Two Techniques for Compiling Lazy Pattern Matching

LUC MARANGET
INRIA Rocquencort

In ML style pattern matching, pattern size is not constrained and ambiguous patterns are allowed. This generality leads to a clear and concise programming style but is challenging in the context of lazy evaluation. A first challenge concerns language designers: in a lazy functional language such as lazy ML or Haskell, the evaluation order of expressions follows actual data dependencies. That is, only the computations that are needed to produce the final result are performed. Once given a proper, non-ambiguous, semantics, pattern matching should be compiled in a similar spirit: any value matching a given pattern should be recognized by performing only the minimal number of elementary tests needed to do so. This challenge was first met by A. Laville.

A second challenge concerns compiler designers. As it stands, Laville's compilation algorithm cannot be incorporated in an actual lazy ML compiler for efficiency and completeness reasons. As a matter of fact, Laville's original algorithm did not fully treat the case of integers in patterns and can lead to explosions both in compilation time and generated code size. This paper provides a complete solution to that second challenge. In particular, the well-known (and size-efficient) pattern matching compilation technique using backtracking automata is here introduced for the first time into the world of lazy pattern matching.

Categories and Subject Descriptors: D.3.4 [**Programming Languages**]: Processor—*Compilers*; D.3.1 [**Programming Languages**]: Formal Definitions and Theory—*Semantics*

General terms: Lazy Functional Languages, Pattern Matching

1. INTRODUCTION

Pattern matching is a key feature of the ML language. It provides a way to discriminate between values of structured types and to access their subparts. Pattern matching enhances the clarity and readability of programs. Compare, for instance, the ML function computing the sum of an integers list with its Lisp counterpart (all examples are in CAML [Weis 1990; Leroy 1995] syntax).

```
let rec sum xs = match xs with
  []    → 0
| y::ys → y+sum ys
```

```
(defun sum (l)
  (if (consp l)
    (+ (car l) (sum (cdr l)))
    0))
```

In ML, patterns can be nested arbitrarily. This means that pattern matching has to be compiled into sequences of simple tests: a complicated pattern such as $((1, x), y::[])$ cannot be recognized by a single test. Usually, pattern matching compilers attempt to "factorize" tests as much as possible, to avoid testing several times the same position in a term.

A pattern matching expression does not specify the order in which tests are performed. When ML is given strict semantics, as in SML [Harper, Milner, and Tofte

Author's address: L. Maranget, INRIA Rocquencourt, BP 105, 78153 Le Chesnay Cedex, France. Luc.Maranget@inria.fr.

1991], all orders are correct, and choosing a particular order is only a matter of code size and run-time efficiency. When ML is given lazy semantics, as in LML [Augustsson 1985], all testing orders are not semantically equivalent. Consider for instance the ML definition:

```
let F x y  = match (x,y) with
  (true,true)   → 1
| (_,false)     → 2
| (false,true)  → 3
```

The patterns can be checked from left to right, as it is usually the case (function $\mathtt{F}_1$, below), or from right to left (function $\mathtt{F}_2$).

```
let F1 x y =
  if x then
    if y then 1 else 2
  else
    if y then 3 else 2
```

```
let  F2 x y =
  if y then
    if x then 1 else 3
  else
    2
```

When the variable `y` is bound to `false`, the test on `x` is useless. This can be avoided by testing `y` before `x`, as in $\mathtt{F}_2$.

Worse, consider the application `F` $\perp$ `false`, where $\perp$ is a non-terminating computation. In strict ML, function arguments are reduced before calling the functions, so that both compilations $\mathtt{F}_1$ $\perp$ `false` and $\mathtt{F}_2$ $\perp$ `false` do not terminate. In lazy ML, function arguments are not evaluated until their values are actually needed. Therefore, $\mathtt{F}_1$ will loop by trying to evaluate `x` $= \perp$, whereas $\mathtt{F}_2$ will give the answer 2. In the spirit of lazy evaluation, a result should be given whenever possible. Thus, a lazy compiler should compile the function `F` as $\mathtt{F}_2$, not as $\mathtt{F}_1$.

In the previous example, there is a good reason to claim that $\mathtt{F}_2$ is more "correct" than $\mathtt{F}_1$. Indeed, both functions output the same result when given the same input, except in one case, where $\mathtt{F}_2$ gives the result 2, whereas $\mathtt{F}_1$ does not terminate. Unfortunately, things are not always that simple and there are pattern matching definitions that cannot be compiled correctly. Consider, for instance, Berry's famous example:

```
let B x y z = match (x,y,z) with
  (true ,false,_)      → 1
| (false,_    ,true)   → 2
| (_    ,true ,false)  → 3
```

One easily checks that there does not exist a correct compilation of `B` and that one cannot order the possible compilation by their termination properties. By "possible" compilation, I mean here some automaton that will examine all or some of the variables `x`, `y` and `z`, one after the other. Let us consider for instance an automaton that examines `x` first. Because such an automaton will always loop by engaging in the non-terminating computation $\perp$, it cannot output the correct value 3 when given the input `x` $= \perp$, `y` $=$ `true` and `z` $=$ `false`, whereas other automata that start by examining `y` can. Similarly no automaton examining `y` first will ever produce the correct result 2 as a compilation of the function call `B false` $\perp$ `true`, whereas some of the automata examining `x` first will.

The task of a lazy ML compiler is thus twofold. When given a pattern matching definition, it must both determine whether this definition can be compiled correctly

or not and produce a correct compilation when possible. This problem has first been solved in the case of non-overlapping patterns by Huet and Lévy [1991]. Given a set of possibly overlapping patterns, Laville [1991] shows how to replace them, when possible, by an equivalent set of non-overlapping patterns, compiled using Huet and Lévy's technique. Laville's method is not complete, since it cannot treat the case of datatypes with infinite signatures (such as the type of integers). As a matter of fact, in presence of infinite signatures, Laville's equivalent set of non-overlapping patterns can also be infinite. Following a similar idea, Puel and Suárez [1993] translate the initial set of overlapping patterns into an equivalent set of "constrained" patterns, which are special patterns encoding the disambiguating rule of pattern matching on overlapping patterns. Then, they compile the pattern matching on the constrained patterns with an extension of Huet and Lévy's technique. Although the Suárez and Puel's approach does consider infinite signatures, its direct application in a compiler remains problematic, since the size of the constrained patterns can be exponential in the size of the initial ambiguous patterns. I experienced such a misbehavior while implementing Suárez and Puel technique, even on very ordinary pattern matching definition.

In this paper, I take a simpler approach: compilation operates directly on overlapping patterns. I first recall the semantics of lazy pattern matching, as given by Laville [1991]. I then present two direct compilation techniques that preserve this semantics whenever possible. These technique are presented as source-to-source transformations. They differ by the nature of the target automata they produce. The first technique produces *tree-like automata*. The main advantage of tree-like automata is simplicity, so that the associated compilation scheme is straightforward and its correctness proof is easy. Unfortunately, the price paid for this simplicity is a potential explosion in size of the output automata. The second and more sophisticated technique produces *automata with failures*. These automata possess a static exception construct that enables some code sharing. The sharing introduced is sufficient to guarantee that the size of the output automaton is linear in the size of the input pattern matching definition. Both compilation schemes have been integrated into the GAML compiler [Maranget 1991] for lazy ML. Finally, I compare these schemes one with another and with previous non-lazy approaches, both from the theoretical and practical points of view.

This paper extends on Maranget [1992a] in many important ways. In particular, the treatment of automata with failures is entirely new (section 4). A thorough presentation of the full compilation algorithm is given (section 5.2), as well as a discussion of the main theoretical an pragmatical aspects of lazy pattern matching (section 6).

2. VALUES AND PATTERNS

Our intention is to model pattern matching as a function on the set of terms representing the results of lazy ML programs.

2.1 Partial values

A constructor is a functional symbol with an arity. A constructor will often be represented by c and its arity by a. Constructors are defined by data type declarations.

Consider for instance the type declaration:

$$\texttt{type tree } \alpha \texttt{ = Leaf } \alpha \texttt{ | Node (tree } \alpha\texttt{) } \alpha \texttt{ (tree } \alpha\texttt{)}$$

This declaration defines the type `tree` α of the binary trees of objects of type α. It introduces the two constructors `Leaf` and `Node`, of respective arities 1 and 3. The set $\{\texttt{Leaf}, \texttt{Node}\}$ is the *signature* of the type `tree` α. Sometimes, to stress upon the fact that there is no other constructor in the signature of `tree` α, the set $\{\texttt{Leaf}, \texttt{Node}\}$ is said to make up a *full signature*. There is at least one predefined type with an infinite signature: the integer type `int`, all integers being seen as nullary constructors.

The distinguished nullary symbol Ω stands for the unknown parts of a value. Intuitively, Ω can be anything, including a non-terminating computation, usually noted as $\bot$. The set $\mathcal{V}_\Omega$ of partial values is the set of terms built from constructors and the symbol Ω:

$$\textit{Partial values} \quad \mathcal{V}_\Omega\colon\ V ::= \Omega \mid c\ V_1\ V_2 \ldots V_a$$

Given a partial value $V = c\ V_1\ V_2 \ldots V_a$, we say that constructor c is the *root constructor* of V. We only consider values that are well typed in the standard sense, with the partial value Ω belonging to all types. For instance, `Node` Ω `1` (`Leaf 2`) has type `tree int`.

A lazy language distinguishes between totally unknown values and partially unknown values. Consider for instance a list of two unknown values, represented as $\Omega::\Omega::[\,]$. We can refer to the length of such a partial list. In a lazy language, we should even be able to compute it. As to the totally unknow value Ω, it does not carry any information at all. This suggests that partial values may be considered as more or less precise approximations of computations results. The definition ordering captures this intuition.

Definition 1. Let U and V be two partial values of the same type. The partial value U is said to be *less defined* than V, written $U \preceq V$, if and only if :

$$\begin{cases} U = \Omega \\ \quad \text{or} \\ \begin{cases} U = c\ U_1 \ldots U_a,\ V = c\ V_1 \ldots V_a \\ \quad \text{and} \\ \text{for all } i \text{ in } [1 \ldots a],\ U_i \preceq V_i \end{cases} \end{cases}$$

Two partial values U and V are said to be *compatible*, written $U \uparrow V$, when they can be refined towards the same partial value, that is, when there exists a common upper bound of U and V. When this is not the case, U and V are *incompatible*, written $U \# V$.

We also consider the set $\mathcal{T}_\Omega$ of well-typed partial terms built from constructors, the symbol Ω and a set of variables (typical element v).

$$\textit{Partial terms}\ \mathcal{T}_\Omega\colon\ M ::= \Omega \mid v \mid c\ M_1\ M_2 \ldots M_a$$

A *substitution*, written σ, is a morphism on partial terms, that is, a function on partial terms such that $\sigma(c\ M_1\ M_2 \ldots M_a) = c\ \sigma(M_1)\ \sigma(M_2) \ldots \sigma(M_a)$ and $\sigma(\Omega) = \Omega$. Thus, a substitution is totally defined by its values on variables. The *domain* of a substitution σ is the set of variables v such that we have $\sigma(v) \neq v$. In practice, only substitutions with finite domains are considered. Such a substitution will

often be written as the environment $[v_1\backslash M_1, \ldots, v_n\backslash M_n]$ binding, for any integer i in $[1 \ldots n]$, the variable v_i to the partial term M_i. Application of an environment to a partial term is written $N[v_1\backslash M_1, \ldots, v_n\backslash M_n]$. For any partial term M, the partial value M_Ω is obtained by substituting Ω for all variables in M (that is, $M_\Omega = M[v_1\backslash\Omega, v_2\backslash\Omega, \ldots, v_n\backslash\Omega]$, where $v_1, v_2, \ldots, v_n$ are the variables of term M).

Given two substitutions σ and ρ with disjoint domains, the union $\sigma \cup \rho$ is the substitution whose domain is the set union of the domains of σ and ρ and that coincides with σ and ρ on their respective domains.

2.2 Patterns

Patterns are strict linear terms, that is, partial terms without Ω such that the same variable does not appear more than once in them. Pattern variables are written as x.

$$\textit{Patterns } \mathcal{P}\colon\ p ::= x \mid c\ p_1\ p_2 \ldots p_a \quad p \text{ is linear}$$

A pattern can be seen as representing a set of (partial) terms sharing a common prefix. Additionally, subterms located under this prefix are bound to pattern variables.

Definition 2. Let p be a pattern and M be a partial term belonging to a common type. The term M is an instance of the pattern p, written $p \preceq M$, if and only if there exists a substitution σ such that $\sigma(p) = M$.

Note that if M is an instance of p, then there exists an unique substitution σ such that $\sigma(p) = M$ and whose domain is minimal. In the following, this unique substitution will be implicitly selected, when I talk about "the" substitution σ such that $\sigma(p) = M$.

The instantiation relation is closely related to the definition ordering.

LEMMA 2.1. *Let p be a pattern and M be a partial term. The following equivalence holds:*

$$p \preceq M \quad \textit{if and only if} \quad p_\Omega \preceq M_\Omega$$

PROOF. Easy induction on p, using the linearity of patterns. □

This lemma justifies using the same operator $\preceq$ for both the instantiation relation and the definition ordering.

