# Optimizing Pattern Matching by Program Transformation 

Emilie Balland, Pierre-Etienne Moreau

## To cite this version:

Emilie Balland, Pierre-Etienne Moreau. Optimizing Pattern Matching by Program Transformation. [Technical Report] 2006, pp.19. inria-00001127v1

## HAL Id: inria-00001127 https://inria.hal.science/inria-00001127v1

Submitted on 21 Feb 2006 (v1), last revised 15 Jun 2006 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Optimizing Pattern Matching by Program Transformation 

Emilie Balland and Pierre-Etienne Moreau

UHP \& LORIA and INRIA \& LORIA


#### Abstract

The compilation of pattern matching constructs is crucial to the efficient implementation of functional languages like ML, Caml, or Haskell as well as rewrite rule based languages such as ASF+SDF, ELAN, Maude, or Stratego for example. Until now, the classical approach was to compute an (optimized) automaton before generating, in a straight-forward way, the corresponding implementation code. Optimizations such as tests-sharing are encoded in the construction of the automaton. While efficient, this leads to algorithms which are often complex and difficult to extend and to maintain. In this paper we present a new compilation and optimization method based on program transformation. The principle is to separate the compilation of pattern matching from the optimization, in order to improve modularity and make extensions simpler. In a first step, the patterns are compiled using a simple, but safe algorithm. Then, optimizations are directly performed on the generated code, using transformation rules. Separating optimization from compilation eases the compilation of extensions, such as new equational theories, or the addition of or-patterns for example. Another contribution of this paper is to define a set of rules which defines the optimization, to show their correction as well as their effectiveness on real programs. The presented approach has been implemented and applied to Tom, a language extension which adds pattern-matching facilities to $C$ and Java.


## 1 Introduction

Pattern matching is an elegant high-level construct which appears in many programming languages. Similarly to method dispatching in object oriented languages, in functional languages like ML [1], or rewrite rule bases languages like ASF+SDF [7|14], ELAN [4], Maude [6] or Tom [12], the notion of pattern matching is essential: it is part of the main execution mechanism.

In this context, it is necessary to have an efficient compilation of this construction. There exists several methods $5 / 2918$ to compile pattern matching. The simplest one is to consider and compile each pattern independently. These kind of algorithms are called one-to-one because they operate on only one pattern and one subject. A more efficient approach consists in considering the system globally and building a discrimination network to efficiently select the pattern that matches. These methods are called many-to-one, and they usually consist of three phases: first constructing an automaton, then optimizing it, and finally generating the implementation code. There are two main approaches to construct a matching automaton. The first is based on decision trees [59], which correspond
to deterministic automata. By introducing some redundancy, the idea is to ensure that every position of a given term is tested at most once. As a counterpart, the size of the decision tree becomes exponential in the number and the size of patterns. The second approach is based on backtracking automata [2], to avoid duplicating code. As a consequence, the efficiency can no longer be linear in the size of the subject: the compromise between speed and memory space is unavoidable [13].

In this paper, we present a new approach for pattern-matching compilation where the optimization phase is kept separated from the compilation phase. This allows us to design algorithms which are simpler to implement, easier to extend, and that can be formally certified 1110. In addition, this work allows to generate efficient implementations.

The presented work takes part of the Tom project, whose goal is to integrate the notion of pattern matching into classical languages such as $C$ and Java. As presented in [12], a Tom program is a program written in a host language and extended by some new instructions like the \%match construct. Therefore, a program can be seen as a list of Tom constructs interleaved with some sequences of characters. During the compilation process, all Tom constructs are dissolved and replaced by instructions of the host-language, as it is usually done by a pre-processor. The following example shows how a simple symbolic computation (addition) over Peano integers can be defined. This supposes the existence of a datastructure where Peano integers are represented by zero and successor: the integer 3 is denoted by suc(suc(suc(zero))) for example.

```
public class PeanoExample {
    Term plus(Term t1, Term t2) {
        %match(t1, t2) {
            x,zero -> { return x; }
            x,\operatorname{suc(y) -> { return suc(plus(x,y)); }}
        }
    }
    void run() {
        System.out.println("plus(1,2) = " + plus(suc(zero),suc(suc(zero))));
    }
}
```

In this example, given two terms $t_{1}$ and $t_{2}$ (that represent Peano integers), the evaluation of plus returns the sum of $t_{1}$ and $t_{2}$. This is implemented by pattern matching: $t_{1}$ is matched by $x, t_{2}$ is possibly matched by the two patterns zero and $\operatorname{suc}(y)$. When zero matches $t_{2}$, the result of the addition is $x$ (with $x=t_{1}$, instantiated by matching). When $\operatorname{suc}(y)$ matches $t_{2}$, this means that $t_{2}$ is rooted by a suc symbol: the subterm $y$ is added to $x$ and the successor of this number is returned. The definition of plus is given in a functional programming style, but the plus function can be used in Java to perform computations. This first example illustrates how the \%match construct can be used in conjunction with the considered native language.

In order to understand the choices we have made when designing the pattern matching algorithm, it is important to consider Tom as a restricted compiler (like a pre-processor)
which does not have any information about the host-language. In particular, the datastructure, against which the pattern matching is performed, is not fixed. As a consequence, we cannot assume that a function symbol like suc or zero is represented by an integer, like it is commonly done in other implementations of pattern matching. In some sense, the data-structure is a parameter of the pattern matching, see [10] for more details.

To allow this flexibility and support several host-languages, a pattern matching problem is compiled into an intermediate language code, called PIL, before being compiled into the selected host-language. The purpose of this work is to add an optimizing phase which performs program transformation on the PIL code. The general architecture of Tom is illustrated in Figure 1.


Fig. 1. General architecture of Tom: the compiler generates an intermediate PIL program which is optimized before being pretty-printed by the backend.

Optimization can been seen independently from the compilation step. In our case, optimization is called source-to-source because the input and the target languages are the same. In order to manipulate the program, we consider an abstract representation, e.g. the abstract syntactic tree (AST). By representing each instruction by a node, this structure is very closed to the initial code. Using program transformation to do source-to-source optimization has several advantages: the transformations can be easily described using rewrite rules, allowing us to test and activate them separately. The order and combination of different rules can be described using a strategy languages a la ELAN [4] or Stratego [15].

Road-map of the paper. In Section 2, we present the intermediate language PIL and its semantics. In Section 3, we present a rewrite system which describe an optimizer. We also define a strategy to apply efficiently these rules. In Section 4 we show that the presented rewrite system, and thus optimizations are correct. Finally, we present some experimental results in several revealing examples.

## 2 PIL language

In this section we define the syntax and the semantics of our intermediate language, called PIL.

### 2.1 Preliminary concepts

We assume the reader to be familiar with the basic definitions of first order term given, in particular, in 3. We briefly recall or introduce the notation for a few concepts that will be used along this paper.

A signature $\mathcal{F}$ is a set of function symbols, each one associated to a natural number by the arity function $(\operatorname{ar}: \mathcal{F} \rightarrow \mathbb{N}) . \mathcal{T}(\mathcal{F}, \mathcal{X})$ is the set of terms built from a given finite set $\mathcal{F}$ of function symbols and a enumerable set $\mathcal{X}$ of variables. Positions in a term are represented as sequences of integers and denoted by $\wedge$ or $\omega$. The set of positions is noted $\Omega$. The empty sequence $\wedge$ denotes the position associated to the root, and it is called the root (or top) position. The subterm of $t$ at position $\omega$ is denoted $t_{\mid \omega}$. The replacement in $t$ of the subterm $t_{\mid \omega}$ by $t^{\prime}$ is denoted $t\left[t^{\prime}\right]_{\omega}$.

The set of variables occurring in a term $t$ is denoted by $\mathcal{V a r}(t)$. If $\operatorname{Var}(t)$ is empty, $t$ is called a ground term and $\mathcal{T}(\mathcal{F})$ is the set of ground terms.
$\operatorname{Symb}(t)$ is a partial function from $\mathcal{T}(\mathcal{F}, \mathcal{X})$ to $\mathcal{F}$, which associates to each term $t$ its root symbol $f \in \mathcal{F}$. Two ground terms $t$ and $u$ of $\mathcal{T}(\mathcal{F})$ are equal, and we note $t=u$, when, for some function symbol $f, \operatorname{Symb}(t)=\operatorname{Symb}(u)=f, f \in \mathcal{F}, t=f\left(t_{1}, \ldots, t_{n}\right)$, $u=f\left(u_{1}, \ldots, u_{n}\right)$, and $\forall i \in[1 . . n], t_{i}=u_{i}$.

