* and *setter* of a field of a record. In this example, we use type parameters in the module definitions.

```
Module "Getter" => (name : string) => (r : Record
) {
    Functor "GetType" : *
    Func "get" -> (r.RecordType) : GetType }

Module "Setter" => (name : string) => (r : Record
) {
    Functor "SetType" : *
    Func "set" -> (r.RecordType) -> SetType : (r.
        RecordType) }
```

These two modules respectively define a functor to retrieve the type of the record field, and a function to get or set its value. Note that in the function definitions *get* and *set* we are calling the functor of the *Record* module to generate the appropriate type for the signature. This is allowed, since the result of a functor is indeed a type.

A record meta-type (i.e. its representation at meta-language level) is recursively defined as a sequence of pairs (*field, type*),

whose termination is given by `EmptyField`. We thus define the following functors:

```
Functor "EmptyRecord" : Record
Functor "RecordField" => string => * => Record :
  Record
```

The first functor defines the end point of a record, which is simply a record without fields. The second functor defines a field as the pair mentioned above followed by other field definitions.

Moreover, we must define two functors that are able to dynamically build the *getter* and *setter* for the field.

```
Functor "GetField" => string => Record : Getter
Functor "SetField" => string => Record : Setter
```

The behaviour of a functor is expressed, as for normal functions, through a rule in the meta-program. A rule that evaluates a functor returns an instantiation of a module or a type. Note that, inside a module instantiation, it is possible to define and implement functions other than those in the module definition, i.e. the module instantiation must implement *at least* all the functors and functions of the definition. For instance, the following is the type rule instantiating the module for `EmptyRecord`:

```
-----
EmptyRecord => Record {
  Func "cons" : unit
  -----
  RecordType => unit
  -----
  cons -> ()
}
```

The function `cons` defines a constructor for the record, which, in the case of an empty record, returns nothing. The module instantiation for a record field evaluates as well `RecordType`, and has a different definition and evaluation of the function `cons` (because it is constructed in a different way):

```
-----
RecordField name type r = Record {
  Func "cons" -> type -> r.RecordType :
    RecordType
  -----
  RecordType => Tuple[type, r.RecordType]
  -----
  cons x xs -> (x, xs)}
```

Note that the return type of `cons` is to be intended as calling `RecordType` of the current module, so as it were `this.RecordType`. The getter of a field must be able to lookup the record data structure in search of the field and generate a function to get the value from it. For this reason,

the functor instantiates two separate modules, depending on the name of the field that we are currently examining.

Listing 1. Module instantiations for getters

```
//Rule 1
name = fieldName
thisRecord := RecordField name type r
-----
GetField fieldName (RecordField name type r) =>
  Getter fieldName thisRecord {
  -----
  GetType => type
  -----
  get (x,xs) -> x}

//Rule 2
name <> fieldName
thisRecord := RecordField name type r
-----
GetField fieldName (RecordField name type r) =>
  Getter fieldName thisRecord{
  Functor "GetAnotherField" : Getter
  -----
  GetAnotherField => GetField fieldName r
  GetAnotherField => g
  -----
  GetType => g.GetType
  GetAnotherField => getter
  getter.get xs -> v
  -----
  get (x,xs) -> v }
```

Analogously, the setter of a field instantiates two separate modules whether the current field is the one we want to set or not. This can be seen in Listing 2

5.3 Functor Result Inlining

If a premise or a conclusion contains a call to a functor, this call is evaluated at compile time, rather than at runtime. Metacasanova has been extended with an interpreter which is able to evaluate the result of the functor calls. The behaviour of the interpreter follows the same logic explained when presenting the code generation steps in Section 4, thus here we do not present the details for brevity. When a rule outputs the instantiation of the module, the generated code will contain only rules of the modules whose conclusion contains a function (i.e. functions that output values, not functors). In this way the generated code will contain a different version of those functions depending on the instantiation parameters of the module.

We now show how to use the implementation of the records given in Section 5.2 for the physical body presented as a case study. The definition of the record type for the physical body is done through a functor (Listing 3).