A pattern p and a partial term M are incompatible, and we write $p \# M$, when M is sufficiently defined to ensure that it is not an instance of p. That is $p \# M$ holds, if and only if

$$\begin{cases} p = c\ p_1 \ldots p_a,\ M = c'\ M_1 \ldots M_{a'} \text{ with } c \neq c' \\ \quad \text{or} \\ \begin{cases} p = c\ p_1 \ldots p_a,\ M = c\ M_1 \ldots M_a \\ \quad \text{and} \\ \text{there exists } i \text{ such that } p_i \# M_i \end{cases} \end{cases}$$

When pattern p and partial term M are not incompatible, we write $p \uparrow M$. The following equivalence properties, holding for any pattern p and any partial term M, directly follow from lemma 2.1:

$$p \# M \quad \text{iff} \quad p_\Omega \# M_\Omega \qquad \text{and} \qquad p \uparrow M \quad \text{iff} \quad p_\Omega \uparrow M_\Omega$$

The partial term M can be a pattern q. If p and q are compatible, then they are also said to be *ambiguous* or *overlapping*. Indeed, as a consequence of lemma 2.1, two patterns are compatible if and only if they admit a common instance.

2.3 The matching function

Pattern matching is usually formalized as a predicate on partial values [Huet and Lévy 1991; Laville 1991]. I prefer a representation as a function over partial values, closer to pattern matching in ML.

A *clause* is a triple (i, p, e), where i is an integer, p is a pattern and e is a partial term, such that all variables in term e are variables of pattern p. Integer i is the *number* of the clause, whereas term e is its *result*. To simplify notations, we shall write clauses as $p^i\!: e_i$. We consider *sets of clauses* meeting the following three conditions:

(1) All clause numbers are distinct.

(2) All patterns belong to a common type.

(3) All results belong to a common type.

Sets of clauses are written $E = \{p^i\!: e_i \mid i \in I\}$, where I is a set of numbers. (To simplify, assume that clause numbers just express the natural textual clause ordering meant by the programmers, when they write one clause after (under) another.) Clause sets are ordered by the ordering on the clause numbers. The pattern matching function takes this ordering into account to resolve possible ambiguities between patterns, as illustrated by the following definiton of the matching predicate with priorities, first introduced by A. Laville:

Definition 3. Let $E = \{p^1\!: e_1,\ p^2\!: e_2, \ldots, p^m\!: e_m\}$ be a set of clauses. Let V be a partial value. The value V *matches* clause number i in E, written $match_i[E](V)$, if and only if the following two conditions are satisfied:

$$p^i \preceq V \quad \text{and} \quad \forall j < i,\ p^j \# V.$$

Notice that the matching predicates defined by two distinct clause numbers are mutually exclusive, because $p^j \# V$ excludes $p^j \preceq V$.

The matching function is a slight extension of Laville's definition.

Definition 4. Let E be a set of clauses. For any partial value V, we define the partial value $match[E](V)$ as follows:

— If the value V matches clause number i, we take $match[E](V) = \sigma(e_i)$, where σ is the substitution such that $\sigma(p^i) = V$.

— Otherwise, V is a *non-matching value* and we take $match[E](V) = \Omega$.

Given an ordered set of patterns $P = \{p^1, p^2, \ldots, p^m\}$, the matching predicate $match[P]$ is defined as the matching function $match[E]$, where E is the set of clauses $\{p^1\!: \texttt{true}, p^2\!: \texttt{true}, \ldots, p^m\!: \texttt{true}\}$. Obviously, any value V matches some pattern in P, if and only if we have $match[E](V) = \texttt{true}$.

It is easy to show that, given a set of clauses E, the $match[E]$ function is monotonic over partial values. Other rules than the textual priority rule (definition 3) can be used to resolve ambiguity in patterns: in particular, the specificity rule [Kennaway 1990]. We do not consider this alternative, since the textual priority ordering mimics the familiar "*if* $condition_1$ *then* $result_1$ *else if* $condition_2$ *then* $result_2$ ..."

construct. Furthermore, both schemes have the same expressive power [Laville 1991].

Pattern matching expressions can also be written as ML programs. If a pattern variable does not appear in the corresponding result expression, then its name is unimportant and the pattern variable is replaced by the symbol "_". Consider, for instance, the set of clauses $E = \{(x_1, \texttt{true}) : \texttt{true}, (\texttt{true}, x_2) : \texttt{true}, (x_3, x_4) : \texttt{false}\}$ and the function $or(V) = match[E](V)$ (Here ","is the pair constructor, the sole constructor of the pair type). In ML syntax, we have:

```
or(V) =  match V with
    (_,true)    → true
|   (true,_)    → true
|   (_,_)       → false
```

There is a finite number of partial values of type bool × bool. From definition 4, we get:

V	$or(V)$
Ω (Ω,Ω) $(\Omega,\texttt{false})$ $(\texttt{true},\Omega)$ $(\texttt{false},\Omega)$	Ω
$(\Omega,\texttt{true})$ $(\texttt{false},\texttt{true})$ $(\texttt{true},\texttt{true})$ $(\texttt{true},\texttt{false})$	true
$(\texttt{false},\texttt{false})$	false

Note that *or* is *not* the "parallel or" function *por*, since, by definition, we have $por(\Omega, \texttt{true}) = por(\texttt{true}, \Omega) = \texttt{true}$.

It may seem that the definition of pattern matching might be simplified by replacing the second condition $\forall j < i,\ p^j \# V$ by the new and less strict condition $\forall j < i,\ p^j \not\preceq V$. Such a change is not advisable though, since it would imply losing monotonicity. Consider, for instance, the pattern matching $match[E](V)$ defined by:

```
match V with 1  → Ω  |  _  →  2
```

the above modified definition of the matching function would give us:

$$match[E](\Omega) = 2 \not\preceq match[E](1) = \Omega$$

2.4 Pattern matching on vectors

In the next section, we shall consider "intermediate" matchings. In these matchings, the value to match and the patterns have a common prefix: the part of the value examined so far. More precisely, let n be a positive integer, let $v_1, v_2 \ldots v_n$ be n variables and let N be a linear partial term whose variables are $v_1, v_2 \ldots v_n$. An intermediate matching is a pattern matching of the format:

```
match N[v1\V1, ..., vn\Vn] with
   N[v1\p1^1, ..., vn\pn^1] → e1
      ⋮
|  N[v1\p1^m, ..., vn\pn^m] → em
```

Obviously, the result of such a matching does not depend on the prefix N, but only on the partial values V_i and patterns p_i^j that are substituted for the variables v_i.

The n partial values may be seen as a vector $\vec{V} = (V_1\ V_2 \dots V_n)$, whereas each clause may be seen as a vector clause consisting of a number i, of a vector of n patterns $\vec{p}^{\,i}$ and of a result term e_i. The set of clauses is replaced by a *pattern matrix* (P) and a *term vector* (E) written as:

$$(P) = \begin{pmatrix} p_1^1 & p_2^1 \dots p_n^1 \\ p_1^2 & p_2^2 \dots p_n^2 \\ & \vdots \\ p_1^m & p_2^m \dots p_n^m \end{pmatrix} \qquad (E) = \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_m \end{pmatrix}$$

In pattern p_j^i, the index i is the clause or row number, whereas j is the column index.

The instantiation and the incompatibility relations on patterns and values trivially extend to vectors:

$$(p_1\ p_2 \dots p_n) \preceq (V_1\ V_2 \dots V_n), \text{ iff for all } i \text{ in } 1 \dots n, \text{ we have } p_i \preceq V_i$$

$$(p_1\ p_2 \dots p_n)\ \#\ (V_1\ V_2 \dots V_n), \text{ iff there exists } i \text{ in } 1 \dots n \text{ such that } p_i\ \#\ V_i$$

It is then straightforward to extend the definition of the matching predicate to vectors of partial values $\vec{V}$ and pattern matrices (P):

$$match_i[(P)](\vec{V}), \text{ if and only if } \begin{cases} \vec{p}^{\,i} \preceq \vec{V} \\ \quad \text{and} \\ \text{for all } j < i, \text{ we have } \vec{p}^{\,j}\ \#\ \vec{V} \end{cases}$$

Finally, once substitutions are extended to operate on vectors in the natural way: $\sigma(\vec{p}) = (\sigma(p_1)\ \sigma(p_2) \dots \sigma(p_n))$, the matching function also extends to vectors: if $\vec{V}$ matches the row number i in matrix (P), then there exists a substitution σ such that $\sigma(\vec{p}^{\,i}) = \vec{V}$ and we take $match[(P),(E)](\vec{V}) = \sigma(e_i)$. Otherwise, the vector $\vec{V}$ does not match any clause in $(P),(E)$ and we take $match[(P),(E)](\vec{V}) = \Omega$.

Another natural use of pattern matching on vectors is getting rid of superfluous tuples. To see this, consider the introductory example of this section, assuming that the prefix N is the n-tuple $(v_1, v_2, \dots, v_n)$. For instance, the function *or* from the previous section can be simplified, provided the value of $or\,(\Omega)$ is of no importance.

$$\texttt{or}(x\ y) = match[\begin{pmatrix} \texttt{_} & \texttt{true} \\ \texttt{true} & \texttt{_} \\ \texttt{_} & \texttt{_} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{true} \\ \texttt{false} \end{pmatrix}](x\ y)$$

Here also, an alternative formulation as a ML definition exists:

```
let or _    true = true
|   or true _    = true
|   or _     _    = false
```

3. FIRST TECHNIQUE

3.1 Tree-like automata

Tree-like automata are the simplest and most natural among sequential automata. Other authors call them decision trees or matching trees. Every state in a tree-like automata is characterized by some position u to examine and some already

examined prefix N; both position and prefix are relative to the currently examined term. Every transition coming out of a state is labelled by some constructor symbol c, that can take place at position u. These automata are described here as ML programs. These particular ML programs are nested simple matchings, simple matching being the natural representation of elementary comparison in ML.

Simple matchings:
F ::= `match` x `with` $p \rightarrow A$ $\{\mid p \rightarrow A\}^*$ $[\mid _ \rightarrow A]$

Simple patterns:
p ::= $c\ x_1\ x_2 \ldots x_a$ and p is linear

Automata:

A ::=	F	nested matching
$\mid$	M	Partial term

A typical simple matching is thus written as follows:

$$F = \left\{ \begin{array}{lll} \texttt{match}\ x\ \texttt{with} & & \\ \quad c_1\ x_1\ x_2 \ldots x_{a_1} & \rightarrow & A_1 \\ \mid \quad \vdots & & \\ \mid\ c_z\ x_1\ x_2 \ldots x_{a_z} & \rightarrow & A_z \\ \mid\ _ & \rightarrow & A_d \end{array} \right.$$

We further enforce that all *recognized constructors* c_1, c_2,..., c_z are different and belong to the same signature. When present, the last clause $_ \rightarrow A_d$ is the *default clause*. Additionally, we require simple matchings to be *complete*: when the constructors recognized by a given simple matching do not make up a full signature, then this simple matching must possess a default clause.

The functions $\mathtt{F}_1$ and $\mathtt{F}_2$ from the introduction are example of tree-like automata, provided the `if ... then ... else ...` construct is translated into the equivalent `match ... with true` $\rightarrow$ `...` $\mid$ `false` $\rightarrow$ `...` construct.

the semantics of tree-like automata can be described by automaton *states* and *transitions* from one state to another. An automaton state is a pair (σ, A), where σ is an environment (i.e., a substitution) and A is an automaton. The following one-step transitions describe the run-time behavior of simple matchings:

If $\sigma(x) = c_k\ V_1 \ldots V_{a_k}$,
$(\sigma, \texttt{match}\ x\ \texttt{with}\ \ldots \mid c_k\ x_1 \ldots x_{a_k} \rightarrow A_k\ \ldots) \hookrightarrow \ldots$
$\ldots \quad (\sigma \cup [x_1 \backslash V_1, \ldots, x_{a_k} \backslash V_{a_k}], A_k)$ (constr)

If $\sigma(x) = c\ V_1 \ldots V_a$ where c is not recognized, then
$(\sigma, \texttt{match}\ x\ \texttt{with}\ \ldots \mid _ \rightarrow A_d) \hookrightarrow (\sigma, A_d)$ (default)

Observe that some condition on variable names is implicitly assumed, so that the union of substitutions in the right-hand side of rule (constr) can be performed safely. Obviously, requiring the pattern variables in an automaton to be all different is enough.

Simple matchings examine the value bound to some variable x. Given the completeness requirement, there are three possible cases:

— If $\sigma(x)$ has a root constructor c_k that is recognized (rule (constr)), the the analysis of the input values will proceed by using the right-hand side of the matched

clause $c_k\ x_1\ x_2 \ldots x_{a_k} \to A_k$, after some appropriate extension of the current environment.

— If $\sigma(x)$ has a root constructor c that is not recognized (rule (default)), then the analysis of the input values will proceed by using the right-hand side of the default clause.

— Otherwise, $\sigma(x)$ is Ω and is not defined enough to make a decision.