A substitution $\sigma$ is an assignment from $\mathcal{X}$ to $\mathcal{T}(\mathcal{F})$, written, when its domain is finite, $\sigma=\left\{x_{1} \mapsto t_{1}, \ldots, x_{k} \mapsto t_{k}\right\}$. It uniquely extends to an endomorphism $\sigma^{\prime}$ of $\mathcal{T}(\mathcal{F}, \mathcal{X})$ : $\sigma^{\prime}(x)=\sigma(x)$ for each variable $x \in \mathcal{X}, \sigma^{\prime}\left(f\left(t_{1}, \ldots, t_{n}\right)\right)=f\left(\sigma^{\prime}\left(t_{1}\right), \ldots, \sigma^{\prime}\left(t_{n}\right)\right)$ for each function symbol $f \in \mathcal{F}$.

Given a pattern $p \in \mathcal{T}(\mathcal{F}, \mathcal{X})$ and a ground term $t \in \mathcal{T}(\mathcal{F})$, $p$ matches $t$, written $p \ll t$, if and only if there exists a substitution $\sigma$ such that $\sigma(p)=t$.

A conditional rewrite rule is an ordered tuple of terms $\langle l, r, c>$ noted $l \rightarrow r$ IF $c$ such that $\operatorname{Var}(r) \subseteq \operatorname{Var}(l)$ and $\operatorname{Var}(c) \subseteq \mathcal{V} \operatorname{Var}(l)$. Given a set $R$ of rewrite rules, called a rewriting system, a term $t \in \mathcal{T}(\mathcal{F})$ is rewritten into $t^{\prime}$ if there exist: a rule $l \rightarrow r$ IF $c \in R$, a position $\omega$, and a substitution $\sigma$ such that $t_{\mid \omega}=\sigma(l)$ and $\sigma(c)$ is true. In that case, $t^{\prime}=t[\sigma(r)]_{\omega}$. When no more rule can be applied, the term obtained is said in normal form and noted $t \downarrow_{R}$.

### 2.2 Syntax

The syntax of PIL is given in Figure 2. As mentioned previously, the data-model of this language is a parameter. In practice, this means that a meta-language has to be used to define on which type of data-structure the expressions and instructions operate. For expository reasons, in the rest of this paper, we consider that first order terms define the data-model. Similarly to functional programming languages, given a signature $\mathcal{F}$ and a set of variables $\mathcal{X}$, the considered PIL language can directly handle terms and perform operations like checking that a given term $t$ is rooted by a $\operatorname{symbol} f$ (is_fsym $(t, f)$ ), or accessing to n-th child of a this term $t\left(\operatorname{subterm}_{f}(t, n)\right)$.

The intermediate language PIL has both functional and imperative flavors: the assignment instruction (let(variable,term,instr)) defines a scoped unmodifiable variable, whereas the sequence instruction (instr;instr) comes from imperative languages. A last particularity of PIL comes from the hostcode(...) instruction which is used to abstract

```
PIL \(\quad:=\langle\) instr \(\rangle\)
symbol \(::=f \in \mathcal{F}\)
variable \(::=x \in \mathcal{X}\)
\(\langle\) term \(\rangle \quad::=t \in \mathcal{T}(\mathcal{F}, \mathcal{X})\)
    \(\mid\) subterm \(_{f}(\langle\) term \(\rangle, n)\)
    \((f \in \mathcal{F} \wedge n \in \mathbb{N})\)
```



Fig. 2. PIL syntax
part of code written in the underlying host-language. This instruction is parametrized by a list of PIL-variables which are used in this part of code. Examples of Tom, and PIL code are given in figure 3

```
Tom code:
%match(Term t) {
    f(a)}=>{{\operatorname{print(...);}
    g(x)}=>{\operatorname{print(\ldotsx...);}
    f(b)}=>{\operatorname{print}(\ldots);
}
Generated PIL code:
if(is_fsym (t,f), let (t t, subtermf(t, 1),
    if(is_fsym(t t a ), hostcode(), nop)),
    nop);
if(is_fsym (t,g), let (t t , subtermg}(t,1)
    let(x, th, hostcode(x)))
    nop) ;
if(is_fsym (t,f), let (t t, subtermf(t, 1),
    if(is_f_sym(t. (t) b), hostcode(), nop)),
    nop)
```

Fig. 3. Example of PIL code generated by Tom

### 2.3 Semantics

We define PIL semantics as in [10 by a big-step semantics à la Kahn. First, we introduce the notion of environment, which models the memory of a program during its evaluation. To represent a substitution, we model an environment by a stack of assignments of terms to variables.

Definition 1. An atomic environment $\epsilon$ is an assignment from $\mathcal{X}$ to $\mathcal{T}(\mathcal{F})$, written $[x \leftarrow t]$. The composition of environments is left-associative, and written $\left[x_{1} \leftarrow t_{1}\right]\left[x_{2} \leftarrow\right.$ $\left.t_{2}\right] \cdots\left[x_{k} \leftarrow t_{k}\right]$. Its application is such that: $\epsilon[x \leftarrow t](y)=t$ if $y \equiv x, \epsilon(y)$ otherwise.

We note $\epsilon_{\mid \mathrm{x}}$, the environment $\epsilon$ whose domain is restricted to $X$. We extend the notion of environment to a morphism $\epsilon$ from PIL to PIL, and we note $\mathcal{E} n v$ the set of all environments.

The reduction relation of this big-step semantics is expressed on tuples $\langle\epsilon, \delta, i\rangle$ where $\epsilon$ is an environment, $\delta$ is a list of pairs (restricted environment, hostcode), and $i$ is an
instruction. Thanks to $\delta$, we can keep track of the executed hostcode blocks within their environment: the restricted environment associated to each hostcode construct gives the instances of all variables which appear in the block. The reduction relation is the following:

$$
\langle\epsilon, \delta, i\rangle \mapsto_{b s} \delta^{\prime}, \text { with } \epsilon \in \mathcal{E} n v, \delta, \delta^{\prime} \in[\mathcal{E} n v,\langle i n s t r\rangle]^{*}, \text { and } i \in\langle i n s t r\rangle
$$

Before giving the semantics for the instructions, we define a rewrite system, denoted $R$, which describes how expressions and terms are evaluated in a PIL program:

$$
R=\left\{\begin{array}{lll}
\operatorname{is\_ fsym}(t, f) & \rightarrow \operatorname{Symb}(t)=f & \text { IF } t \in \mathcal{T}(\mathcal{F}) \\
\operatorname{eq}\left(t_{1}, t_{2}\right) & \rightarrow t_{1}=t_{2} & \text { IF } t_{1}, t_{2} \in \mathcal{T}(\mathcal{F}) \\
\operatorname{subterm}_{f}(t, i) \rightarrow t_{\mid i} & \text { IF } \operatorname{Symb}(t)=f \wedge i<\operatorname{arity}(f)
\end{array}\right.
$$

Proposition 1. Given $t \in\langle$ term $\rangle$ (resp. $t \in\langle$ expr $\rangle$ ), if $t$ contains no variable, $t \downarrow_{R} \in \mathcal{T}(\mathcal{F})$ (resp. $t \downarrow_{R} \in \mathbb{B}$ ).

Proof. We reason by induction on el $\in\langle$ expr $\rangle \cup\langle$ term $\rangle$ :

- when $e$ is a term from $\mathcal{T}(\mathcal{F}, \mathcal{X})$, no rule can be applied. Since it contains no variable, we have $e \downarrow_{R}=t \in \mathcal{T}(\mathcal{F})$,
- when $e=\operatorname{subterm}_{f}(t, i)$ : by induction $t \downarrow_{R} \in \mathcal{T}(\mathcal{F})$, the third rule applies and we obtain $e \downarrow_{R}=\left(t_{\mid i}\right) \downarrow_{R} \in \mathcal{T}(\mathcal{F})$
- when $e=b \in \mathbb{B}$ : no rule applies, $e$ is in normal form and $e \downarrow_{R} \in \mathbb{B}$,
- when $e=\mathrm{eq}\left(t_{1}, t_{2}\right)$ : by induction, $t_{1}, t_{2} \in \mathcal{T}(\mathcal{F})$. Therefore, one application of $R$ gives the normal form $t_{1}=t_{2}$ which is true or false, and we have $e \downarrow_{R} \in \mathbb{B}$,
- when $e=$ is_fsym $(t, f)$ : same reasoning.