Listing 2. Module instantiations for setters

```

name = fieldName
thisRecord := RecordField name type r
-----
SetField fieldName (RecordField name type r) =>
  Setter fieldName thisRecord{

-----
SetType => type

-----
set (x,xs) v -> (v,xs)}

name <> fieldName
thisRecord := RecordField name type r
-----
SetField fieldName (RecordField name type r) =>
  Setter fieldName thisRecord{
  Functor "SetAnotherField" : Setter

-----
SetAnotherField => SetField fieldName r

SetAnotherField => s
-----
SetType => s.SetType

SetAnotherField => setter
setter.set xs v -> xs'
-----
set (x,xs) v -> (x,xs') }

```

Listing 3. Functor for physical body

```

Functor "PhysicalBodyType" : Record

EmptyRecord => empty
RecordField "Velocity" Vector2 empty => velocity
RecordField "Position" Vector2 velocity => body
-----
PhysicalBodyType => body

```

The rule in Listing 3 is evaluated at compile time by the interpreter that generates one module for each field of the PhysicalBody, containing the constructor. For example, for the field Velocity the interpreter will generate³

```

Func "cons" -> Vector2 -> unit : Tuple[Vector2,
  unit]

-----
cons x xs -> (x,xs)

```

This because the functor will call the evaluation rule for RecordField with the argument (RecordField "Velocity" Vector2 (EmptyRecord)). This rule generates the function cons by evaluating the result of the functors EmptyRecord.RecordType and RecordField.RecordType, which respectively produce unit and Tuple[Vector2,unit].

³Note that here we give a high-level representation of the generated rules that are actually directly generated as C# code.

Instantiating a physical body will just require to build a function that returns the type of the physical body, which is obtained by calling the functor PhysicalBodyType.

```

Func "PhysicalBody" : PhysicalBodyType.RecordType

-----
PhysicalBody -> PhysicalBodyType.cons((Vector2.
  Zero,(Vector2.Zero,()))

```

Defining the setter and getter of a field, requires to use the functor GetField to generate the appropriate getter function. After the module has been correctly generated, we can use the getter for the field. For example, in order to get the position field, we use the following function.

```

Func "getPos" : Vector2

GetField "Position" PhysicalBodyType => getter
getter.get PhysicalBody -> p
-----
getPos -> p

```

The result of the premise GetField will be evaluated at compile time through the code in Listing 1 and will instantiate a module containing the following function definition and rule.

```

Func "get" -> Tuple[Vector2,Tuple[Vector2,unit]]
  : Vector2

-----
get (x,xs) -> x

```

Note that the second premise of getPos will immediately call the get generated in this step. The case of setPos is analogous except the setter takes an additional argument.

Reading Velocity analogously uses a functor call to generate a getter:

```

Func "getVel" : Vector2

GetField "Velocity" PhysicalBodyType => getter
getter.get PhysicalBody -> p
-----
getVel -> p

```

This time the functor will generate two different functions in two separate modules. The first time the record is processed, Rule 2 in Listing 1 will be activated (because the first field in the Record is Position). This rule will instantiate an additional module when evaluating the functor call in its premise, which in turn is able to get the Velocity field. The rule for get in the first module will contain in its premise a call to get of the second module (Listing 4).

We want to point out that this optimization has been presented on the specific case of records, but can be generalized for any situations where you would use a symbol table. Indeed any symbol table can be expressed with the representation above as a sequence of pair where the first

item is the value of the current variable, and the second item is the continuation of the symbol table.

Listing 4. Getter rule for velocity

```
//Code for module1
Func "get" -> Tuple[Vector2, Tuple[Vector2, unit]]
: Vector2

module2.get xs -> v
-----
get (x, xs) -> v

//Code for module2 generated by evaluating the
funcion in the premise of Rule 2
Func "get" -> Tuple[Vector2, unit] : Vector2
-----
get (x, xs) -> x
```