A *final state* is any state from which no transition is possible. The result of an automaton A in the initial environment σ is a value written $\mathcal{RUN}(\sigma, A)$. This value is computed as follows: first, repetitively apply the transition rules, starting from the initial state (σ, A), until some final state (σ', A') is reached, thereby computing a *complete transition*. Second, consider the final automaton A'. If A' is the right-hand side of a clause (up to some variable names), then we say that the execution of A is *successful* and we take $\mathcal{RUN}(\sigma, A) = \sigma'(A')$. Otherwise, the execution of A has failed and we take $\mathcal{RUN}(\sigma, A) = \Omega$.

3.2 Compilation

Compilation is defined as a function $\mathcal{C}$, mapping general pattern matching expressions to nested simple pattern matching expressions. More precisely, $\mathcal{C}$ takes three arguments and a typical call is written $\mathcal{C}(\vec{x}, (P), (E))$. The first argument $\vec{x}$ is a linear vector of n variables $(x_1\ x_2 \ldots x_n)$, it abstracts the partial value vectors given as arguments to the compiled matching. In other words, given an automaton $A = \mathcal{C}((x_1\ x_2 \ldots x_n), (P), (E))$ and an environment $\sigma = [x_1 \backslash V_1, \ldots, x_n \backslash V_n]$, $\mathcal{C}((V_1\ V_2 \ldots V_n), (P), (E))$ is defined as the result of the execution of A in the initial environment σ.

$$\mathcal{C}((V_1\ V_2 \ldots V_n), (P), (E)) = \mathcal{RUN}([x_1 \backslash V_1, \ldots, x_n \backslash V_n], \mathcal{C}(\vec{x}, (P), (E)))$$

The second argument (P) is a pattern matrix of size n by m. Initially, (P) is the pattern matrix to be compiled. Later on, it represents the yet unprocessed subparts of the initial patterns. Finally, the third argument (E) is a term vector of size m. Initially, it is made of the right-hand sides of the compiled clauses. Later on, it represents the result terms that still can be reached at some compilation stage.

The function $\mathcal{C}$ is inductively defined as follows:

(1) If (P) is empty ($n = 0$ and $m = 0$), then the compilation is finished and matching will always fail (since there are no rows to match here):

$$\mathcal{C}((), (), ()) = \Omega$$

(2) If the first row of matrix (P) is made of variables only, then matching will always succeed and give the first component of vector (E) as its result. Note that this case includes the case where (P) is made of empty rows ($n = 0$, $m > 0$).

$$\mathcal{C}((x_1\ x_2 \ldots x_n), \begin{pmatrix} y_1 & y_2 & \cdots & y_n \\ p_1^2 & p_2^2 & \cdots & p_n^2 \\ & \vdots & & \\ p_1^m & p_2^m & \cdots & p_n^m \end{pmatrix}, \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_m \end{pmatrix}) = e_1[y_1 \backslash x_1, \ldots, y_n \backslash x_n]$$

(3) Matching can also progress by examining one of the variables x_i such that the column number i in (P) possesses at least one non-variable pattern. Until otherwise

stated, the choice of this variable is arbitrary and the result of compilation depends *a priori* on it. When the variable x_i is chosen, we shall say that compilation progresses by following index i or that i is the *discriminating* index. To be more specific, assume that the first variable x_1 is chosen. Let $\Sigma = \{c_k \mid 1 \leq k \leq z\}$ be the set of the root constructors of the patterns in the first column of (P). To each constructor c_k, of arity a_k, a new matrix (P_k) is associated. Matrix (P_k) contains the rows of (P) that can match a value vector of the format $((c_k\ U_1 \ldots U_{a_k})\ V_2 \ldots V_n)$. More precisely, the following table shows how each row of the new matrix (P_k) and each component of the new vector (E_k) are built:

p_1^i	row in (P_k) $\quad$ (E_k)
y	$_ \cdots _ \quad p_2^i \cdots p_n^i \quad e_i[y \backslash x_1]$
$c_k\ q_1^i \ldots q_{a_k}^i$	$q_1^i \cdots q_{a_k}^i \quad p_2^i \cdots p_n^i \qquad e_i$
$c_{k'}\ q_1^i\ q_2^i \ldots q_{a_{k'}}^i \quad (k' \neq k)$	No row number i

If the set of constructors Σ is not a full signature, then some rows in matrix (P) can match value vectors of the format $((c\ U_1 \ldots U_a)\ V_2 \ldots V_n)$, where c is a constructor that does not belong to Σ. To treat these cases, a default matrix (P_d) is built:

p_1^i	row in (P_d) $\quad$ (E_d)
y	$p_2^i \cdots p_n^i \quad e_i[y \backslash x_1]$
$c\ q_1^i\ q_2^i \ldots q_a^i$	No row number i

Of course, the original ordering of the rows is preserved in the new matrices (P_k) and (P_d), so that the clauses are arranged by increasing numbers. Compilation then proceeds by emitting an elementary test on variable x_1. At run-time, this test will direct analysis towards the appropriate subproblem.

$$\mathcal{C}((x_1\ x_2 \ldots x_n), (P), (E)) = \left\{ \begin{array}{ll} \texttt{match}\ x_1\ \texttt{with} & \\ \quad c_1\ y_1 \ldots y_{a_1} & \to \mathcal{C}((y_1 \ldots y_{a_1}\ x_2 \ldots x_n), (P_1), (E_1)) \\ \mid\ c_2\ y_1 \ldots y_{a_2} & \to \cdots \\ & \vdots \\ \mid\ c_z\ y_1 \ldots y_{a_z} & \to \mathcal{C}((y_1 \ldots y_{a_z}\ x_2 \ldots x_n), (P_z), (E_z)) \\ \mid\ _ & \to \mathcal{C}((x_2 \ldots x_n), (P_d), (E_d)) \end{array} \right.$$

Compilation always terminates, since the total number of constructors contained in the patterns of (P) strictly decreases at each recursive call to $\mathcal{C}$.

As an example of a full compilation, consider the function `or` from the end of section 2.4. The initial call to function $\mathcal{C}$ is as follows:

$$\mathcal{C}((\texttt{x y}), \begin{pmatrix} \texttt{_} & \texttt{true} \\ \texttt{true} & \texttt{_} \\ \texttt{_} & \texttt{_} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{true} \\ \texttt{false} \end{pmatrix})$$

There are two possible compilations for the `or` function, depending on which of

the variables `x` or `y` is examined first:

$$
\begin{array}{l}
\texttt{match x with} \\
\quad \texttt{true} \to \mathcal{C}((\texttt{y}), \begin{pmatrix} \texttt{true} \\ \texttt{_} \\ \texttt{_} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{true} \\ \texttt{false} \end{pmatrix}) \\
\mid \texttt{_} \quad \to \mathcal{C}((\texttt{y}), \begin{pmatrix} \texttt{true} \\ \texttt{_} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{false} \end{pmatrix}) \\
\qquad \vdots \\
\qquad \vdots
\end{array}
$$

$$
\texttt{or}_1\texttt{(x y)} = \begin{cases} \texttt{match x with} \\ \quad \texttt{true} \to \texttt{(match y with true} \to \texttt{true} \mid \texttt{_} \to \texttt{true)} \\ \mid \texttt{_} \quad \to \texttt{(match y with true} \to \texttt{true} \mid \texttt{_} \to \texttt{false)} \end{cases}
$$

$$
\begin{array}{l}
\texttt{match y with} \\
\quad \texttt{true} \to \mathcal{C}((\texttt{x}), \begin{pmatrix} \texttt{_} \\ \texttt{true} \\ \texttt{_} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{true} \\ \texttt{false} \end{pmatrix}) \\
\mid \texttt{_} \quad \to \mathcal{C}((\texttt{x}), \begin{pmatrix} \texttt{true} \\ \texttt{_} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{false} \end{pmatrix}) \\
\qquad \vdots \\
\qquad \vdots
\end{array}
$$

$$
\texttt{or}_2\texttt{(x y)} = \begin{cases} \texttt{match y with} \\ \quad \texttt{true} \to \texttt{true} \\ \mid \texttt{_} \quad \to \texttt{(match x with true} \to \texttt{true} \mid \texttt{_} \to \texttt{false)} \end{cases}
$$

The automata $\texttt{or}_1$ and $\texttt{or}_2$ are syntactically different, they also differ semantically. Indeed, we get here: $\texttt{or}_1(\Omega\ \texttt{true}) = \Omega$ —since simple matching of variable `x` immediately fails—, whereas $\texttt{or}_2(\Omega\ \texttt{true}) = \texttt{true}$. For all other values of the boolean vector $\vec{V}$, we get $\texttt{or}_1(\vec{V}) = \texttt{or}_2(\vec{V}) = \texttt{or}(\vec{V})$. Thus, $\texttt{or}_2$ is the only automaton that implements correctly the function `or`.

3.3 Correctness

A compilation is correct when the execution of the output automaton implements exactly the input pattern matching function.

Definition 5. Given a pattern matrix (P) and a result vector (E), compilation $\mathcal{C}(\vec{x}, (P), (E))$ is correct, if and only if the following equality holds for all value vectors $\vec{V}$:

$$\mathcal{C}(\vec{V}, (P), (E)) = match[(P), (E)](\vec{V})$$

As shown by the automaton $\texttt{or}_1$ from the previous section, the $\mathcal{C}$ compilation scheme may not be correct. In this case, the problem can be traced back to the untimely examination of the first component of $(\Omega\ \texttt{true})$, while $(\Omega\ \texttt{true})$ clearly

matches the first clause of the function `or`. The idea is to select carefully column indices, in order to avoid such a situation.

Definition Directions. Let (P) be a pattern matrix of width n. The column index d such that $1 \leq d \leq n$ is a *direction* for the matching by (P), written $d \in Dir(P)$, if and only if the following two conditions are met:

(1) The vector $\vec{\Omega}$ does not match any row in (P).

(2) There is no vector $\vec{V} = (V_1\ V_2 \ldots V_n)$, such that $\vec{V}$ matches a pattern row in (P) and $V_d = \Omega$.

The first condition above is just here to rule out the trivial case where all value vectors match the first row of matrix (P).

Directions are computable from the matrix (P). Consider the set $Dir_i(P)$ of *directions towards row number* i, defined as:

$$Dir_i(P) = \{\, d \in [1 \ldots n] \mid match_i[(P)](\vec{V}) \Rightarrow V_d \succ \Omega \,\}$$

Then, $Dir(P)$ is the intersection of the sets $Dir_i(P)$. For instance, at the critical compilation step for the function *or*, we have:

$$(P) = \begin{pmatrix} _ & \texttt{true} \\ \texttt{true} & _ \\ _ & _ \end{pmatrix} \quad \text{and} \quad \begin{cases} match_1[(P)](\vec{V}) \Leftrightarrow \texttt{true} \preceq V_2 \\ match_2[(P)](\vec{V}) \Leftrightarrow \texttt{true} \preceq V_1 \wedge \texttt{true} \,\#\, V_2 \\ match_3[(P)](\vec{V}) \Leftrightarrow \texttt{true} \,\#\, V_1 \wedge \texttt{true} \,\#\, V_2 \end{cases}$$

Thus, we get $Dir_1(P) = \{2\}$, $Dir_2(P) = \{1, 2\}$, $Dir_3(P) = \{1, 2\}$ and hence $Dir(P) = \{2\}$. In section 5.1, I give a full description of a general and efficient method for computing directions.

It is now time to examine again the arbitrary choice of the variable to be tested at the critical compilation step 3.

Definition 7. Let E be a set of clauses. A given compilation of the matching by E is done by *following directions*, when, at each inductive application of step 3, there exists a direction d in pattern matrix (P) and that compilation goes on by following index d.

Directions gives a sufficient condition to assert the correctness of a given compilation.

Proposition Correct compilation. *Compilations that follow directions produce correct automata.*

Proof. In fact given any compilation step $A = \mathcal{C}(\vec{x}, (P), (E))$ and any substitution σ, we show $\mathcal{RUN}(\sigma, A) = match[(P), (E)](\sigma(\vec{x}))$, a property which is slightly stronger than the correctness identity itself.

The base cases are easy. If A is Ω, then we get $\mathcal{RUN}(\sigma, A) = \Omega$. Moreover, by definition of $\mathcal{C}$, (P) has no row to be matched. Otherwise, A is a clause right-hand side and any vector matches the first row of (P). Now, consider the case of the inductive step 3, where some elementary test is emitted. We adopt the notations from section 3.2 and assume that column index 1 is a direction of (P). There are three subcases to consider, depending on $\sigma(x_1)$.

If $\sigma(x_1)$ is Ω, then, on the one hand, the simple matching of $\sigma(x_1)$ will always fail and we get $\mathcal{RUN}(\sigma, A) = \Omega$. On the other hand, since i is a direction of (P), no value vector with an undefined first component matches some row of (P) and we get $match[(P),(E)](\sigma(x)) = \Omega$.