| $\overline{\langle\epsilon, \delta, \text { nop }\rangle \mapsto_{b s} \delta}$ | (nop) |
| :--- | :--- |
| $\frac{\langle\epsilon[x \leftarrow t], \delta, i\rangle \mapsto_{b s} \delta^{\prime} \quad \epsilon(u) \rightarrow_{R *} t}{\langle\epsilon, \delta, \operatorname{let}(x, u, i)\rangle \mapsto_{b s} \delta^{\prime}}$ | (let) |
| $\frac{\left\langle\epsilon, \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime} \quad \epsilon(e) \rightarrow_{R *} \top}{\left\langle\epsilon, \delta, \text { if }\left(e, i_{1}, i_{2}\right)\right\rangle \mapsto_{b s} \delta^{\prime}}$ | (iftrue) |
| $\frac{\left\langle\epsilon, \delta, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime} \quad \epsilon(e) \rightarrow_{R *} \perp}{\left\langle\epsilon, \delta, \text { if }\left(e, i_{1}, i_{2}\right)\right\rangle \mapsto_{b s} \delta^{\prime}}$ | (iffalse) |
| $\frac{\left\langle\epsilon, \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}}{\left\langle\epsilon, \delta^{\prime}, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}}\left\langle\epsilon, \delta, i_{1} ; i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}$ | (seq) |
| $\frac{\text { (hostcode) }}{\langle\epsilon, \delta, \text { hostcode(list) }\rangle \mapsto_{b s} \delta::\left[\epsilon_{\mid l i s t}, \text { hostcode(list)] }\right.}$ |  |

Fig. 4. Big-step semantics for PIL

### 2.4 Properties

Definition 2. A program $\pi \in$ PIL is said to be well-formed when it satisfies the following properties:

- each expression $\operatorname{subterm}_{f}(t, n)$ is such that $t$ belongs to $\langle$ term , is_fsym $(t, f) \equiv \top$ and $n \in[1 . . \operatorname{ar}(f)]$,
(In practice, we verify that each expression of the form $\operatorname{subterm}_{f}(t, n)$ belongs to the then part of an instruction if (is_fsym $(t, f), \ldots)$ )
- each variable appearing in a sub-expression is previously initialized by a let construct, or in the evaluation environment,
- in the construction let $(v, u, i)$, the instruction block $i$ must not contain an instruction $\operatorname{let}\left(v, u^{\prime}, i^{\prime}\right)$ (i.e. a variable cannot be redefined).

Proposition 2. For all $\epsilon \in \mathcal{E} n v$ and $\delta \in[\mathcal{E} n v,\langle\text { instr }\rangle]^{*}$, the reduction of a well-formed instruction $i \in\langle$ instr $\rangle$ in the environment $\epsilon, \delta$ is unique.

Proof. The proof can be done by induction on program structures, see [10] for details.
Definition 3. Given $\pi_{1}$ and $\pi_{2}$ two well-formed PIL programs, they are semantically equivalent, noted $\pi_{1} \sim \pi_{2}$, iff:

$$
\forall \epsilon, \delta, \exists \delta^{\prime} \text { s.t. }\left\langle\epsilon, \delta, \pi_{1}\right\rangle \mapsto_{b s} \delta^{\prime} \text { and }\left\langle\epsilon, \delta, \pi_{2}\right\rangle \mapsto_{b s} \delta^{\prime}
$$

Given a well-formed program $\pi \in \operatorname{PIL}$, a position $\omega$, an environment $\epsilon$, and a hostcode list $\delta$, we define $\Phi$ the function that gives the environment $\epsilon^{\prime}$ during the evaluation of $\pi_{\mid \omega}$ in the reduction tree of $\pi:\langle\epsilon, \delta, \pi\rangle \mapsto_{b s} \delta^{\prime}$. (Note that $\pi_{\mid \omega}$ can be either an expression, a term or an instruction).

Definition 4. The function $\Phi \in \operatorname{PIL} \times \Omega \times \mathcal{E} n v \times[\mathcal{E} n v,\langle$ instr $\rangle] \rightarrow \mathcal{E} n v$ is defined such that for a given $\pi, \omega, \epsilon, \delta$ and with $\epsilon^{\prime}=\Phi(\pi, \omega, \epsilon, \delta)$ :

- if $\pi_{\mid \omega}=i \in\langle$ instr $\rangle, \epsilon^{\prime}$ is the environment $\epsilon$ in which the rule corresponding to $i$ (let, nop, seq, iftrue, iffalse or hostcode) is applied,
- if $\pi_{\mid \omega}=e \in\langle$ expr $\rangle$, e appears inevitably in an instruction if noted $i$ and we have $\epsilon^{\prime}=\Phi(\pi, \epsilon, \delta, i)$,
- if $\pi_{\mid \omega}=t \in\langle$ term $\rangle$, we must distinguish two cases. Either $t$ appears in an expression $e$, and we have $\epsilon^{\prime}=\Phi(\pi, \epsilon, \delta, e)$. Either $t$ appears in an instruction let $\left({ }_{-}, t,-\right)$ noted $i$, and we have $\epsilon^{\prime}=\Phi(\pi, \epsilon, \delta, i)$.

Note that $\Phi$ is defined only for $\pi_{\mid \omega}$ which are reachable at run time.
Let us consider again the program defined in Figure 3. We have:
$-\pi=\operatorname{if}\left(\right.$ is_fsym $^{\prime}(t, f), \operatorname{let}\left(t_{1}, \operatorname{subterm}_{f}(t, 1), \pi^{\prime}\right)$, nop $)$ with
$-\pi^{\prime}=\operatorname{if}\left(\pi^{\prime \prime}\right.$, hostcode(), nop) and
$-\pi^{\prime \prime}=$ is_fsym $\left(t_{1}, a\right)$.

Starting from the environments $\epsilon=[t \leftarrow f(a)]$ and $\delta=\{ \}$, we can construct the following derivation tree for $\pi$. From this derivation tree, we have: $\Phi\left(\pi, \epsilon, \delta, \pi^{\prime \prime}\right)=\epsilon\left[t_{1} \leftarrow a\right]$.

For layout purpose, we note $\Delta$ the list $\delta::[\{ \}$, hostcode ()$]$, and $\epsilon^{\prime}=\epsilon\left[t_{1} \leftarrow a\right]$ :

$$
\begin{aligned}
& \frac{\epsilon^{\prime}\left(\text { is_fsym }\left(t_{1}, a\right)\right) \rightarrow_{R e *} T \overline{\left\langle\epsilon^{\prime}, \delta, \text { hostcode }()\right\rangle \mapsto_{b s} \Delta}}{\frac{\left\langle\epsilon\left[t_{1} \leftarrow a\right], \delta, \operatorname{if}\left(\pi^{\prime \prime}, \text { hostcode }(), \text { nop }\right)\right\rangle \mapsto_{b s} \Delta}{\langle\epsilon}(\text { let })} \text { (iftruestcode) } \\
& \frac{\epsilon(\text { is_fsym }(t, f)) \rightarrow_{R *} T}{\left\langle\epsilon, \delta, \operatorname{if}\left(\operatorname{is\_ fsym}(t, f), \operatorname{let}\left(t_{1}, \operatorname{subterm}_{f}(t, 1), \pi^{\prime}\right), \text { nop }\right)\right\rangle \mapsto_{b s} \Delta}(\text { iftrue })
\end{aligned}
$$

Definition 5. Given $\pi_{1}$ and $\pi_{2}$ two well-formed PIL programs, they are semantically equivalent, noted $\pi_{1} \sim \pi_{2}$, when:

$$
\forall \epsilon, \delta, \exists \delta^{\prime} \text { s.t. }\left\langle\epsilon, \delta, \pi_{1}\right\rangle \mapsto_{b s} \delta^{\prime} \text { and }\left\langle\epsilon, \delta, \pi_{2}\right\rangle \mapsto_{b s} \delta^{\prime}
$$

Definition 6. Given a program $\pi$, two expressions $e_{1}$ and $e_{2}$ are said $\pi$-equivalent, and noted $e_{1} \sim_{\pi} e_{2}$, if for all starting environment $\epsilon$, $\delta, \epsilon_{1}\left(e_{1}\right) \downarrow_{R}=\epsilon_{2}\left(e_{2}\right) \downarrow_{R}$ where $\epsilon_{1}=$ $\Phi\left(\pi, \epsilon, \delta, e_{1}\right)$ and $\epsilon_{2}=\Phi\left(\pi, \epsilon, \delta, e_{2}\right)$.
Definition 7. Given a program $\pi$, two expressions $e_{1}$ and $e_{2}$ are said $\pi$-incompatible, and noted $e_{1} \perp_{\pi} e_{2}$, if for all starting environment $\epsilon$, $\delta, \epsilon_{1}\left(e_{1}\right) \downarrow_{R} \wedge \epsilon_{2}\left(e_{2}\right) \downarrow_{R}=\perp$ where $\epsilon_{1}=\Phi\left(\pi, \epsilon, \delta, e_{1}\right)$ and $\epsilon_{2}=\Phi\left(\pi, \epsilon, \delta, e_{2}\right)$.