6 Evaluation

An extensive evaluation of Casanova implemented in Metacasanova, which we omit for brevity, can be found in [9]. The implementation of Casanova operational semantics in Metacasanova is almost 5 times shorter than the corresponding F# implementation in the hard-coded compiler. In addition to Casanova, we have implemented a subset of the C language called C--. This language supports if-then-else, while-loop, and for statements, as well as local scoping of variables. The total length of the language definition in Metacasanova is 353 lines of code. The corresponding C# code to implement the operational semantics of the language is 3123 lines, thus the code reduction with Metacasanova is roughly 8.84 times. For comparison, in Table 2 it is possible to see the code length to implement three different statements, both in Metacasanova and C#. We tested C-- against Python by computing the average running time to compute the factorial of a number. The choice of Python is due to the fact that both Casanova and C-- exhibits behaviours of dynamic languages, as explained in Section 4.3. C-- results to be 50 times slower than Python. This result is worse than what we obtained when evaluating Casanova, because in order to emulate the interruptible rule mechanism of Casanova in Python you must rely on coroutines that are slower than a program containing simple statements. Moreover, we tested the performance improvement of the optimization using Functors to represent records against the standard one using dynamic symbol tables. The test was run using records with a number of fields ranging from 1 to 10 and updating from 10000 to 1000000 instances of such records. In Table 1, which for brevity shows only the result for 1000000 instances, we can see that the optimization using Functors leads to a performance increase on average of about 11 times, with peaks of almost 30 times. The gain increases with the number of fields, thus Functors are particularly effective for records with high number of fields. Figure 1 shows a chart of the

Table 1. Running time with the functor optimization and the dynamic table with 1000000 records.

FIELDS	Functors (ms)	Dynamic Table (ms)	Gain
1	9.47E-04	7.29E-04	0.77
2	9.51E-04	1.78E-03	1.87
3	9.50E-04	3.33E-03	3.51
4	9.60E-04	5.43E-03	5.66
5	9.65E-04	8.03E-03	8.32
6	9.71E-04	1.11E-02	11.44
7	9.75E-04	1.47E-02	15.12
8	9.82E-04	1.89E-02	19.28
9	9.92E-04	2.37E-02	23.86
10	1.00E-03	2.87E-02	28.62
Average gain			11.84

Table 2. Code length implementation of C-- and run-time performance

Statement	Metacasanova	C#
if-then-else	4	103
while	7	73
For	11	81

C--	Python
1.26ms	$2.36 \cdot 10^{-2}$ ms

overall performance of the two techniques (the data points are taken from Table 1). The horizontal axis contains the amount of fields per record, while the vertical axis contains the number of records that are being updated. We can see that the performance of the dynamic table degrades considerably when increasing the number of fields, and that the higher the amount of records is, the steeper the curve is. On the other hand, the performance of the implementation with Functors is almost constant, regardless of the amount of fields or records that are being updated. Moreover, note that the performance of the dynamic table is improved by the fact that we are using a dictionary implemented in .NET, which can access the entries in $O(\log n)$. If the symbol table were represented as a meta-data structure in the language the performance would be even worse, since it would have to be encoded as a list of pairs with the field name and its value, and its manipulation would be affected by the evaluation rules that should implement this behaviour. Furthermore, the dynamic lookup should be done also to ensure that the types of the record fields are used consistently (for example to prevent that a record is constructed with incompatible values for its fields), while using the functors in Metacasanova embeds the type system of the language in the meta-type system, whose type safety is checked at compile-time rather than at runtime, and this contributes to further increase the performance.

7 Conclusion

In this work we presented the architecture of a Metacom-
piler called Metacasanova, whose meta-language is based on
the operational semantics. In Section 2 we discussed how
it is possible to capture repetitive patterns in designing a
compiler for DSL through a meta-compiler. We presented the
meta-compiler Metacasanova and its meta-language. Meta-
casanova has been evaluated by re-implementing the DSL
for games Casanova, and by implementing a subset of the C
language, called C $\rightarrow$. Our results show that implementing
the language semantics in Metacasanova is up to 8 times
shorter than with a hard-coded compiler. The additional ab-
straction layer of the meta-compiler leads to a performance
decay that, in the case of C $\rightarrow$, makes the language 50 times
slower than Python, and in the case of Casanova on the same
order but still 3 times slower. We identified the problem in
the fact that the meta-language is unaware of the type sys-
tem and memory model of the implemented language, thus
all type checks and field lookups must be done dynamically
at runtime adding the overhead of a dynamic lookup table.
We have proposed a language extension based on *Modules*
and *Functors* that allows the meta-language to embed the
type system of the implemented language in the meta-type
system, and to inline the lookups directly into the generated
code. This optimization leads to a performance improvement
factor of 11, which peaks to 30 in presence of many updates
and data structures with many fields.