If $\sigma(x_1)$ is of the form $c\ V_1 \ldots V_a$, where the constructor c is not one of the root constructors appearing in the first column of (P), then, by the rule (default) of simple matching, we get the following equality:

$$\mathcal{RUN}(\sigma, \mathcal{C}(\vec{x}, (P), (E))) = \mathcal{RUN}(\sigma, \mathcal{C}((x_2 \ldots x_n), (P_d), (E_d)))$$

Therefore, by induction hypothesis, it follows that:

$$\mathcal{RUN}(\sigma, \mathcal{C}(\vec{x}, (P), (E))) = match[(P_d), (E_d)](\sigma(x_2) \ldots \sigma(x_n))$$

Moreover, for any valid row number, the following equivalences hold:

$$\vec{p}^{\,i} \preceq ((c\ V_1 \ldots V_a) \ldots U_n) \Leftrightarrow p_1^i = y^i \text{ and } (p_2^i \ldots p_n^i) \preceq (U_2 \ldots U_n)$$

$$\vec{p}^{\,i} \# ((c\ V_1 \ldots V_a) \ldots U_n) \Leftrightarrow \begin{cases} p_1^i = c_k\ q_1^i \ldots q_{a_k}^i \\ \text{or} \\ p_1^i = y^i \text{ and } (p_2^i \ldots p_n^i) \# (U_2 \ldots U_n) \end{cases}$$

Thus, in this case, the matching predicates defined by (P) and (P_d) are equivalent and correctness follows.

The last case $\sigma(x_1) = c_k\ V_1 \ldots V_{a_k}$, where c_k is one of the root constructors appearing in the first column of (P) is proved by a very similar reasoning. □

As an application of proposition 3.1, in the case of the function *or*, the automaton $\texttt{or}_2$ can be stated as correct without testing it on all partial values, since it always performs simple matchings by following directions.

4. SECOND TECHNIQUE

Simplicity is the main advantage of tree-like automata. Unfortunately, these automata have a major drawback: the size of output automata can be exponential in the size of the input programs [Maranget 1992a; Sekar, Ramesh, and Ramakrishnan 1992]. This drawback is particularly annoying when one thinks of integrating the $\mathcal{C}$ pattern matching compilation scheme in an operational ML compiler. One has to notice though, that compilers as mature as SML/NJ [Appel and MacQueen 1991; Baudinet and MacQueen 1986] and Caml V3.1 [Weis 1990] use a similar compilation technique. Nevertheless, there already exists a different kind of automata for implementing pattern matching, such automata are the targets of the LML [Augustsson 1985] and Caml Light [Leroy 1995; Leroy 1990] compilers. One salient feature of these automata is a backtracking construct. Standard pattern matching compilation schemes rely on it to guarantee that the size of output automaton is linear in the size of the input program. The main contribution of the present paper is to show how lazy pattern matching can be compiled correctly in this second framework, while preserving the highly desirable linear bound on the size of output automata.

4.1 Automata with failures

From the syntactic point of view, there is little change between tree-like and automata with failures: the right-hand side of a clause can now be a new `default` primitive.

Simple matchings:
F ::= `match` x `with` $p \rightarrow A$ {| $p \rightarrow A$}* [| _ $\rightarrow A$]

Simple patterns:
p ::= $c\ x_1\ x_2 \ldots x_a$ and p is linear

Automata:

A ::=	F	nested matching
|	M	partial term
|	`default`	default primitive

Semantically, when execution comes across a `default` primitive, control is transferred to the closest enclosing default clause. This behavior is best described by extended automaton states and transition rules. An automaton state is now a triple (ρ, σ, A), where A is an automaton, σ is a value environment and ρ is a control environment, a control environment being a list of partial automaton states $<\sigma, A>$, where σ is a value environment and A is an automaton. The following extended one-step transitions describe the run-time behavior of automata with failures.

If $\sigma(x) = c_k\ V_1 \ldots V_{a_k}$,
$(\rho, \sigma, \texttt{match}\ x\ \texttt{with}\ \ldots \mid c_k\ x_1 \ldots x_{a_k} \rightarrow A_k \ \ldots \mid _ \rightarrow A_d) \hookrightarrow \quad \ldots$
$\ldots \quad (<\sigma, A_d>.\rho, \sigma \cup [x_1 \backslash V_1, \ldots, x_{a_k} \backslash V_{a_k}], A_k)$ (constr)

If $\sigma(x) = c\ V_1 \ldots V_a$ where c is not recognized,
$(\rho, \sigma, \texttt{match}\ x\ \texttt{with}\ \ldots \mid _ \rightarrow A_d) \hookrightarrow (\rho, \sigma, A_d)$ (default)

$(<\phi, A>.\rho, \sigma, \texttt{default}) \hookrightarrow (\rho, \phi, A)$ (raise)

Observe that control environments are stacks. That is, the transition (constr) pushes the current default clause on the stack, while the transition (raise) pops the closest enclosing default clause from the stack.

We still enforce the completeness requirement on simple matchings and the unicity requirement on pattern variables. Furthermore, we enforce the more specific condition that the transition (raise) never fails, that is, that the `default` primitive always finds a non-empty control environment. This run-time condition can be guaranteed by a simple syntactic verification on automata: every `default` primitive must be a subterm of the right-hand side of a non-default clause of some simple matching that possesses a default clause.

The result $\mathcal{RUN}(\rho, \sigma, A)$ of the execution of some automaton A in the initial environments σ and ρ is a value defined as in the case of tree-like automata. First, compute the final state (ρ', σ', A'). Then, either A' is the right hand side of a clause and we take $\mathcal{RUN}(\rho, \sigma, A) = \sigma'(A')$; otherwise, the execution has failed and we take $\mathcal{RUN}(\rho, \sigma, A) = \Omega$.

4.2 Compilation

Thanks to the `default` primitive, automata with failures can directly implement pattern matching with priorities: try to match against the first pattern, in case of failure, try to match against the second pattern,... In practice, one seeks to "factorize" matching attempts, by trying to match many patterns at the same time.

The new compilation function $\mathcal{D}$ takes four arguments. A typical call is written $\mathcal{D}(\vec{x},(P),(E),d)$. The first three arguments are as in section 3.2: $\vec{x}$ is a variable vector of size n, (P) is a pattern matrix of width n and height m and (E) is a term vector of size m. The fourth argument d is itself an automaton and may be seen as the "failure continuation" of the compiled matching. Control is to be transferred to the automaton d in case of matching failure. The initial call to $\mathcal{D}$ is thus written as follows:

$$\mathcal{D}((x),\begin{pmatrix} p^1 \\ p^2 \\ \vdots \\ \\ p^m \end{pmatrix},\begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ \\ e_m \end{pmatrix},\Omega)$$

The function $\mathcal{D}$ is then defined inductively:

(1) base case: $n=0$. If (P) has no rows ($m=0$), then the output automata reduces to the default automaton d.

$$\mathcal{D}((),(),(),d)=d$$

Otherwise, there is at least one empty row in (P) and the compilation result is the first component of (E).

$$\mathcal{D}((),\begin{pmatrix} \\ \vdots \\ \end{pmatrix},\begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_m \end{pmatrix},d)=e_1$$

(2) If the matrix (P) has at least one column ($n>0$), then select one discriminating index arbitrarily. Let us assume for the moment that the first column of (P) is chosen. There are then two subcases:

(a) If the pattern p_1^1 is a variable y^1, then let k be highest row number such that all patterns p_1^1, p_1^2, ..., p_1^k are variables y^1, y^2, ..., y^k. The matrix (P) is cut horizontally into two new submatrices (Q) and (R) of respective heights k and $k-m$:

$$(Q)=\begin{pmatrix} y^1 & p_2^1 & \cdots & p_n^1 \\ y^2 & p_2^2 & \cdots & p_n^2 \\ & \vdots & & \\ y^k & p_2^k & \cdots & p_n^k \end{pmatrix} \qquad (R)=\begin{pmatrix} p_1^{k+1}(\neq y) & p_2^{k+1} & \cdots & p_n^{k+1} \\ p_1^{k+2} & p_2^{k+2} & \cdots & p_n^{k+2} \\ & \vdots & & \\ p_1^m & p_2^m & \cdots & p_n^m \end{pmatrix}$$

Two nested inductive calls to $\mathcal{D}$ are then performed, where the main argument to the outer call is (Q) without its first column and the main argument to the inner call is (R).

$$\mathcal{D}(\vec{x},\,(P)\,,(E),\,d)$$

$$\Big\|$$

$$\mathcal{D}((x_2\ x_3 \ldots x_n),\ \begin{pmatrix} p_2^1 \ldots p_n^1 \\ p_2^2 \ldots p_n^2 \\ \vdots \\ p_2^k \ldots p_n^k \end{pmatrix},\ \begin{pmatrix} e_1[y^1 \backslash x_1] \\ e_2[y^2 \backslash x_1] \\ \vdots \\ e_k[y^k \backslash x_1] \end{pmatrix},\ \mathcal{D}(\vec{x},\ (R),\ \begin{pmatrix} e_{k+1} \\ e_{k+2} \\ \vdots \\ e_m \end{pmatrix},\ d)\,)$$

(b) If the pattern p_1^1 has a root constructor, then let $k+1$ be the lowest row number such that the pattern p_1^{k+1} is a variable. As above, (P) is cut into two new submatrices of respective heights k and $m-k$:

$$(Q) = \begin{pmatrix} p_1^1\ p_2^1 \cdots p_n^1 \\ p_1^2\ p_2^2 \cdots p_n^2 \\ \vdots \\ p_1^k\ p_2^k \cdots p_n^k \end{pmatrix} \qquad (R) = \begin{pmatrix} p_1^{k+1}(=y)\ p_2^{k+1} \cdots p_n^{k+1} \\ p_1^{k+2} \quad p_2^{k+2} \cdots p_n^{k+2} \\ \vdots \\ p_1^m \quad p_2^m \cdots p_n^m \end{pmatrix}$$

The result vector (E) is also cut in the same way:

$$(E') = \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_k \end{pmatrix} \qquad (F) = \begin{pmatrix} e_{k+1} \\ e_{k+2} \\ \vdots \\ e_m \end{pmatrix}$$

Let c_1, c_2, …, c_z be the root constructors that appear in the first column of (Q). The rows of (Q) are grouped according to the root constructor of their first pattern, yielding the new matrices (Q_1), (Q_2), …, (Q_z). The decomposition process is exactly the same as in the case of tree-like automata —section 3.2, case 3 of the definition of $\mathcal{C}$—, it also produces the result vectors (E_1), (E_2), …, (E_z) from (E'). An important point is that no row of (Q) will ever be duplicated here, since there are no variables in the first column of (Q). A simple matching on variable x_1 is finally emitted:

$$\mathcal{D}((x_1\ x_2 \ldots x_n), (P), (E), d)$$

$$\Big\|$$

$$\begin{array}{ll} \texttt{match } x_1 \texttt{ with} & \\ \quad c_1\ y_1 \ldots y_{a_1} & \to \mathcal{D}((y_1 \ldots y_{a_1}\ x_2 \ldots x_n), (Q_1), (E_1), \texttt{default}) \\ \mid\ c_2\ y_1 \ldots y_{a_2} & \to \cdots \\ & \quad\ \vdots \\ \mid\ c_z\ y_1 \ldots y_{a_z} & \to \mathcal{D}((y_1 \ldots y_{a_z}\ x_2 \ldots x_n), (Q_z), (E_z), \texttt{default}) \\ \mid\ _ & \to \mathcal{D}((x_1\ x_2 \ldots x_n), (R), (F), d) \end{array}$$

At run-time, this simple matching will examine the root constructor of the value of x_1 and will select the appropriate subproblem (Q_1), (Q_2),… (Q_z) or the default subproblem (R), depending on whether this root constructor is one of the recognized constructors $c_1, c_2, \ldots c_z$ or not. The failure continuation of the inductive calls (Q_1), (Q_2),… (Q_z) is the `default` construct, so that control will be transferred to the default subproblem (R), in case any selected subproblem (Q_i) fails.

Compilation always terminates, since the sum of the sizes of the patterns in (P) strictly decreases at every inductive call. Moreover, since no pattern ever gets duplicated, the total number of non-default clauses in the simple matchings produced is bounded by the sum of the sizes of the initial patterns. Taking the default clauses into account, the size of the output automaton is bounded by twice the size of the input program.

Readers already familiar with [Augustsson 1985; Wadler 1987] will notice the similarities between the $\mathcal{D}$ scheme and the algorithms proposed by these authors. Nevertheless, there are two differences. First, the $\mathcal{D}$ scheme chooses an arbitrary discriminating column at step (2), whereas Augustsson and Wadler alogrithms systematically select the leftmost column. Second, our presentation is different and (in our opinion) slightly simpler: when confronted with a pattern matrix whose discriminating column is a mixture of variable and non-variable patterns, the algorithms described by Augustsson and Wadler explicitly build a list of pattern matrices whose discriminating columns are entirely made of variable or non-variable patterns (by the so-called mixture rule). In the $\mathcal{D}$ scheme, the repetition of inductive rule (2)-a or (2)-b incrementally performs a similar partition in a more flexible way, since the discriminating column may change from one submatrix to the next.