We can now define two conditions which are sufficient to determine whether two expression are $\pi$-equivalent or $\pi$-incompatible. Propositions 3 and 4 are interesting because they can be easily used in practice.

Proposition 3. Given a program $\pi$ and two expressions $e_{1}, e_{2} \in\langle$ expr $\rangle$, we have $e_{1} \sim_{\pi} e_{2}$ if: $\forall \epsilon, \delta, \Phi\left(\pi, \epsilon, \delta, e_{1}\right)=\Phi\left(\pi, \epsilon, \delta, e_{2}\right)$ (cond1) and $e_{1}=e_{2}$ (cond2).

Proof. Given a program $\pi$ and two expressions $e_{1}, e_{2}$, for all $\delta, \epsilon$, because of cond1 and cond2, $\Phi\left(\pi, \epsilon, \delta, e_{1}\right)\left(e_{1}\right)=\Phi\left(\pi, \epsilon, \delta, e_{2}\right)\left(e_{2}\right)$ therefore $\Phi\left(\pi, \epsilon, \delta, e_{1}\right)\left(e_{1}\right) \downarrow_{R}=$ $\Phi\left(\pi, \epsilon, \delta, e_{2}\right)\left(e_{2}\right) \downarrow_{R}$.

Proposition 4. Given a program $\pi$ and two expressions $e_{1}, e_{2} \in\langle$ expr $\rangle$, we have $e_{1} \perp_{\pi}$ $e_{2}$ if: $\forall \epsilon, \delta, \Phi\left(\pi, \epsilon, \delta, e_{1}\right)=\Phi\left(\pi, \epsilon, \delta, e_{2}\right)$ (cond1) and incompatible $\left(e_{1}, e_{2}\right)$ (cond2), where incompatible defined as follows:

$$
\begin{aligned}
& \text { incompatible }\left(e_{1}, e_{2}\right)=\text { match } e_{1}, e_{2} \text { with } \\
& \mid \perp, \top \quad \rightarrow \top \\
& \mid \top, \perp \quad \rightarrow \top \\
& \mid \text { is_fsym }\left(t, f_{1}\right) \text {,is_fsym }\left(t, f_{2}\right) \rightarrow T \text { if } f_{1} \neq f_{2} \\
& \mid-, \quad \rightarrow \perp
\end{aligned}
$$

Proof. Given a program $\pi$ and two expressions $e_{1}, e_{2} \in\langle$ expr $\rangle$, for all $\epsilon, \delta$, we want to prove that cond $1 \wedge$ cond $2 \rightarrow \epsilon_{1}\left(e_{1}\right) \downarrow_{R} \wedge \epsilon_{2}\left(e_{2}\right) \downarrow_{R}=\perp$ with $\epsilon_{1}=\Phi\left(\pi, \epsilon, \delta, e_{1}\right)$ and $\epsilon_{2}=$ $\Phi\left(\pi, \epsilon, \delta, e_{2}\right)$. We will detail every case for which cond2 is true (because when cond 2 is false, the proposition is trivially verified):

- case $1 e_{1}=\top, e_{2}=\perp$ : as $e_{1}, e_{2}$ are yet evaluated and reduced, $\epsilon_{1}\left(e_{1}\right) \downarrow_{R} \wedge \epsilon_{2}\left(e_{2}\right) \downarrow_{R}=$ $e_{1} \wedge e_{2}=\perp$
- case $2 e_{1}=\top, e_{2}=\perp$ : similar to case 1
- case $3 e_{1}=$ is_fsym $\left(t, f_{1}\right), e_{2}=$ is_fsym $\left(t, f_{2}\right)$ with $f_{1} \neq f_{2}$ : because of cond1, $\epsilon_{1}(t)=$ $\epsilon_{2}(t)=u$ then $\epsilon_{1}\left(e_{1}\right) \downarrow_{R}=\operatorname{Symb}(u)=f_{1}$ and $\epsilon_{2}\left(e_{2}\right) \downarrow_{R}=\operatorname{Symb}(u)=f_{2}$. Since every ground term has an unique head symbol and $f_{1} \neq f_{2}, \epsilon_{1}\left(e_{1}\right) \downarrow_{R} \wedge \epsilon_{2}\left(e_{2}\right) \downarrow_{R}=\perp$.


## 3 Optimizations

An optimization is a transformation which reduces the size of code (space optimization) or the time of execution (time optimization). In the case of PIL, the presented optimizations reduce the number of assignments (let) and tests (if) that are executed at run time. When manipulating abstract syntax trees, an optimization can be easily be described by a rewriting system. Its application consists in rewriting an instruction into an equivalent one, using a conditional rewrite rule of the form $i_{1} \rightarrow i_{2}$ IF $c$.

Definition 8. An optimization rule $i_{1} \rightarrow i_{2}$ IF $c$ rewrites a well-formed program $\pi$ into a program $\pi^{\prime}$ if there exists a position $\omega$ and a substitution $\sigma$ such that $\sigma\left(i_{1}\right)=\pi_{\mid \omega}$, $\pi^{\prime}=\pi\left[\sigma\left(i_{2}\right)\right]_{\omega}$ and $\sigma(c)$ is verified. If $c=e_{1} \sim e_{2}$ (resp. $c=e_{1} \perp e_{2}$ ), we say that $\sigma(c)$ is verified when $\sigma\left(e_{1}\right) \sim_{\pi_{\mid \omega}} \sigma\left(e_{2}\right)$ (resp. $\sigma\left(e_{1}\right) \perp_{\pi_{\mid \omega}} \sigma\left(e_{2}\right)$ ).

Since the notion of equivalence or incompatibility is dependent on the context of a program. In a rule, the context is given by the left part of the rule.

### 3.1 Reducing the number of assignments

This kind of optimization is standard, but useful to eliminate useless assignments. In the context of pattern matching, this improves the construction of substitutions, when a variable from the left-hand side is not used in the right-hand side for example.

Constant propagation. This first optimization removes the assignment of a variable defined as a constant. Since no side-effect can occur in a PIL program, it is possible to replace all occurrences of the variable by the constant (written $i[v / t]$ ).

$$
\text { ConstProp: } \operatorname{let}(v, t, i) \rightarrow i[v / t] \text { IF } t \in \mathcal{T}\left(\mathcal{F}_{0}\right)
$$

Dead variable elimination. Using a simple static analysis, this optimization eliminates useless assignments:

$$
\text { DeadVarElim: let }(v, t, i) \rightarrow i \text { IF use }(v, i)=0
$$

Definition 9. use is a function, defined recursively, which computes how many times the value of a variable $v$ may be used to evaluate a term, an expression or an instruction $t$ :

```
use \((v, t)=\) match \(t\) with
    \(|v \in \mathcal{X} \quad \rightarrow 1 \quad|\) eq \(\left(t_{1}, t_{2}\right) \quad \rightarrow \quad\) use \(\left(v, t_{1}\right)+\operatorname{use}\left(v, t_{2}\right)\)
    \(\mid v^{\prime} \in \mathcal{X} \quad \rightarrow 0\) if \(v \neq v^{\prime} \mid t \in \mathcal{T}(\mathcal{F}) \quad \rightarrow \quad 0\)
    \(|\operatorname{nop} \quad \rightarrow 0 \quad| i_{1} ; i_{2} \quad \rightarrow \quad\) use \(\left(v, i_{1}\right)+\operatorname{use}\left(v, i_{2}\right)\)
    \(\left|\operatorname{is\_ fsym}\left(t,{ }_{-}\right) \quad \rightarrow \operatorname{use}(v, t) \quad\right| \operatorname{let}\left({ }_{-}, t, i\right) \quad \rightarrow \quad\) use \((v, t)+\operatorname{use}(v, i)\)
    \(\left|\operatorname{subterm}_{f}\left(t,,_{-}\right) \rightarrow \operatorname{use}(v, t) \quad\right|\) hostcode(list) \(\quad \rightarrow \quad \operatorname{count}(\) list, \(v)\)
    \(\mid \operatorname{if}\left(e, i_{1}, i_{2}\right) \rightarrow \operatorname{use}(v, e)+\max \left(\operatorname{use}\left(v, i_{1}\right)\right.\), use \(\left.\left(v, i_{2}\right)\right)\)
```

where count $(l, v)$ counts $v$ apparitions in list. Note that the value computed by use is an upper bound.