References

- [1] Mohamed Abbadi. 2017. *Casanova 2, A domain specific language for general game development*. Ph.D. Dissertation. Università Ca' Foscari, Tilburg University.
- [2] Mohamed Abbadi, Francesco Di Giacomo, Agostino Cortesi, Pieter Spronck, Giulia Costantini, and Giuseppe Maggiore. 2015. Casanova: a simple, high-performance language for game development. In *Joint International Conference on Serious Games*. Springer, 123–134.
- [3] Alfred V Aho, Ravi Sethi, and Jeffrey D Ullman. 1986. *Compilers, Principles, Techniques*. Addison wesley Boston.
- [4] Erwin Book, Dewey Val Shorre, and Steven J Sherman. 1970. The cwic/360 system, a compiler for writing and implementing compilers. *ACM SIGPLAN Notices* 5, 6 (1970), 11–29.
- [5] Martin Bravenboer, Karl Trygve Kalleberg, Rob Vermaas, and Eelco Visser. 2008. Stratego/XT 0.17. A language and toolset for program transformation. *Science of computer programming* 72, 1 (2008), 52–70.
- [6] WR Campbell. 1978. A compiler definition facility based on the syntactic macro. *Comput. J.* 21, 1 (1978), 35–41.
- [7] Luca Cardelli. 1996. Type systems. *Comput. Surveys* 28, 1 (1996), 263–264.
- [8] Krzysztof Czarnecki, Ulrich W Eisenecker, G Goos, J Hartmanis, and J van Leeuwen. 2000. Generative programming. *Edited by G. Goos, J. Hartmanis, and J. van Leeuwen* 15 (2000).
- [9] Francesco Di Giacomo, Mohamed Abbadi, Agostino Cortesi, Pieter Spronck, and Giuseppe Maggiore. 2017. Building Game Scripting DSL's with the Metacasanova Metacompile. In *INTETAIN 2016, Utrecht, The Netherlands, June 28–30*. Springer, 231–242.

Figure 1. Execution time of the different memory models

- [10] Plotkin G.D. 1981. *A structural approach to operational semantics*. Technical Report. Computer science department, Aarhus University.
- [11] F. Di Giacomo, M. Abbadi, A. Cortesi, P. Spronck, G. Costantini, and G. Maggiore. 2017. High performance encapsulation and networking in Casanova 2. *Entertainment Computing* 20 (2017), 25 – 41.
- [12] Gilles Kahn. 1987. Natural semantics. *STACS 87* (1987), 22–39.
- [13] Samuel N Kamin. 1998. Research on domain-specific embedded languages and program generators. *Electronic Notes in Theoretical Computer Science* 14 (1998), 149–168.
- [14] O. Kiselyov. 2016. Free and Freer Monads: Putting Monads Back into Closet. <http://okmij.org/ftp/Computation/free-monad.html>. (2016).
- [15] Jan Willem Klop et al. 1992. Term rewriting systems. *Handbook of logic in computer science* 2 (1992), 1–116.
- [16] Marjan Mernik, Jan Heering, and Anthony M Sloane. 2005. When and how to develop domain-specific languages. *ACM computing surveys (CSUR)* 37, 4 (2005), 316–344.
- [17] Mikael Pettersson. 1996. A compiler for natural semantics. In *Compiler Construction*. Springer, 177–191.
- [18] Tim Sheard and Simon Peyton Jones. 2002. Template meta-programming for Haskell. In *Proceedings of the 2002 ACM SIGPLAN workshop on Haskell*. ACM, 1–16.
- [19] Anthony M Sloane. 2002. Post-design domain-specific language embedding: A case study in the software engineering domain. In *System Sciences, 2002. HICSS. Proceedings of the 35th Annual Hawaii International Conference on*. IEEE, 3647–3655.
- [20] Arie Van Deursen, Paul Klint, Joost Visser, et al. 2000. Domain-specific languages: An annotated bibliography. *Sigplan Notices* 35, 6 (2000), 26–36.
- [21] Markus Voelter, Sebastian Benz, Christian Dietrich, Birgit Engelmann, Mats Helander, Lennart CL Kats, Eelco Visser, and Guido Wachsmuth. 2013. *DSL engineering: Designing, implementing and using domain-specific languages*. dslbook.org.
- [22] Daniel Weise and Roger Crew. 1993. Programmable syntax macros. In *ACM SIGPLAN Notices*, Vol. 28. ACM, 156–165.