As an example, consider the following function definition:

```
let or'  false false  = false
|   or'  true  _      = true
|   or'  _     true   = true
```

The initial call to function $\mathcal{D}$ is as follows:

$$\mathcal{D}((x\ y), \begin{pmatrix} \texttt{false} & \texttt{false} \\ \texttt{true} & \texttt{_} \\ \texttt{_} & \texttt{true} \end{pmatrix}, \begin{pmatrix} \texttt{false} \\ \texttt{true} \\ \texttt{true} \end{pmatrix}, \Omega)$$

There are now two possibiities, depending on whether compilation proceeds left-to-right or right-to-left. In the first case, `x` is examined first:

$$\begin{array}{l} \texttt{match x with} \\ \quad \texttt{false} \to \mathcal{D}((y), (\texttt{false}), (\texttt{false}), \texttt{default}) \\ \mid\ \texttt{true} \ \to \mathcal{D}((y), (\texttt{_}), (\texttt{true}), \texttt{default}) \\ \mid\ \texttt{_} \quad\ \to \mathcal{D}((x\ y), (\texttt{_ true}), (\texttt{true}), \Omega) \end{array}$$

This first left-to-right compilation finally yields the automaton $\texttt{or'}_1$:

$$\begin{array}{l} \texttt{match x with} \\ \quad \texttt{false} \to (\texttt{match y with false} \to \texttt{false} \mid \texttt{_} \to \texttt{default}) \\ \quad \texttt{true} \ \to \texttt{true} \\ \mid\ \texttt{_} \quad\ \to (\texttt{match y with true} \to \texttt{true} \mid \texttt{_} \to \Omega) \end{array}$$

The second compilation, which examines `y` before `x`, produces the following automaton $\texttt{or'}_2$:

$$
\begin{array}{l}
\texttt{match y with} \\
\quad \texttt{false} \rightarrow \mathcal{D}((x), (\ \texttt{false}\ ), (\ \texttt{false}\ ), \texttt{default}) \\
\mid\ _ \quad\ \rightarrow \mathcal{D}((x\ y), \begin{pmatrix} \texttt{true} & \texttt{_} \\ \texttt{_} & \texttt{true} \end{pmatrix}, \begin{pmatrix} \texttt{true} \\ \texttt{true} \end{pmatrix}, \Omega) \\
\qquad\qquad\qquad\qquad\qquad\qquad \vdots
\end{array}
$$

```
match y with
  false → (match x with false → false | _ → default)
| _     → (match x with
             true → true
           | _    → (match y with true → true | _ → Ω))
```

By explicit computation of all cases, one easily shows that the function `or'`, the automaton $\texttt{or'}_1$ and the automaton $\texttt{or'}_2$ coincide on all possible value vectors $\vec{V} =$ (`x y`), except for the case $\vec{V} =$ (`true` Ω), where we get $\texttt{or'}(\vec{V}) = \texttt{or'}_1(\vec{V}) = \texttt{true}$ and $\texttt{or'}_2(\vec{V}) = \Omega$. That is, the automaton $\texttt{or'}_1$ implements correctly the function `or'`, whereas the automaton $\texttt{or'}_2$ does not. Observe that $\vec{V} =$ (`true` Ω) matches the second clause of function `or'` and that the non-correctness of the automaton $\texttt{or'}_2$ can be traced back to the untimely examination of the second component of $\vec{V}$, whereas the automaton $\texttt{or'}_1$ does not attempt to examine the undefined component of a matching vector.

4.3 Which directions ?

As suggested by the example of the function `or'`, the correctness of the $\mathcal{D}$ scheme will rely on following directions of the compiled matching, just as in the case of the $\mathcal{C}$ scheme. However, a complication arises here: in a call $\mathcal{D}(\vec{x}, (P), (E), d)$ the argument pattern matrix (P) may not fully account for the currently compiled matching. This difficulty appears when (P) has been produced by previously cutting the pattern matrix argument. I now illustrate it on two examples.

First, consider the following matching definition:

$$(P) = \begin{pmatrix} \texttt{true} & \texttt{true} & \texttt{true} \\ \texttt{_} & \texttt{_} & \texttt{false} \\ \texttt{_} & \texttt{_} & \texttt{_} \end{pmatrix}, \quad (E) = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$$

Then, consider the compilation $\mathcal{D}((x\ y\ z), (P), (E), \Omega)$. Assuming that the first simple matching is performed on `z`, we get:

```
match z with
  true  → D((x y), ( true true ), ( 1 ), default)
| false → 2
| _     → 3
```

Let (P') be (`true true`). Obviously, there is no vector matching (P') with one Ω component. In other words, it is safe to examine either `x` or `y` first to recognize the vectors $(x\ y)$ that match (P'). Nevertheless, this is no longer the case when the whole matching is considered, since (Ω `false true`) and (`false` Ω `true`) do match the third row of (P). More precisely, not only must the subautomaton $\mathcal{D}((x\ y), (P'), (1), \texttt{default})$ recognize the vectors (V_1 V_2 `true`) that match the first row of (P), but it also has to recognize the vectors (V_1 V_2 `true`) that are incompat-

ible with the first row of (P). In termzs of pattern matrices, all the rows that may still be matched have to be considered; that is, we must compute the directions of the following matrix:

$$(P'') = \begin{pmatrix} \texttt{true} & \texttt{true} \\ _ & _ \end{pmatrix}$$

We get here $\mathrm{Dir}[(P'')] = \emptyset$, thereby confirming that neither `x` nor `y` can be examined safely. Observe that (P'') is exactly the submatrix that the $\mathcal{C}$ scheme would use in this situation: the matrix (P'') such that a vector $(V_1\ V_2)$ matches a row of (P'') if and only if the vector $(V_1\ V_2\ \texttt{true})$ matches a row of (P).

Now, consider another matching definition:

$$(P) = \begin{pmatrix} \texttt{true} & \texttt{true} & _ \\ \texttt{false} & \texttt{false} & \texttt{true} \\ _ & _ & \texttt{true} \end{pmatrix}, \quad (E) = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$$

Suppose that we first apply the inductive step 2-(a) to the third column of (P), obtaining the following partition:

$$(Q) = \begin{pmatrix} \texttt{true} & \texttt{true} \end{pmatrix}, \quad (F) = \begin{pmatrix} 1 \end{pmatrix}$$

and

$$(R) = \begin{pmatrix} \texttt{false} & \texttt{false} & \texttt{true} \\ _ & _ & \texttt{true} \end{pmatrix}, \quad (G) = \begin{pmatrix} 2 \\ 3 \end{pmatrix}$$

This partition leads to the following inductive calls:

$$\mathcal{D}((x\ y\ z), (P), (E), \Omega) = \mathcal{D}((x\ y), (Q), (F), \mathcal{D}((x\ y\ z), (R), (G), \Omega))$$

Then, applying the inductive step (2)-b to the third column of (R), we get:

$$\mathcal{D}((x\ y\ z), (R), (G), \Omega) = \begin{cases} \texttt{match z with} & & \\ \quad \texttt{true} & \to & \mathcal{D}((x\ y), (R'), \begin{pmatrix} 2 \\ 3 \end{pmatrix}, \texttt{default}) \\ \mid\ _ & \to & \Omega \end{cases}$$

where

$$(R') = \begin{pmatrix} \texttt{false} & \texttt{false} \\ _ & _ \end{pmatrix}$$

In order to compute $A' = \mathcal{D}((x\ y\ z), (R), (G), \texttt{default})$ further, we now need to emit a simple matching either on `x` or on `y`. But neither 1 nor 2 are directions of (R'), since the vectors $(\Omega\ \texttt{true})$ and $(\texttt{true}\ \Omega)$ match the second row of (R'). However, if we consider the whole matching we are compiling, we immediately see that $(\Omega\ \texttt{true}\ \texttt{true})$ and $(\Omega\ \texttt{true}\ \texttt{true})$ do not match the third row of (P), since these two vectors are not defined enough to be incompatible with the first row of (P). In other words, the input to $\mathcal{D}((x\ y), (R'), \ldots)$ is not just any value vector: it can be characterized briefly as the set of value vectors that have been recognized to be incompatible with the rows of (Q) by the subautomaton $\mathcal{D}((x\ y), (Q), \ldots)$. In terms of directions, we have to consider the directions towards the second and third row of the matrix (R''), which is made by appending the two rows of (R') to

the single row of (Q):

$$(R'') = \begin{pmatrix} \texttt{true} & \texttt{true} \\ \texttt{false} & \texttt{false} \\ \texttt{_} & \texttt{_} \end{pmatrix}$$

Here also, observe that (R'') is exactly the pattern matrix that the $\mathcal{C}$ scheme would use. That is, a vector $(V_1\ V_2)$ matches a row of (R''), if and only if the vector $(V_1\ V_2\ \texttt{true})$ matches a row of (P).

4.4 Correctness

The intuition given by the examples of the previous section can be summarized as follows: directions should not be compiled on the matrix (P) but on another matrix that accounts for the totality of the matching being compiled. A tentative solution is thus for the full $\mathcal{D}$ scheme to take an extra pattern matrix (Q) as argument. This pattern matrix should not be cut during recursive calls, it should only be decomposed in the same way the $\mathcal{C}$ scheme decomposes its main argument.

In practice, it is convenient to partition the matrix (Q) into two matrices (P') and (P''), in order to make a distinction between the rows that are already known to be incompatible with the currently examined values and the rows that can still be reached. Moreover, because the rule 2-(a) may delete a column in (P) without examining the corresponding input value, a discrepancy in width may appear between (P) on the one hand, and (P') and (P'') on the other hand. Finally, the full $\mathcal{D}$ compilation scheme takes three extra arguments and a typical call to $\mathcal{D}$ is now written $\mathcal{D}(\vec{x}, (P), (E), \vec{x}\,', (P'), (P''), d)$. The new arguments are the variable vector $\vec{x}\,'$ of size n', the pattern matrix (P') of size $n' \times m'$ and the pattern matrix (P'') of size $n' \times m''$. In some sense, made precise in the safety conditions below, $\vec{x}\,'$ include $\vec{x}$ and (P'') includes (P).

Definition 8. Consider a call $\mathcal{D}(\vec{x}, (P), (E), \vec{x}\,', (P'), (P''), d)$. Two safety conditions are defined as follows:

(1) The vector $\vec{x}$ is a subvector of the vector $\vec{x}\,'$. That is, there exists an injective mapping $\mathcal{I}$ from $[1 \dots n]$ into $[1 \dots n']$ such that, for any integer i in $[1 \dots n]$, we have $x_i = x'_{\mathcal{I}(i)}$.

(2) The matrix (P) is a submatrix of the matrix (P''). That is, we have $m \leq m''$ and, for any column index i in $[1 \dots n]$ and any row number j in $[1 \dots m]$, we have $p_i^j = p''^j_{\mathcal{I}(i)}$. Furthermore, for any row number j in $[1 \dots m]$ and any column index i'' that is not the image of some integer by the mapping $\mathcal{I}$, the pattern $p''^j_{i''}$ is a variable.

As we shall see later in the correctness proof, the main goal of the safety condition is to guarantee that any vector matching the row number i in (P'') (where $i < m$) also matches the row number i in (P).

The initial call to $\mathcal{D}$ is now written as follows:

$$\mathcal{D}((x), \begin{pmatrix} p^1 \\ p^2 \\ \vdots \\ p^m \end{pmatrix}, \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_m \end{pmatrix}, (x), (), \begin{pmatrix} p^1 \\ p^2 \\ \vdots \\ p^m \end{pmatrix}, \Omega)$$

Now, we return to the inductive cases of the definition of $\mathcal{D}$ (section 4.2), still making the arbitrary choice of examining the first column of (P).