Inlining. Since no side-effect can occur in a PIL program, an assigned variable cannot be modified. Therefore, when a variable is used only once, we can replace its occurrence by the corresponding term:

$$
\text { Inlining: let }(v, t, i) \rightarrow i[v / t] \text { IF use }(v, i)=1
$$

Fusion. The following rule merges two successive let which assign a same value to two different variables. This kind of optimization rarely applies on human written code, but in the context of pattern matching compilation, rules are compiled independently from others: the variables corresponding to the subject is defined for each rule. By merging the bodies, this allows to recursively perform some optimizations on subterms.

$$
\text { LetFusion: } \operatorname{let}\left(v_{1}, t, i_{1}\right) ; \operatorname{let}\left(v_{2}, t, i_{2}\right) \rightarrow \operatorname{let}\left(v_{1}, t, i_{1} ; i_{2}\left[v_{2} / v_{1}\right]\right)
$$

Note that the assignments must be the same for the two variables, to be sure that the values at run time are the same. We also suppose that use $\left(v_{1}, i_{2}\right)=0$. Otherwise, it would require to replace $v_{1}$ by a fresh variable in $i_{2}$.

### 3.2 Reducing the number of tests

The key technique to optimize pattern matching consists in merging branches, and thus tests that correspond to patterns with identical prefix. Usually, the discrimination between branches is performed by a switch instruction. In Tom, since the data-structure is not fixed, we cannot assume that a symbol is represented by an integer, an thus, contrary to standard approaches, we have to use an if statement instead. This restriction prevents us from selecting a branch in constant time. The two following rules define the fusion and the interleaving of conditional blocks.

Fusion. The fusion of two conditional adjacent blocks reduces the number of tests. This fusion is possible only when the two conditions are $\pi$-equivalent, This notion of $\pi$-equivalence means that the evaluation of the two conditions in a given program are the same (see Definition (6):

$$
\text { IfFusion: if }\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right) \rightarrow \operatorname{if}\left(c_{1}, i_{1} ; i_{2}, i_{1}^{\prime} ; i_{2}^{\prime}\right) \operatorname{IF} c_{1} \sim c_{2}
$$

To evaluate $c_{1} \sim c_{2}$ (i.e. $c_{1} \sim_{\pi} c_{2}$ with $\pi$ the redex of the rule), we use Proposition 3. The condition cond1 (i.e. $\Phi\left(\pi, \epsilon, \delta, e_{1}\right)=\Phi\left(\pi, \epsilon, \delta, e_{2}\right)$ ) is trivially verified because the semantics of the sequence instruction preserves the environment $\left(\forall \delta, \epsilon, \Phi\left(\pi, \epsilon, \delta, i_{1} ; i_{2}\right)=\right.$ $\left.\Phi\left(\pi, \epsilon, \delta, i_{1}\right)=\Phi\left(\pi, \epsilon, \delta, i_{2}\right)\right)$ and then $\forall \delta, \epsilon, \Phi\left(\pi, \epsilon, \delta, \sigma\left(c_{1}\right)\right)=\Phi\left(\pi, \epsilon, \delta, \sigma\left(c_{2}\right)\right)$. We just have to verify cond2 (i.e. $e_{1}=e_{2}$ ) which is easier.

Interleaving. As matching code consists of a sequence of conditional blocks, we would like to optimize blocks with $\pi$-incompatible conditions (see Definition 7). Some parts of the code cannot be executed at the same time, so swapping statically their order does not change the program behavior.

As we want to keep only one of the conditional block, we determine what instructions must be executed in case of success or failure of the condition and we obtain the two following transformation rules:

$$
\begin{aligned}
& \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right) \rightarrow \operatorname{if}\left(c_{1}, i_{1} ; i_{2}^{\prime}, i_{1}^{\prime} ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right)\right) \text { IF } c_{1} \perp c_{2} \\
& \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right) \rightarrow \operatorname{if}\left(c_{2}, i_{1}^{\prime} ; i_{2}, \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right), i_{2}^{\prime}\right) \text { IF } c_{1} \perp c_{2}
\end{aligned}
$$

As for the equivalence in the IfFusion rule, to evaluate $c_{1} \perp c_{2}$, we just have to verify cond2 in Proposition 4 (i.e. incompatible $\left(e_{1}, e_{2}\right)$ ). A drawback of these two rules is that some code is duplicated ( $i_{2}^{\prime}$ in the first rule, and $i_{1}^{\prime}$ in the second one). As we want to maintain linear the size of the code, we consider specialized instances of these rules with respectively $i_{2}^{\prime}$ and $i_{1}^{\prime}$ equal to nop, and eliminate the extra nop instructions. This gives us the following rules:

```
If Interleaving: \(\operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right) \rightarrow \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime} ; \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right)\right) \operatorname{IF} c_{1} \perp c_{2}\)
    if \(\left(c_{1}, i_{1}\right.\), nop \() ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right) \rightarrow \operatorname{if}\left(c_{2}, i_{2}, \operatorname{if}\left(c_{1}, i_{1}\right.\right.\), nop \(\left.) ; i_{2}^{\prime}\right)\) IF \(c_{1} \perp c_{2}\)
```

These two rules reduce the number of tests at run time because one of the tests is moved into the "else" branch of the other. The second rule can be instantiated and used to swap blocks. When $i_{1}^{\prime}$ and $i_{2}^{\prime}$ are reduced to the instruction nop, the second rule can be simplified into:

$$
\operatorname{if}\left(c_{1}, i_{1}, \operatorname{nop}\right) ; \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right) \rightarrow \operatorname{if}\left(c 2, i_{2}, \operatorname{if}\left(c_{1}, i_{1}, \text { nop }\right)\right) \text { IF } c_{1} \perp c_{2}
$$

As the two conditions are $\pi$-incompatible, we have the following equivalence:

$$
\operatorname{if}\left(c_{2}, i_{2}, \operatorname{if}\left(c_{1}, i_{1}, \operatorname{nop}\right)\right) \equiv \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right) ; \operatorname{if}\left(c_{1}, i_{1}, \text { nop }\right) \text { IF } c_{1} \perp c_{2}
$$

After all, we obtain the following rule corresponding to the swapping of two conditional adjacent blocks. This rule does not optimize the number of tests but is useful to join blocks subject to be merged thanks to a smart strategy.

$$
\text { IfSwapping: if }\left(c_{1}, i_{1}, \text { nop }\right) ; \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right) \rightarrow \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right) ; \text { if }\left(c_{1}, i_{1}, \text { nop }\right) \text { IF } c_{1} \perp c_{2}
$$

| DeadVarElim | $\operatorname{let}(v, t, i)$ | $\rightarrow i$ | IF use $(v, i)=0$ |
| :--- | :--- | :--- | :--- |
| ConstProp | $\operatorname{let}(v, t, i)$ | $\rightarrow i[v / t]$ | IF $t \in \mathcal{T}\left(\mathcal{F}_{0}\right)$ |
| Inlining | $\operatorname{let}(v, t, i)$ | $\rightarrow i[v / t]$ | IF use $(v, i)=1$ |
| LetFusion | $\operatorname{let}\left(v_{1}, t, i_{1}\right) ; \operatorname{let}\left(v_{2}, t, i_{2}\right)$ | $\rightarrow \operatorname{let}\left(v_{1}, t, i_{1} ; i_{2}\left[v_{2} / v_{1}\right]\right)$ |  |
| IfFusion | $\operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right)$ | $\rightarrow \operatorname{if}\left(c_{1}, i_{1} ; i_{2}, i_{1}^{\prime} ; i_{2}^{\prime}\right)$ | IF $c_{1} \sim c_{2}$ |
| IfInterleaving $\operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}\right.$, nop $)$ | $\rightarrow \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime} ; \operatorname{if}\left(c_{2}, i_{2}\right.\right.$, nop $\left.)\right)$ | IF $c_{1} \perp c_{2}$ |  |
| IfSwapping | if $\left(c_{1}, i_{1}\right.$, nop $) ;$ if $\left(c_{2}, i_{2}\right.$, nop $) \rightarrow \operatorname{if}\left(c_{2}, i_{2}\right.$, nop $) ;$ if $\left(c_{1}, i_{1}\right.$, nop $)$ IF $c_{1} \perp c_{2}$ |  |  |

Fig. 5. System of optimization rules: OptSys

### 3.3 Application strategy

Given the rules presented in Section 3.1 and Section 3.2, Figure 5 defines the rewrite systems called OptSys.

Without strategy, this system is clearly not confluent and not terminating. For example, without condition, the IfSwapping rule can be applied indefinitely because of the symmetry of incompatibility. The confluence of the system is not necessary as long as the programs obtained are semantically equivalent to the source program but the termination is an essential criterion. Moreover, the strategy should apply the rules to obtain a program, as efficient as possible. Let us consider again the program given in Figure 3, and let us suppose that we interleave the two last patterns. This would result in the following suboptimal program:

```
if(is_fsym(t,f), let(t (t, subterm
if(is_fsym(t,g), let(t t , subtermg}(t,1), let(x, t t , hostcode(x)))
    if(is_fsym(t,f), let(t t , subterm
```

IfSwapping and IfFusion rules can longer be applied to share the is_fsym $(t, f)$ tests. This order of application is not optimal. It is intuitive that the interleaving rule must be applied at the end, when no more optimization or fusion is possible.