2-(a) If p_1^1 is a variable, then we get the following two nested inductive calls:

$$\mathcal{D}(\vec{y},(Q),\begin{pmatrix}e_1\\ \vdots\\ e_k\end{pmatrix},\vec{y}\,',(Q'),(Q''),\mathcal{D}(\vec{z},(R),\begin{pmatrix}e_{k+1}\\ \vdots\\ e_m\end{pmatrix},\vec{z}\,',(R'),(R''),d))$$

where we have (k being the highest row number such that p_k^1 is a variable):

$$\vec{y}=(x_2\dots x_n)\qquad (Q)=\begin{pmatrix}p_2^1\cdots p_n^1\\ \vdots\\ p_2^k\cdots p_n^k\end{pmatrix}\qquad (R)=\begin{pmatrix}p_1^{k+1}\cdots p_n^{k+1}\\ \vdots\\ p_1^m\quad\cdots p_n^m\end{pmatrix}$$

The new arguments for the outermost call are easily defined, we take $\vec{y}\,'=\vec{x}\,'$, $(Q')=(P')$ and $(Q'')=(P'')$. As to the innermost call, we first take $\vec{z}\,'=\vec{x}\,'$. Then, the matrix (R'') is built by removing the k top rows of (P''). Finally, these row are added at the bottom of (P'), yielding the matrix (R').

$$(R')=\begin{pmatrix}{p'}_1^1 & {p'}_2^1 & \cdots & {p'}_{n'}^1\\ & \vdots & & \\ {p'}_1^{m'} & {p'}_2^{m'} & \cdots & {p'}_{n'}^{m'}\\ {p''}_1^1 & {p''}_2^1 & \cdots & {p''}_{n'}^1\\ & \vdots & & \\ {p''}_1^k & {p''}_2^k & \cdots & {p''}_{n'}^k\end{pmatrix}\qquad (R'')=\begin{pmatrix}{p''}_1^{k+1} & {p''}_2^{k+1} & \cdots & {p''}_{n'}^{k+1}\\ & \vdots & & \\ & & & \\ {p''}_1^{m''} & {p''}_2^{m''} & \cdots & {p''}_{n'}^{m''}\end{pmatrix}$$

2-(b) If p_1^1 has a root constructor, let $i'=\mathcal{I}(1)$ be the component index in $\vec{x}\,'$ such that $x'_{i'}=x_1$. Let also c_1, c_2, …, c_z be the root constructors appearing in the first row of (P). Without loss of generality, I describe the inductive call $\mathcal{D}(\vec{y},(Q_1),\vec{y}\,',(Q'_1),(Q''_1),\texttt{default})$, which is to be performed when constructor c_1 has been recognized (see case 2-(b) in section 4.2). Matrices (P') and (P'') are decomposed according to their column number i', in order to extract pattern rows whose pattern number i' can match values with c_1 as root constructor (see section 3.2, case 3, for a complete description of this decomposition procedure). This process yields the new matrices (Q'_1) and (Q''_1). Finally we take $\vec{y}\,'=(y_1\dots y_{a_1}\ {x''}_1\dots {x''}_{i'-1}\ x''_{i'+1}\dots x''_{n''})$.
As to the inductive call $\mathcal{D}(\vec{x},(R),(F),\vec{x},(R'),(R''),d)$, which is to be performed in the default case, the matrices (R') and (R'') are the same as the ones for the innermost inductive call in case 2-(a) just above. That is, the first k rows of (P') are transferred at the bottom of (P'') (remember that here, the integer k is the highest row number such that p_k^1 is not a variable).

One easily checks that the safety conditions are met in the initial call. To see that they are preserved across inductive calls, first notice that, whenever some rows are removed from the top of (P), the same number of rows are removed from the top of (P'') and that whenever a column is removed from (P) and not from (P''), this column consists solely of variables. Furthermore, in step 2-(b), (P) and (P'') are decomposed according to the columns number 1 and $i'=\mathcal{I}(1)$ respectively and

we know from the safety conditions that p_1^j equals $p''^j_{i'}$ for any row number j valid both in (P) and (P''), so that (Q_1), (Q_2), ..., (Q_z) are submatrices of (Q''_1), (Q''_2), ..., (Q''_z).

The correctness property is simple to formulate: a compilation is correct when it produces an automaton $A = \mathcal{D}(\vec{x}, (P), (E), \ldots, \Omega)$, whose result in environments $\sigma = [x_1 \backslash V_1, \ldots, x_n \backslash V_n]$ and ρ equates the value $match[(P), (E)](\vec{V})$. The correctness proof will rely on a stronger property and requires a few new concepts and notations.

NOTATION 4.1. *Let (P') and (P'') be two pattern matrices of same width n' and of respective heights m' and m''. The new matrix $(P'@P'')$ is defined as the $n' \times (m' + m'')$ matrix built by concatenating the rows of (P') and the rows of (P''). Furthermore a* direction *towards (P'') in $(P'@P'')$ is any column index d, such that there exists no value vector $\vec{V}$ whose component V_d is the undefined term Ω and that matches some row in $(P'@P'')$ whose number is strictly greater than m'.*

At any compilation stage, we write $A = \mathcal{D}(\vec{x}, (P), \vec{x}\,', (P'), (P''), d)$ for the produced automaton and consider some arbitrary value environment σ and arbitrary control environment ρ. Then, we define the following two strong correctness properties:

(α) If there exists some integer i in $[1 \ldots m]$, such that $\sigma(\vec{x}\,')$ matches row number $m' + i$ in $(P'@P'')$ —and then we know that there exists some substitution ϕ, such that $\phi(\vec{p}^{\,''i}) = \sigma(\vec{x}\,')$ and $match[(P), (E)](\sigma(\vec{x})) = \phi(e_i)$—, then the automaton state (ρ, σ, A) leads to the final state (ρ', σ', e'_i) and we have the equality $\sigma'(e'_i) = \phi(e_i)$.

(β) If there exists some integer i in $[m + 1 \ldots m'']$, such that $\sigma(\vec{x}\,')$ matches row number $m' + i$ in $(P'@P'')$, then the automaton state (ρ, σ, A) leads to the intermediate state (ρ, σ', d).

It should be noticed that the correctness condition (α) uses the safety conditions, when it implicitly asserts that any vector matching the row number $m' + i$ in $(P'@P'')$ with $i \leq m$ also matches the row number i in (P).

LEMMA 4.1. *Any call $\mathcal{D}(\vec{x}, (P), \vec{x}\,', (P'), (P''), d)$ obeys the correctness properties (α) and (β), provided that compilation follows some direction towards (P'') in $(P'@P'')$ at every critical compilation step 2-(b).*

PROOF. Proof is by induction on the definition of $\mathcal{D}$.

(1) If (P) is empty ($n = 0$, $m = 0$), then the output automaton A reduces to the failure continuation d and the intermediate state (ρ, σ, d) is already reached, regardless of σ. Therefore, one just has to check that case (α) is not possible here, which is obvious since an empty matrix does not even have a single row to match. If (P) has at least one empty row ($n = 0$, $m > 0$), then we get $A = e_1$ and the final state (ρ, σ, e_1) is already reached, regardless of σ. Moreover, the vector $\sigma(\vec{x}\,')$ cannot match any row in (P'') but the first, because the first row in (P'') is made of variables only (by the safety condition). Therefore, cases (α) with $i > 1$ and (β) are not possible.

(2) If (P) possesses at least one pattern, then assume —for instance— that compilation progresses by decomposing (P) along its first row and that we have $\mathcal{I}(1) = 1$ (that is, $x'_1 = x_1$).

(a) If p_1^1 is a variable y^1, then we have the following nested calls, (see case 2-(a) in the previous descriptions of $\mathcal{D}$):

$$\mathcal{D}(\vec{y},(Q),(F),\ldots,\mathcal{D}(\vec{x},(R),(G),\ldots,d))$$

Let then σ be an environment such that $\sigma(\vec{x}\,')$ matches row number $m'+i$ ($1 \leq i \leq m''$) in $(P'@P'')$. Let ϕ be a substitution such that $\phi(\vec{p}\,''^{\,i}) = \sigma(\vec{x}\,')$. There are three possible cases:
If $i \leq k$ —remember that k is the smallest row number such that p_{k+1}^1 is not a variable—, then the transition $(\rho,\sigma,A) \hookrightarrow^* (\rho',\sigma',f_i')$ holds, where we have $\sigma'(f_i') = \phi(f_i)$, by direct application of the hypothesis (α) to the outermost call $\mathcal{D}(\vec{y},(Q),(F),\ldots)$. Moreover, we have $f_i = e_i[y^i \backslash x_1]$ (by definition of $\mathcal{D}$), $\sigma'(x_1) = \sigma(x_1)$ (environments can only grow during transitions), and $\sigma(x_1) = \phi(y^i)$ (by hypothesis). Therefore, we finally get $\sigma'(f_i') = \phi(e_i)$.
If i belongs to $[k+1\ldots m]$, then, by application of the induction hypothesis (β) to the outermost call $\mathcal{D}(\vec{y},(Q),(F),\ldots)$, followed by an application of the induction hypothesis (α) to the innermost call $\mathcal{D}(\vec{x},(R),(G),\ldots)$, we get the transitions:

$$(\rho,\sigma,A) \hookrightarrow^* (\rho,\sigma',\mathcal{D}(\vec{x},(R),(G),\ldots,d)) \hookrightarrow^* (\rho'',\sigma'',g_i')$$

By induction hypothesis, we also have $\phi(g_i) = \sigma''(g_i')$. Thus, we directly get: $\phi(e_i) = \sigma''(g_i')$, since $g_i = e_i$ by construction of the vector (G).
Finally, if i is strictly greater than m, then, by a double application of the hypothesis (β), we get:

$$(\rho,\sigma,A) \hookrightarrow^* (\rho,\sigma',\mathcal{D}(\vec{x},(R),(G),\ldots,d)) \hookrightarrow^* (\rho,\sigma'',d)$$

(b) If p_1^1 has a root constructor, then, keeping notations from the definition of $\mathcal{D}$, consider an environment σ, such that $\sigma(\vec{x}\,')$ matches the row number $m'+i$ in $(P'@P'')$. It is important to notice that, by hypothesis, the index 1 is a direction in $(P'@P'')$, so that $\sigma(x_1)$ has a root constructor c.
If i belong to $[1\ldots k]$, then the constructor c must be recognized by the elementary matching emitted at this compilation stage — let c be c_1 for instance. By definition of one-step transitions and application of the hypothesis (α) to the call $A_1 = \mathcal{D}(\ldots,(Q_1),(E_1),\ldots,\texttt{default})$, we get the transitions:

$$(\rho,\sigma,A) \hookrightarrow (\rho'',\sigma'',A_1) \hookrightarrow^* (\rho',\sigma',e_i')$$

Where $\phi(e_i) = \sigma'(e_i')$.
If i belongs to $[k+1\ldots m]$, then there are two subcases. If c is recognized by the elementary test emitted here —assume $c = c_1$ for instance—, then the correctness result follows by a one-step transition (constr), an application of the hypothesis (β) to the call $\mathcal{D}(\ldots,(Q_1),(E_1),\ldots,\texttt{default})$, a one-step transition (raise) and an application of the hypothesis (α) to $\mathcal{D}(\vec{x},(R),(F),\ldots,d)$. Otherwise, we get correctness more directly, by a one-step transition (default) and an application of the hypothesis (α) to the call $\mathcal{D}(\vec{x},(R),(F),\ldots,d)$.
If i is strictly greater than m, then the reasoning is exactly as above, but for the last step which is replaced by an application of the hypothesis (β) to the default inductive call $\mathcal{D}(\vec{x},(R),(F),\ldots,d)$.

□

The lemma above can be summarized as follows: matching vectors are recognized by the automata that the $\mathcal{D}$ scheme produces when it follows directions. In order to prove correctness, we need to show the converse property: when the automata produced by the $\mathcal{D}$ scheme recognize a vector, this vector is a matching vector.

Therefore, we consider the following properties (α') and (β'). At any compilation stage, we write $A = \mathcal{D}(\vec{x}, (P), \ldots, d)$ for the produced automaton and we consider some arbitrary value environment σ and control environment ρ. Then, we state:

(α') If the execution of A in the initial environments σ and ρ succeeds, then $\sigma(\vec{x}\,')$ matches some row number i in (P) and we get:

$$\mathcal{RUN}(\rho, \sigma, \sigma(\vec{x})) = match[(P), (E)](\sigma(\vec{x}))$$

(β') If the transition $(\rho, \sigma, A) \hookrightarrow^* (\rho, \sigma', d)$ holds, then $\sigma(\vec{x})$ is incompatible with all the rows of (P).

LEMMA 4.2. *Any call $\mathcal{D}(\vec{x}, (P), \ldots, d)$ obeys the properties (α') and (β'), regardless of the choice of discriminating indices.*

PROOF. The proof is a simple induction over the definition of $\mathcal{D}$. Notice that it does not need the extra arguments $\vec{x}\,'$, (P') and (P''). □

PROPOSITION 4.1. *The $\mathcal{D}$ scheme is correct, provided it follows a direction towards (P'') in $(P'@P'')$ at every step 2-(b).*

PROOF. Direct corollary of the properties (α), (β), (α') and (β'). □

5. BUILDING A CORRECT AND EFFICIENT PATTERN MATCHING COMPILER

The ultimate goal of this paper is to introduce a pattern matching compiler that outputs a correct automaton whenever there exists one. Obviously, a first step is to compute directions, the next section shows how to achieve this intermediate goal.