The second matter is to ensure the termination. The IfSwapping rule is the only rule that does not decrease the size or the number of assignments of a program. To limit its application for interesting cases, we define a condition which ensures that a swapping is performed only if it enables a fusion. This condition can be implemented in two ways, either in using a context, or in defining a total order on conditions. Given a sequence of if instructions, the first approach consists in swapping two of them when the right one can be fusioned with an if which earlier in the sequence. This approach is not efficient since we have to find two elements in a list that can react via the IfFusion rule. The second approach is more efficient since it consists in considering a total order on conditions and perform the swapping only if two successive if have some conditions $c 1$ and $c_{2}$ such that $c_{1} \perp c_{2}$ and $c_{1}<c_{2}$. Similarly to a swap-sort algorithm, this approach ensures the termination of the algorithm. In this way, we obtain a new IfSwapping rule:

$$
\operatorname{if}\left(c_{1}, i_{1}, \operatorname{nop}\right) ; \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right) \rightarrow \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right) ; \operatorname{if}\left(c_{1}, i_{1}, \text { nop }\right) \text { IF } c_{1} \perp c_{2} \wedge c_{1}<c_{2}
$$

Using basic strategy operators such as Innermost(s) (which applies $s$ as many times as possible, starting from the leaves), $s 1 \mid s 2$ (which applies $s 1$ or $s 2$ indifferently), repeat $(s)$ (which applies $s$ as many times as possible, returning the last unfailing result), and $r 1 ; r 2$ (which applies $s 1$, and then $s 2$ if $s 1$ did not failed), we can easily define a strategy which describes how the rewrite system OptSys should be applied to normalize a PIL program:

```
Innermost( repeat(ConstProp | DeadVarElim | Inlining| LetFusion| IfFusion| IfSwapping);
    repeat(IfInterleaving))
```

Starting from the program given in Figure 3, we can apply the rule IfSwapping, followed by a step of IfFusion, and we obtain:

```
if(is_fsym(t,f), let(t}\mp@subsup{t}{1}{},\mp@subsup{\operatorname{subterm}}{f}{(t,1),if(is_fsym(t 
    ; let(t
if(is_fsym(t,g), let(th},\mp@subsup{\operatorname{subterm}}{g}{(t,1),let(x, t 
```

Then, we can apply a step of Inlining to remove the second instance of $t_{1}$, a step of LetFusion, and a step of Interleaving (is_fsym $\left(t_{1}, a\right)$ and is_fsym $\left(t_{1}, b\right)$ are $\pi$ incompatible). This results in the following program:

```
if \(\left(\right.\) is_fsym \((t, f), \operatorname{let}\left(\mathrm{t}_{1}, \operatorname{subterm}_{\mathrm{f}}(\mathrm{t}, 1)\right.\),
    if(is_fsym( \(\left.t_{1}, a\right)\), hostcode(), if(is_fsym( \(\left.t_{1}, b\right)\), hostcode(), nop))), nop) ;
if(is_fsym \((t, g), \operatorname{let}\left(x, \operatorname{subterm}_{g}(\mathrm{t}, 1)\right.\), hostcode(x)), nop)
```

Since is_fsym $(t, f)$ and is_fsym $(t, g)$ are $\pi$-incompatible, we can apply a step of If Interleaving, and get the irreducible following program:

```
if(is_fsym \((t, f)\),
    let \(\left(t_{1}, \operatorname{subterm}_{f}(t, 1)\right.\), if(is_fsym( \(\left.t_{1}, a\right)\), hostcode(), if(is_fsym( \(\left.t_{1}, b\right)\), hostcode(), nop))),
    if \(\left(\right.\) is_fsym \((t, g), \operatorname{let}\left(x, \operatorname{subterm}_{g}(\mathrm{t}, 1)\right.\), hostcode \(\left.(\mathrm{x})\right)\), nop)
```


## 4 Properties

When performing optimization by program transformation, it is important to ensure that the generated code has some expected properties. The use of formal methods to describe our optimization algorithm allows us to give proofs. In the section we show that each transformation rule is correct, in the sense that the the optimized program has the same observational behavior as the original. We also show that the optimized code is both more efficient, and smaller than the initial program.

### 4.1 Correction

Definition 10. A transformation rule $r$ is correct if for all well-formed program $\pi$, $r$ rewrites $\pi$ in $\pi^{\prime}$ (Definition 8) implies that $\pi \sim \pi^{\prime}$ (Definition 5).

From this definition, we prove that every rule given in Section 3 is correct. For that, two conditions have to be verified:

1. $\pi^{\prime}$ is well-formed,
2. $\forall \epsilon, \delta$, the derivations of $\pi$ and $\pi^{\prime}$ lead to the same result $\delta^{\prime}$.

The first condition is quite easy to verify. The second one is more interesting: we consider a program $\pi$, a rule $l \rightarrow r$ IF $c$, a position $\omega$, and a substitution $\sigma$ such that $\sigma(l)=\pi_{\mid \omega}$. We have $\pi^{\prime}=\pi[\sigma(r)]_{\omega}$. We have to compare the derivations of $\pi$ and $\pi^{\prime}$ in a the context $\epsilon, \delta$.

- when $\omega=\wedge$, we have to compare the derivation tree of $\pi=\sigma(l)$ and $\pi^{\prime}=\sigma(r)$,
- otherwise, we consider the derivation of $\pi$ (resp. $\pi^{\prime}$ ): there is a step which needs in premise the derivation of $\pi_{\mid \omega}$ (resp. $\left.\pi[\sigma(r)]_{\omega}\right)$. This is the only difference between the two trees.

In both cases, we have to verify that $\pi_{\mid \omega}=\sigma(r)$ and $\sigma(l)$ have the same derivation in a context which is either equal to $\epsilon, \delta$ if $\omega=\wedge$ or given by the including instruction rule in the derivation tree of $\pi$ (or $\pi^{\prime}$ ). To simplify the proof, we consider $l, r$ and $c$ instead of $\sigma(l), \sigma(r)$ and $\sigma(c)$.

Correction of Inlining. In the Inlining rule, the condition use $(v, i)=1$ is not a required condition for correction. We have to prove that $l=\operatorname{let}(v, u, i) \sim r=i[v / u]$. To prove that $\pi \sim \pi^{\prime}$, we have to verify that for a given $\epsilon, \delta, l$ and $r$ have the same derivation.

The derivation of $l$ consists in applying the rule (let) :

$$
\frac{\langle\epsilon[v \leftarrow t], \delta, i\rangle \mapsto_{b s} \delta^{\prime} \quad \epsilon(u) \rightarrow_{R *} t}{\langle\epsilon, \delta, \operatorname{let}(v, t, i)\rangle \mapsto_{b s} \delta^{\prime}}(l e t)
$$

We consider that the derivation of $r$ is a given $\delta^{\prime \prime}$ :

$$
\langle\epsilon, \delta, i[v / u]\rangle \mapsto_{b s} \delta^{\prime \prime}
$$

We have to prove that $\delta^{\prime}=\delta^{\prime \prime}$. As these programs are well-formed, in $i$, there is no new declaration for variables contained in $u$ or for $v$ therefore $\Phi(i, \epsilon[v \leftarrow t], \delta, v)(v)=$ $\epsilon[v \leftarrow t](v) \rightarrow_{R *} t$. Furthermore, $\Phi(i, \epsilon, \delta, u)(u)=\epsilon(u) \rightarrow_{R *} t$. We can conclude that every evaluation of $v$ or $u$ in $i$ are equal to $t$ thus replacing $v$ by $u$ in $i$ doesn't change the program behavior, $\delta^{\prime}=\delta^{\prime \prime}$.

Correction of ConstantProp. The proof is very similar to the correction of Inlining but as we have the condition $u \in \mathcal{T}\left(\mathcal{F}_{0}\right)$, we know that $\forall \epsilon, \epsilon(u) \rightarrow_{R *} u$ thus we can directly deduce that every evaluation of $v$ or $u$ in $i$ are equal to $u$.

Correction of DeadVarElim. In the DeadVarElim rule, use $(v, i)=0$ i.e. $v$ is never evaluated in $i$ therefore $\langle\epsilon[v \leftarrow t], \delta, i\rangle$ and $\langle\epsilon, \delta, i\rangle$ have the same reduction $\delta^{\prime}$.