5.1 Efficient computation of directions

Given a pattern matrix (P), a row number i and a column index d, we want to know whether d is a direction towards row number i in (P) or not. This decision problem can be expressed as the unused match case detection problem: is the last row of the matrix $(Q_{(d,i)})$ below satisfiable or not? That is, does there exist a value vector $\vec{V}$ such that the predicate $match_i[(Q_{(d,i)})]$ holds on $\vec{V}$? The matrix $(Q_{(d,i)})$ is the submatrix of (P) obtained by deleting the column d and the rows after row i:

$$\left(Q_{(d,i)}\right) = \begin{pmatrix} p_1^1 \ldots p_{d-1}^1 & p_{d+1}^1 \ldots p_n^1 \\ p_1^2 \ldots p_{d-1}^2 & p_{d+1}^2 \ldots p_n^2 \\ & \vdots \\ p_1^i \ldots p_{d-1}^i & p_{d+1}^i \ldots p_n^i \end{pmatrix}$$

LEMMA 5.1. *Let (P) be a pattern matrix. Let i be a row number and d be a column index in (P). Let $(Q_{(d,i)})$ be as described above. Then, the index d is not a direction for the matching by the row number i in matrix (P), if and only if the pattern p_d^i is a variable and the last clause of matrix $(Q_{(d,i)})$ is satisfiable.*

PROOF. In the case where $p_d^i = c\ q_1 \ldots q_a$ is not a variable, then any value vector $(V_1\ V_2 \ldots V_n)$ that matches the row number i in matrix (P) is such that V_d is an

instance of $c\ q_1 \ldots q_a$ and we have $V_d \succ \Omega$. Otherwise, let $V_1, \ldots, V_{d-1}, V_{d+1}, \ldots, V_n$ be any $n-1$ partial values. The following equality between matching predicates can be shown by expanding definitions:

$$match_i[(P)](V_1 \ldots V_{d-1}\ \Omega\ V_{d+1} \ldots V_n) = match_i[(Q_{(d,i)})](V_1 \ldots V_{d-1}\ V_{d+1} \ldots V_n)$$

□

The following algorithm solves the unused match case detection problem in the general case. Given a pattern matrix (P), of size $n \times m$, it computes the truth value of the formula $\mathcal{F}(P) = \exists \vec{V}\ match_m[(P)](\vec{V})$. This algorithm closely follows the compilation scheme $\mathcal{C}$ itself (see section 3.2):

(1) If the first row of (P) contains only variables (and this includes the case where $n = 0$), then the value of $\mathcal{F}(P)$ depends on the number m of rows in (P). If $m = 1$, then $\mathcal{F}(P)$ is true, since any instance of $\vec{p}^{\,1}$ matches the last (and only) row of matrix (P). Otherwise, $\mathcal{F}(P)$ is false.

(2) In all the other cases, let us choose a column index. Assume that the index 1 is chosen. Let $\Sigma = \{c_1, c_2, \ldots, c_z\}$ be the set of the root constructors of the patterns in the first column of (P). To each constructor c_k in Σ, a new pattern matrix (P_k) is associated, exactly as it is done in the $\mathcal{C}$ scheme. If the set Σ is not a full signature or if Σ is empty, then a default matrix (P_d) is also considered. There are two subcases:

(a) If $p_1^m = x$, then let $\vec{V}$ be a value vector that matches the last row of (P). If V_1 has a root constructor, then the matching by matrix (P) is equivalent to the matching by one of the matrices (P_k) or (P_d). Otherwise, if $V_1 = \Omega$, then, because the matching predicate is monotonic, any value vector $\vec{U} = (U_1\ V_2 \ldots V_n)$ such that $U_1 \succ \Omega$ matches row m as $\vec{V}$ does. Therefore, $\mathcal{F}(P)$ is true, if and only at one of the formulas $\mathcal{F}(P_1)$, $\mathcal{F}(P_2)$, $\ldots$, $\mathcal{F}(P_z)$ or $\mathcal{F}(P_d)$ is.

(b) If $p_1^m \succ x$, then let c_k be the root constructor of p_1^m. We easily show the equality $\mathcal{F}(P) = \mathcal{F}(P_k)$.

Regarding the efficiency of this algorithm, it can be observed that the number of calls to function $\mathcal{F}$ is bounded by the number of calls to function $\mathcal{C}$, when compilation is done by making the same choices at critical steps. As shown by the example given in [Maranget 1992a; Sekar, Ramesh, and Ramakrishnan 1992], the number of calls to function $\mathcal{C}$ can be exponential in the size of (P). Although I do not know whether this upper bound is indeed reached or not in the worst cases, experiments showed that a naive implementation of function $\mathcal{F}$ may lead to important computations. Fortunately, this misbehavior can be avoided in practice by using the following three heuristics:

(1) The matrix (P) itself can be reduced. Let $\vec{p}^{\,i}$ and $\vec{p}^{\,j}$ be two rows in (P), such that $i < m$, $j < m$ and $\vec{p}^{\,i} \preceq \vec{p}^{\,j}$. For any value vector $\vec{V}$ such that $\vec{p}^{\,i} \# \vec{V}$, we necessarily have $\vec{p}^{\,j} \# \vec{V}$, by definition of the compatibility relation. That is, the pattern vector $\vec{p}^{\,j}$ is useless for the computation of $\mathcal{F}(P)$ and the matrix (P) can be simplified by only retaining the pattern rows that are minimal for the definition ordering. This simplification of matrix (P) is particularly worthwhile when some pattern rows contain a lot of variables.

(2) When there is a default matrix (P_d), it is tested first. This amounts to making the assumption that, if there exists a value vector that matches the last row of (P), then its components are likely not to appear inside matrix (P).

(3) I also attempt to minimize the size and number of the matrices (P_1), (P_2), $\ldots$,(P_z), by a good choice of the column to examine at step 2-(a). For each column, characterized by its index i, let $z(i)$ be the number of different root constructors in column i and $v(i)$ be the number of variables in column i. Let then $r(i)$ be the total number of rows in the matrices $(P_1), (P_2), \ldots, (P_{z(i)})$, we have $r(i) = z(i) \times v(i) + m$ or $r(i) = z(i) \times v(i) + m - v(i)$, depending on whether there is a default matrix (P_d) or not. I select a column with a minimal $r(i)$. If there are several columns such that $r(i)$ is minimal, then I favor one with a minimal number of different root constructors $z(i)$. Another, simpler, choice is first to minimize $v(i)$ (in a attempt to limit row duplication) and then to minimize $z(i)$ (in an attempt to limit the number of matrices generated). Other size measures have been tested, including the surface of matrices (number of rows $\times$ number of columns). Choosing a good measure is not easy and this heuristic is less efficient than the two others.

Regarding the efficiency of the computation of the set of directions $Dir(P)$, it is usually not necessary to compute all the $Dir_i(P)$ sets. First, if there is a column in (P) which contains no variable, then, by lemma 5.1, the index of this column is a direction. Such a direction is an *obvious direction* and knowing just one direction is enough to apply the compilation schemes. Otherwise, there is no obvious direction in (P) and the set $Dir(P)$ is computed. The cost of this computation can be lowered by the following observation: if some index d does not belong to $Dir_i(P)$, then, for any other row number j, we need not check whether index d belongs to $Dir_j(P)$ or not, since we already know that d is not a direction for the whole matrix (P). Of course, the $Dir_i(P)$ sets are examined following increasing row numbers i, so that index checkings are avoided when matrices $Q_{(d,i)}$ are large. That is, we compute $\mathcal{D}_m = Dir(P)$, where $\mathcal{D}_1 = Dir_1(P)$ and $\mathcal{D}_{i+1} = \{\, d \in \mathcal{D}_i \mid d \in Dir_{i+1}(P) \,\}$. If matrix (P) has no direction, then there exists a row number max, such that $\mathcal{D}_{max} \neq \emptyset$ and $\mathcal{D}_{max+1} = \emptyset$. The column indices in $\mathcal{D}_{max}$ are called *partial directions*. Partial directions will be used later in section 6, in the process of "not so incorrectly" compiling pattern matching definition that cannot be compiled correctly.

Previous approaches to lazy pattern matching compilation involve the explicit computation of the set of value vectors matching the rows of matrix (P). Let $\mathcal{M}$ be this set. We have $\mathcal{M} = \bigcup_{i=1}^{m} \mathcal{M}_i$, where $\mathcal{M}_i = \{\vec{V} \mid match_i[(P)](\vec{V})\}$. In Laville [1991] the set $\mathcal{M}$ is described by its *minimal generators*, that is, by the subset of its least defined elements. In Puel and Suárez [1993] each set $\mathcal{M}_i$ is represented by a normalized constrained pattern that can be seen as the disjunctive normal form of the following characteristic proposition:

$$\mathcal{X}_i(V_1, V_2, \ldots V_n) = (\bigwedge_{j=1}^{i-1} \bigvee_{k=1}^{n} p_k^j \,\#\, V_k) \bigwedge (\bigwedge_{k=1}^{n} p_k^i \preceq V_k)$$

Direct implementation of these two representations for the set $\mathcal{M}$ leads to data structures whose size grow exponentially with the size of the input matrix (P). Our approach, by directly computing directions, avoids such an exponential space behavior.

Finally it is worth noticing that the algorithm proposed here can be used to solve two classical problems on pattern matching, the so-called exhaustiveness and redundancy problems. The redundancy problem is just another name for the unused clause problem. The exhaustiveness clause problem can be formulated as follows: given a set of clauses $E = \{p^1 \colon e_1, p^2 \colon e_2, \ldots, p^m \colon e_m\}$, does there exist a value V that does not match any clause in E? In other words, does there exists a value V which is incompatible with all the p^i's? One easily sees that the exhaustiveness of the matching by E is equivalent to the satisfiability of the last clause of the matching by $E' = \{p^1 \colon e_1, p^2 \colon e_2, \ldots, p^m \colon e_m, p^{m+1} \colon e_{m+1}\}$ where p^{m+1} is a variable.

5.2 The compilation algorithm

Both correctness criteria (propositions 3.1 and 4.1) apply at compile-time, so that incorrect automata are ruled out before they are completely built. For instance, when compiling the `G` function given in the introduction, the $\mathcal{C}$ and $\mathcal{D}$ compilation scheme find that this function cannot be compiled correctly as soon as the first call to the compilation scheme, since this first call is performed with a (P) matrix that has no directions.

The existence of these compile-time correctness criteria is not enough, though. If a matrix with many directions is discovered, the search for a correct automaton *a priori* necessitates to try all the compilations that follow each of these possible directions. Namely, if, from that stage, compilation goes on following a first direction and that a matrix without a direction is discovered later, we cannot be sure that such a situation would also have occurred, if another direction had been selected in the first place. Therefore, to detect non-compilable matchings, a straightforward compilation algorithm should check all existing directions. Such a technique can be extremely expensive. The key idea for solving this problem is that backtracking as described above is useless, because discovering a matrix without a direction at any compilation stage characterizes a pattern matching function that cannot be implemented correctly.

More specifically, the matching functions that can be implemented correctly are described by invoking the *sequentiality* theory, as introduced by Kahn and Plotkin [1978] and presented by Curien [1986] or Huet and Lévy [1991][Part II]. From here, I assume some familiarity with basic notions on terms such as *occurrences* and *subterms*. The subterm located at occurrence u in some term M is written M/u.

Definition Sequential functions [Kahn and Plotkin 1978]. Let $\mathcal{P}$ be a monotonic predicate over partial values. An occurrence u is an *index* of $\mathcal{P}$ in U, if and only if we have $U/u = \Omega$ and that, for all V such that $V \succeq U$, $\mathcal{P}(V)$ holds implies implies $V/u \succ \Omega$. Then, $\mathcal{P}$ is sequential at U, such that $\mathcal{P}(U)$ does not hold, if and only if whenever there exists $V \succeq U$ such that $\mathcal{P}(V)$ holds, it follows that there exists an index of $\mathcal{P}$ in U. Finally, $\mathcal{P}$ is sequential, if and only if it is sequential at every partial value on which $\mathcal{P}$ does not hold.

Our first step is to introduce matched values in the compilation process, and not only matched vectors. Thus, the $\mathcal{D}$ compilation scheme will now take an extra argument. This new argument N is a linear term and appears in first position. At any compilation stage, it is to be understood as the already *explored prefix* of the examined value. Initially, nothing has been explored yet and the prefix N is a

simple variable. The first call to $\mathcal{D}$ is as follows:

$$\mathcal{D}(x,(x),\begin{pmatrix}p^1\\p^2\\\vdots\\p^m\end{pmatrix},\begin{pmatrix}e_1\\e_2\\\vdots\\e_m\end{pmatrix},(x),(),\begin{pmatrix}p^1\\p^2\\\vdots\\p^m\end{pmatrix},\Omega)$$

Then, the argument N is inductively computed as follows (see the inductive definition of $\mathcal{D}$, that is, the cases 2-(a) and 2-(b) in sections 4.2 and 4.4):

2-(a) If p_1^1 is a variable, then we get the following two nested inductive calls:

$$\mathcal{D}(N,(x_2\dots x_n),(Q),(F),\vec{x}\,',(P'),(P''),\mathcal{D}(N,\vec{x},(R),(G),\vec{x}\,',(R'),(R''),d))$$

Observe that N does not change here, since the examination of the input term does not progress here.

2-(b) If p_1^1 has a root constructor, then let us assume that a simple matching of the variable $x_1 = x'_1$ is emitted and let c_1, c_2, …, c_z be the recognized constructors. Let us then consider —for instance— the inductive call to be performed when constructor c_1 has been recognized:

$$\mathcal{D}(N[x'_1\backslash c_1\ y_1\dots y_{a_1}],(y_1\dots y_{a_1}\ x_2\dots x_n),(Q_1),\dots,\texttt{default})$$

As to the default inductive call, it now yields:

$$\mathcal{D}(N,\vec{x},(R),(F),\vec{x}\,',(R'),(R''),d)$$

In the following two lemmas, we consider a set of clauses $E = \{q^1\colon e_1,\dots,q^k\colon e_k\}$ and an arbitrary call $\mathcal{D}(N,\vec{x},(P),(E),\vec{x}\,',(P'),(P''),d)$ occurring during the compilation of the matching by E.