Correction of LetFusion. In this rule, $l=\operatorname{let}\left(v_{1}, t, i_{1}\right) ; \operatorname{let}\left(v_{2}, t, i_{2}\right)$ and $r=$ let $\left(v_{1}, t, i_{1} ; i_{2}\left[v_{2} / v_{1}\right]\right)$. There is no condition $c$. Similarly to the correction of IfSwapping, we construct the derivation of $l$ :

$$
\frac{\frac{\overline{\left\langle\epsilon\left[v_{1} \leftarrow t\right], \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}}}{\left\langle\epsilon, \delta, \operatorname{let}\left(v_{1}, t, i_{1}\right)\right\rangle \mapsto_{b s} \delta^{\prime}}(l e t) \frac{\overline{\left\langle\epsilon\left[v_{2} \leftarrow t\right], \delta^{\prime}, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}}}{\left\langle\epsilon, \delta, \operatorname{let}\left(v_{1}, t, i_{1}\right) ; \operatorname{let}\left(v_{2}, t, i_{2}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}}(\text { let })}{\left\langle\epsilon, \delta^{\prime}, \operatorname{let}\left(v_{2}, t, i_{2}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}}(\text { seq })
$$

Then, we derive the right part $r$ :

$$
\frac{\overline{\left\langle\epsilon\left[v_{1} \leftarrow t\right], \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}} \quad \overline{\left\langle\epsilon\left[v_{1} \leftarrow t\right], \delta^{\prime}, i_{2}\left[v_{2} / v_{1}\right]\right\rangle \mapsto_{b s} \delta^{\prime \prime}}}{\frac{\left\langle\epsilon\left[v_{1} \leftarrow t\right], \delta, i_{1} ; i_{2}\left[v_{2} / v_{1}\right]\right\rangle \mapsto_{b s} \delta^{\prime \prime}}{\left\langle\epsilon, \delta, \operatorname{let}\left(v_{1}, t, i_{1} ; i_{2}\left[v_{2} / v_{1}\right]\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}}(\text { let })}(\text { seq })
$$

Since $\operatorname{let}\left(v_{1}, t, i_{1}\right) ; \operatorname{let}\left(v_{2}, t, i_{2}\right)$ is in a well-formed program $\pi$, when evaluating $i_{2}$, $v_{2}$ is equal to $t$. Furthermore, we assume that no variable $v_{1}$ occurs in $i_{2}$. Therefore, if $\left\langle\epsilon, \delta^{\prime}, \operatorname{let}\left(v_{2}, t, i_{2}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}$ then we also have $\left\langle\epsilon\left[v_{1} \leftarrow t\right], \delta^{\prime}, i_{2}\left[v_{2} / v_{1}\right]\right\rangle \mapsto_{b s} \delta^{\prime \prime}$. As a consequence, $\delta^{\prime \prime}$ is the same for the two derivations.

Correction of IfFusion. Since $c_{1}$ and $c_{2}$ are equivalent, $c_{1}$ and $c_{2}$ are either true or false. Thus, we have to distinguish two cases.

Case 1: $\epsilon\left(c_{1}\right) \rightarrow_{R *} \top$ and $\epsilon\left(c_{2}\right) \rightarrow_{R *} \top$
First, we calculate the derivation of $l$ :

$$
\frac{\overline{\left\langle\epsilon, \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}} \quad \epsilon\left(c_{1}\right) \rightarrow_{R *} \top}{\frac{\left\langle\epsilon, \delta, \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right)\right\rangle \mapsto_{b s} \delta^{\prime}}{}}(\text { iftrue }) \frac{\overline{\left\langle\epsilon, \delta^{\prime}, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}} \epsilon\left(c_{2}\right) \rightarrow_{R *} \top}{\left\langle\epsilon, \delta, \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, i_{2}^{\prime}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}} \text { (iftrue) }
$$

Then, we give the derivation of $r$ :

$$
\frac{\overline{\left\langle\epsilon, \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}} \overline{\left\langle\epsilon, \delta^{\prime}, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}}}{\frac{\left\langle\epsilon, \delta, i_{1} ; i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}}{\left\langle\epsilon, \delta, \text { if }\left(c_{1}, i_{1} ; i_{2}, i_{1}^{\prime} ; i_{2}^{\prime}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}} \quad \epsilon\left(c_{1}\right) \rightarrow_{R *} T}(\text { iftrue })
$$

Case 2: $\epsilon\left(c_{1}\right) \rightarrow_{R *} \perp$ and $\epsilon\left(c_{2}\right) \rightarrow_{R *} \perp$, the proof is similar.

Correction of IfSwapping. In this rule, $l=\operatorname{if}\left(c_{1}, i_{1}, \operatorname{nop}\right) ; \operatorname{if}\left(c_{2}, i_{2}\right.$, nop $)$ and $r=$ if $\left(c_{2}, i_{2}\right.$, nop $)$; if ( $c_{1}, i_{1}$, nop). Since $c_{1}$ and $c_{2}$ are $\pi$-incompatible, three cases have to be studied:

Case 1: $\epsilon\left(c_{1}\right) \rightarrow_{R *} \top$ and $\epsilon\left(c_{2}\right) \rightarrow_{R *} \perp$

We now consider the program $\operatorname{if}\left(c_{2}, i_{2}\right.$, nop $) ; \operatorname{if}\left(c_{1}, i_{1}\right.$, nop $)$. Starting from the same environment $\epsilon$ and $\delta$, we show that the derivation leads to the same state $\delta^{\prime}$, and thus prove that if ( $c_{1}, i_{1}$, nop $)$;if ( $c_{2}, i_{2}$, nop) and $\operatorname{if}\left(c_{2}, s_{2}\right.$, nop); if ( $c_{1}, s_{1}$, nop) are equivalent:

$$
\frac{\overline{\langle\epsilon, \delta, \text { nop }\rangle \mapsto_{b s} \delta}(\text { nop }) \epsilon\left(c_{2}\right) \rightarrow_{R *} \perp}{\left.\frac{\left\langle\epsilon, \delta, \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right)\right\rangle \mapsto_{b s} \delta}{\left\langle\epsilon, \delta, \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right) ; \operatorname{if}\left(c_{1}, i_{1}, \text { nop }\right)\right\rangle \mapsto_{b s} \delta^{\prime}} \frac{\overline{\left\langle\epsilon, \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}}}{\left\langle\epsilon\left(c_{1}\right) \rightarrow_{R *} T\right.} \text { (iftrue) }\right) \text { if }\left(c_{1}, i_{1}, \text { nop }\right) \mapsto_{b s} \delta^{\prime}}(\text { seq })
$$

Since $\pi$ and $\pi^{\prime}$ are well-formed, their derivation in a given context are unique (Proposition(2). $\left\langle\epsilon, \delta, i_{1}\right\rangle \mapsto_{b s} \delta^{\prime}$ is part of these derivation trees, so it is unique, and $\delta^{\prime}$ is identical in both derivations.

Case 2: $\epsilon\left(c_{1}\right) \rightarrow_{R *} \perp$ and $\epsilon\left(c_{2}\right) \rightarrow_{R *} \top$, the proof is similar.
Case 3: $\epsilon\left(c_{1}\right) \rightarrow_{R *} \perp$ and $\epsilon\left(c_{2}\right) \rightarrow_{R *} \perp$, the proof is similar.

Correction of IfInterleaving. As for the previous rule, since $c_{1}$ and $c_{2}$ are $\pi$ incompatible, we have to consider three cases. We present the second one which is the most interesting.

Case 2: $\epsilon\left(c_{1}\right) \rightarrow_{R *} \perp$ and $\epsilon\left(c_{2}\right) \rightarrow_{R *} \top$

$$
\begin{aligned}
& \frac{\overline{\left\langle\epsilon, \delta, i_{1}^{\prime}\right\rangle \mapsto_{b s} \delta^{\prime}} \quad \epsilon\left(c_{1}\right) \rightarrow_{R *} \perp}{\frac{\left\langle\epsilon, \delta, \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right)\right\rangle \mapsto_{b} \delta^{\prime}}{}(\text { iffalse }) \frac{\overline{\left\langle\epsilon, \delta^{\prime}, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}} \quad \epsilon\left(c_{2}\right) \rightarrow_{R *} T}{\left\langle\epsilon, \delta^{\prime}, \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}}(\text { iftrue })} \begin{array}{l}
\left\langle\epsilon, \delta, \operatorname{if}\left(c_{1}, i_{1}, i_{1}^{\prime}\right) ; \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}
\end{array} \\
& \frac{\overline{\left\langle\epsilon, \delta, i_{1}^{\prime}\right\rangle \mapsto_{b s} \delta^{\prime}} \frac{\overline{\left\langle\epsilon, \delta^{\prime}, i_{2}\right\rangle \mapsto_{b s} \delta^{\prime \prime}}}{\left\langle\epsilon, \delta^{\prime}, \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right)\right\rangle \rightarrow_{R *} \top} \text { (iftrue) }}{\frac{\left\langle\epsilon, \delta, i_{1}^{\prime} ; \operatorname{if}\left(c_{2}, i_{2}, \operatorname{nop}\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}}{\left\langle\epsilon, \delta, \operatorname{if}\left(c_{1}, i_{1} ; i_{1}^{\prime}, i_{1}^{\prime} ; \operatorname{if}\left(c_{2}, i_{2}, \text { nop }\right)\right)\right\rangle \mapsto_{b s} \delta^{\prime \prime}} \quad \epsilon\left(c_{1}\right) \rightarrow_{R *} \perp} \text { (iffalse) }
\end{aligned}
$$

The proof of Case 1 and Case 3 are similar.

### 4.2 Time and space reduction

In this section we show that the optimized code is both more efficient, and smaller than the initial program. For that, we consider two measures:

- the size of a program $\pi$ is the number of instructions which constitute the program,
- the efficiency of a program $\pi$ is determined by the number of tests and assignments which are performed at run-time.

It is quite easy to verify that each transformation rule does not increase the size of the program: DeadVarElim, ConstProp, Inlining, and LetFusion decrease the size of a program, whereas IfFusion, IfInterleaving and IfSwapping maintain the size of the transformed program.

It is also clear that no transformation can reduce the efficiency of a given program:

- each application of DeadVarElim, ConstProp, and Inlining reduces by one the number of assignment that can be performed at run time,
- IfFusion reduces by one the number of tests,
- If Interleaving also decreases the number of tests when the first alternative is chosen. Otherwise, there is no optimization,
- IfSwapping does not modify the efficiency of a program.

The program transformation presented in Section 3 is an optimization which improves the efficiency of a given program, without increasing its size. Similarly to [8, this result is interesting since it allows to generate efficient pattern matching implementations whose size is linear in the number and size of patterns.

## 5 Experimental Results

The presented algorithm (rules and strategy) has been implemented and integrated into Tom (see Figure 1). We have selected several representative programs and measured the effect of optimization in practice:

|  | Fibonnacci | Eratosthene | Gomoku | Langton | Nspk | Structure |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| without optimization | $21.3 s$ | $174.0 s$ | $70.0 s$ | $15.7 s$ | $1.7 s$ | $12.3 s$ |
| with optimization | $20.0 s$ | $2.8 s$ | $30.4 s$ | $1.4 s$ | $1.2 s$ | $11.3 s$ |

Fibonacci computes several times the $18^{\text {th }}$ Fibonacci number. Eratosthene computes primes numbers up to 1000 , using associative list matching. The improvement comes from the Inlining rules which avoids computing a substitution unless the rule applies (i.e. the conditions are verified). Gomoku looks for five pawn on a go board, using list matching. This example contains more than 40 patterns and illustrates the interest of test-sharing. Langton is a program which computes the $1000^{t h}$ iteration of a cellular automaton, using pattern matching to implement the transition function. This example contains more than 100 (ground) patterns. The optimized program is optimal in the sense that a pair
(position,symbol) is never tested more than once. Nspk implements the verification of the Needham-Schroeder Public-Key Protocol. Structure is a prover for the Calculus of Structures where the inference is performed by pattern matching and rewriting. All these examples are available on the Tom web page ${ }^{[1]}$. Some of these examples have been implemented in OCaml. Fibonacci, Eratosthene, and Langton are respectively executed in 4.0, 0.7 and 2.3 seconds.

These benchmarks show that the proposed approach is effective in practice and allows Tom to become competitive with state of the art implementations such as OCaml. We should remind that Tom is not dedicated to a unique language. In particular, the fact that data-structure can be user-defined prevents us from using the switch instruction. Moreover, all the presented examples have been implemented in Java and executed on a PowerMac 2 GHz .

## 6 Conclusion and future works

In this paper, we have presented a new approach to compile efficiently pattern-matching. This method is based on well-attested program optimization methods. Separating compilation and optimization in order to keep modularity and to facilitate extensions is longestablished in compilation community. Furthermore, using program transformation to realize optimization is an elegant way as the correction proof is facilitated contrary to complex algorithms.

As yet, the algorithms proposed to compile pattern-matching [5|9|2|8 were adapted to very much specific forms of matching and assumed properties on target language and term representation. These algorithms hardly conform to evolutions, such as matching modulo AC or non-linear patterns. All these restrictions lead us to give an other approach for Tom compiler by compiling simply pattern-matching and optimizing separately. The set of optimization rules presented in this paper is not specific to pattern-matching and can be applied to any program. However, some rules such as IfSwapping reply to specific problems related to pattern-matching compilation and formalizing them is an other contribution to this paper. The section 4 shows how simply the correction of the system can be proved and justifies the choice of a method based on program transformations. Moreover, as the Tom compiler is bootstrapped, the optimization rules have been implemented in a very natural way using pattern-matching so this work demonstrates how Tom tool is well adapted for transforming programs and more generally for compiler construction.

This paper shows that using program transformation rules to optimize pattern-matching is an efficient solution, in respect to algorithms based on automata. The implementation of this work conjugated with the validation tool of Tom [10] give us an efficient adaptative certified compiler of pattern-matching.

## References

1. Andrew W. Appel and David B. MacQueen. Standard ML of new jersey. In J. Maluszyński and M. Wirsing, editors, Proceedings of the Third International Symposium on Programming
[^0]Language Implementation and Logic Programming, number 528, pages 1-13. Springer Verlag, 1991.
2. Lennart Augustsson. Compiling pattern matching. In Proceedings of a conference on Functional Programming Languages and Computer Architecture, pages 368-381, 1985.
3. Franz Baader and Tobias Nipkow. Term Rewriting and all That. Cambridge University Press, 1998.
4. Peter Borovanský, Claude Kirchner Hélène Kirchner, Pierre-Etienne Moreau, and Christophe Ringeissen. An overview of ELAN. In WRLA'98: Proceedings of the 2nd International Workshop on Rewriting Logic and its Applications, volume 15, Pont-à-Mousson (France), 1998. Electronic Notes in Theoretical Computer Science.
5. Luca Cardelli. Compiling a functional language. In LFP'84: Proceedings of the 1984 ACM Symposium on LISP and functional programming, pages 208-217, 1984.
6. Manuel Clavel, Steven Eker, Patrick Lincoln, and José Meseguer. Principles of Maude. In Proceedings of the first international workshop on rewriting logic, volume 4, Asilomar (California), 1996. Electronic Notes in Theoretical Computer Science.
7. A. van Deursen, J. Heering, and P. Klint, editors. Language Prototyping: An Algebraic Specification Approach, volume 5 of AMAST Series in Computing. World Scientific Publishing Co., 1996.
8. Fabrice Le Fessant and Luc Maranget. Optimizing pattern matching. In ICFP'01: Proceedings of the sixth ACM SIGPLAN International Conference on Functional Programming, pages 2637. ACM Press, 2001.
9. Albert Gräf. Left-to-right tree pattern matching. In RTA'91: Proceedings of the 4th international conference on Rewriting Techniques and Applications, volume 488 of LNCS, pages 323-334. SV, 1991.
10. Claude Kirchner, Pierre-Etienne Moreau, and Antoine Reilles. Formal validation of pattern matching code. In Pedro Barahone and Amy Felty, editors, Proceedings of the 7th ACM SIGPLAN international conference on Principles and practice of declarative programming, pages 187-197. ACM, July 2005.
11. David Lacey, Neil Jones, Eric Van Wyk, and Carl Christian Frederikson. Proving correctness of compiler optimizations by temporal logic. Higher-Order and Symbolic Computation, 17(2), 2004.
12. Pierre-Etienne Moreau, Christophe Ringeissen, and Marian Vittek. A Pattern Matching Compiler for Multiple Target Languages. In G. Hedin, editor, 12th Conference on Compiler Construction, Warsaw (Poland), volume 2622 of LNCS, pages 61-76. Springer-Verlag, May 2003.
13. R. C. Sekar, R. Ramesh, and I. V. Ramakrishnan. Adaptive pattern matching. SIAM Journal on Computing, 24(6):1207-1234, 1995.
14. Mark G. J. van den Brand, J. Heering, P. Klint, and P. A. Olivier. Compiling language definitions: the asf+sdf compiler. ACM Trans. Program. Lang. Syst., 24(4):334-368, 2002.
15. Eelco Visser, Zine el Abidine Benaissa, and Andrew Tolmach. Building program optimizers with rewriting strategies. In ICFP'98: Proceedings of the third ACM SIGPLAN International Conference on Functional Programming, pages 13-26, 1998.


[^0]:    ${ }^{1}$ http://tom.loria.fr