LEMMA 5.2. *For any substitution σ, the following equivalence holds:*

The vector $(\sigma(x'_1)\ \sigma(x'_2)\dots\sigma(x'_{n'}))$ matches some row in $(P'@P'')$

$$\Updownarrow$$

The partial value $\sigma(N)$ matches some clause in E

PROOF. The proof is by induction on the definition of $\mathcal{D}$. The key point is that the rows of $(P'@P'')$ are exactly the subpatterns q_1^i, q_2^i, …, $q_{n'}^i$ of one of the initial patterns q^i that admits N as prefix, (that is, $q^i = N[x'_1\backslash q_1^i,\ x'_2\backslash q_2^i,\dots,x'_{n'}\backslash q_{n'}^i]$). □

LEMMA 5.3. *At any inductive step during the compilation, consider a column index d in the matrix $(P'@P'')$. Let u_d be the occurrence of the variable x'_d in the prefix N. Then, u_d is an index for the matching predicate defined by E, if and only if d is a direction in $(P'@P'')$.*

PROOF. This lemma is a quite direct consequence of the previous lemma and of the fact that the variables of the linear term N are exactly the components of $\vec{x}\,'$. □

COROLLARY 5.1. *If some matrix $(P'@P'')$ without a direction is discovered during compilation of the clause set E, then the matching predicate defined by E is not sequential.*

Our second step is to show that the the recognition of matching values by an automaton defines a sequential predicate. Now, consider an extended version of the transition rules of automata with failures. An extended automaton state is a four-tuple (N, ρ, σ, A), where N is a linear term. Initially, the term N is the same single variable as the one given as first argument to the compilation function $\mathcal{D}$. Later on, during transitions, the term N records the part of the input term that has already been examined, in the sense made clear by the following extended transition rules:

If $\sigma(x) = c_k\ V_1 \ldots V_{a_k}$,
$(N, \rho, \sigma, \texttt{match}\ x\ \texttt{with}\ \ldots\ |\ c_k\ x_1 \ldots x_{a_k}\ \rightarrow\ A_k\ \ldots\ |\ _\ \rightarrow\ A_d) \hookrightarrow\quad \ldots$
$\ldots\quad (N[x \backslash c_k\ x_1 \ldots x_{a_k}], <\sigma, A_d>.\rho, \sigma \cup [x_1 \backslash V_1, \ldots, x_{a_k} \backslash V_{a_k}], A_k)$ (constr)

If $\sigma(x) = c\ V_1 \ldots V_a$ where c is not recognized,
$(N, \rho, \sigma, \texttt{match}\ x\ \texttt{with}\ \ldots\ |\ _\ \rightarrow\ A_d) \hookrightarrow (N[x \backslash c\ \Omega \ldots \Omega], \rho, \sigma, A_d)$ (default)

$$(N, <\phi, A>.\rho, \sigma, \texttt{default}) \hookrightarrow (N, \rho, \phi, A) \qquad \text{(raise)}$$

Now, consider some automaton A, which is a simple matching of the variable x. To compute the result of the execution of A on the value U, we just compute the complete transition starting from the initial state $(x, (), [x \backslash U], A)$, thereby reaching some final state (N', ρ', σ', A'). Then, we say that U has been recognized, if and only if A' is some clause right-hand side. Observe that, at any intermediate state $(N'', \rho'', \sigma'', A'')$, we have $U \succeq N''_\Omega$. More precisely, for any variable x of occcurence u_x in N'', we have $U/u_x = \sigma''(x)$. Also observe that the recognized prefix increases at each one-step transition.

LEMMA 5.4. *The execution of the automaton A on the value U defines a sequential predicate.*

PROOF. Consider some value U which is not recognized by A. There exists a complete transition $(x, (), [x \backslash U], A) \hookrightarrow^* (N', \rho', \sigma', A')$, where A' is not a clause right-hand side. Observe that the execution of A on any value $V \succeq U$ will yield the following complete transition $(x, (), [x \backslash V], A) \hookrightarrow^* (N', \tau', \phi', A') \hookrightarrow^* (N'', \tau'', \phi'', A'')$, because N' is also a prefix of V. We now have two subcases:

(1) If A' equals Ω, then we necessarily get $A'' = A'$ and A cannot recognize V.

(2) If A' is some simple matching, then let u_x be the occurrence in N' of the variable x tested by A'. Since (N', ρ', σ', A') is a final state, we have $\sigma'(x) = \Omega$, that is, $U/u_x = \Omega$. If A recognize V, then A'' cannot be not a simple matching and the one-step transition starting from (N', τ', ϕ', A') must examine $\phi'(x)$. Thus, the subterm N''/u has a root constructor and so does $V/u \succeq N''_\Omega/u$.

□

Note that the above proof also applies to the case of tree-like automata.

I can now formulate the main result of this section and present two efficient, correct and complete algorithms for compiling pattern matching.

PROPOSITION 5.1. *Consider the two following compilation algorithms: use the $\mathcal{C}$ or $\mathcal{D}$ scheme following directions of appropriate matrices whenever an elementary test is emitted. When a matrix with no directions is discovered, then fail.*

Given a clause set E, both algorithms produce an automaton that implements correctly the matching by E, if and only if the matching predicate defined by the patterns of E is sequential in the Kahn-Plotkin sense.

PROOF. The direct sense of the proof comes from the fact that discovering a directionless matrix during compilation implies the existence of a value U, such that the matching predicate defined by E is not sequential at U (corollary 5.1). Conversely, automata of both kinds define sequential predicates (lemma 5.4). Therefore, in case the compilation algorithm succeeds, the matching predicate defined by E is sequential, by correctness property 3.1 or 4.1. □

6. CONCLUSION

In this paper, I have fully described two pattern matching compilation schemes. Moreover, I showed that both schemes can serve as a basis for a realistic and correct compilation of lazy pattern matching, this result is entirely new for the second scheme. The main advantage of lazy pattern matching is maximization of the termination properties of programs. More precisely, a standard supercombinator-based compiler extended with the correct compilation of lazy pattern matching described here implements a correct reduction strategy for ML considered as a rewriting system [Maranget 1992b]. The implemented strategy is correct in the sense that it terminates and produces a result whenever any other strategy does.

When designing a semantics or at least a precise definition for lazy functional languages such as LML or Haskell, the denotation of pattern matching expressions should, in my view, be based upon a matching predicate very similar to the Laville's definition (definition 3). This definition of matching is appealing because of its simplicity and generality. The compilers described in this paper are correct relatively to this definition of matching, but they are not complete, since they only accept sequential pattern matching predicates. First, observe that the sequential functions are widely accepted as the class of functions that can be computed naturally on a sequential computer. On this topic, see, for instance, Berry and Lévy [1979] and Plotkin [1977]. In other words, it would be problematic to design a pattern matching compiler that would be both correct and complete relatively to Laville's semantics, unless the target automata of this new hypothetical compiler run on a parallel computer, real or simulated. Thus, the compilation algorithm 5.1 goes as far as can be expected of any compiler that produces code for ordinary, sequential computers. Also observe that the compilers presented here detect all non-sequential matchings and can thus react appropriately when given such a matching as input. An appropriate reaction is usually to issue a warning message before producing a non-correct automaton.

Published language definitions [Johnsson 1984; Hudak, P. and Peyton Jones, S. L. and Wadler, P. 1992] present a different semantics for pattern matching: they select one particular sequential semantics (i.e., left-to-right). This approach lacks the generality and elegance of Laville's solution. First, there is no natural correspondence between this left-to-right matching order and terminating rewriting strategies, as there is one in the case of the lazy matching order. I already gave an intuition for this argument in the introduction: the left-to-right semantics for the function `F` does not gives a non-Ω result whenever possible. Therefore, a left-to-right semantics can be seen as incorrect, at least, it is not as expressive as possible.

Moreover, in my opinion, including a particular compilation technique for pattern matching in the language definition is not useful to the programmers and may even prove undesirable. Programmers do not need such a precision and relying on non-obvious features does not usually lead to a very clear programming style. In a slightly different context, observe that the definition of the C language does not specify the evaluation order of the arguments to primitives such as "+".

Furthermore, it is important not to overspecify the language definition. Doing so constraints the compiler designers too much and prevents them from introducing some optimizations. As a matter of fact, the correct compilation of the lazy semantics for pattern matching has some practical benefits. Let us consider a slight generalization of the $\mathcal{C}$-based compilation algorithm as introduced under the name of "adaptive" compilation by Sekar, Ramesh, and Ramakrishnan [1992]: when a pattern matrix (P) with directions is discovered, then select one direction, otherwise select column index 1. These authors showed that such a compilation technique produces automata that are always smaller than the automata compiled using the standard left-to-right technique. Although their measure of automata size is not strictly correlated to output code size, experimental comparisons between adaptive and left-to-right versions of the $\mathcal{C}$ and $\mathcal{D}$ compilation algorithm show that the adaptive versions of compilation schemes usually produce shorter code.

Sekar, Ramesh, and Ramakrishnan [1992] also showed the semantic correctness of their compiler: it produces automata that detect matching values whenever left-to-right automata do. My view of correctness is different: I only require that the sequential pattern matchings are compiled correctly, so that I leave room for more aggressive optimizations. Consider the following "ultra-adaptive" algorithm: when a pattern matrix (P) with directions is discovered, then select one direction, otherwise select some column index following heuristics aimed at reducing automata size. A first and efficient heuristic, useful both in $\mathcal{C}$ and $\mathcal{D}$-based ultra-adaptive algorithms, consist in selecting a *partial direction*, a partial direction being some column index that is a direction towards a maximal number of consecutive rows in (P), starting from the top (see the end of section 5.1). If there are more than one partial direction, they are further discriminated by using techniques that differ according to the basic compilation scheme used. In the $\mathcal{C}$-based algorithm, column indices are further selected by following some of the heuristics introduced in the case 3 at the end of section 5.1. A simple and reasonably satisfactory choice is here to first minimize the number of variables in the selected columns (in an attempt to avoid row duplication) and then, if necessary, to minimize the number of different root constructors (in an attempt to locally minimize automaton breadth). Similar heuristics (and their limitations) are described by Baudinet and MacQueen [1986] and by Sekar, Ramesh, and Ramakrishnan [1992]. In the $\mathcal{D}$-based algorithm, the number of transitions between variable and non-variable patterns is minimized (in an attempt to locally minimize control transfer possibilities). The same heuristics apply to select one direction, when matrix (P) admits several directions.

The choice of ultra-adaptive pattern matching has some additional benefits in term of running time performance. When the compiled pattern matching defines a sequential matching predicate, correct tree-like automata perform exactly the tests needed to identify a matching value (otherwise, they would not be correct [Puel and Suárez 1993]), these tests being performed only once. In this sense, correct tree-like automata have an optimal run-time behavior. As to correct automata

with failures, although they also perform exactly the tests needed to identify a matching value, these tests can be performed several times. Nevertheless, correct automata with failures still do not perform useless tests, so that they tend to run faster than incorrect automata. Furthermore, some of the heuristics also tend to lower running time. This is the case of the partial directions in both schemes (some of the matching values are found by performing only the elementary tests that are indeed necessary) and of minimizing control transfer possibilities in the $\mathcal{D}$ scheme (tests are less likely to be performed several times). However, it is fair to say that this kind of optimization lies at the fine tuning level, since performing tests accounts for a small part of the running time of a typical ML program. As a consequence, the gain in code size obtained by using some $\mathcal{D}$-based algorithm outweighs the gain in run-time efficiency obtained by using the corresponding $\mathcal{C}$-based algorithm.

This conclusion is confirmed by the following experiments. The following table describes the size (in kilo-bytes and percentage) of the GAML compiler [Maranget 1991] targeted for the sparc architecture (GAML is written in GAML), when compiled using various pattern matching compilation techniques:

	trees ($\mathcal{C}$)	failures ($\mathcal{D}$)
left-to-right	1176 (110 %)	1080 (101 %)
lazy (ultra-adaptive)	1128 (105 %)	1072 (100 %)

From the experiment just described above, we get four versions of the GAML compiler, other experiments show that there is no significant difference in run-time performance between these compilers.

Thus, using the tree-based technique —as the SML/NJ and CAML V3.1 compilers do— means paying a significant penalty in code size over the failure-based technique, this penalty not being compensated by a significant gain in run-time efficiency. As a final conclusion, these figures show the practical benefits of using automata with failures instead of tree-like automata. These benefits can be enjoyed whether the lazy semantics for pattern matching is preferred or not.

REFERENCES

Appel, A. W. and MacQueen, D. B. 1991. Standard ML of New Jersey. In *International Symposium on Programming Language Implementation and Logic Programming*, pp. 1–13. Springer-Verlag. Lecture Notes in Computer Science 583.

Augustsson, L. 1985. Compiling pattern matching. In J.-P. Jouannaud (Ed.), *Functional Programming Languages and Computer Architecture*, pp. 368–381. Springer-Verlag. Lecture Notes in Computer Science 201.

Baudinet, M. and MacQueen, D. B. 1986. Tree pattern matching for ML. Available electonically at `ftp://ftp.princeton.edu/dist/ml/papers/85-note-baudinet.ps`.

Berry, G. and Lévy, J.-J. 1979. A survey of some syntactic results in the lambda-calculus. In *Mathematical Foundations of Computer Science*. Springer-Verlag. Lecture Notes in Computer Science 74.

Curien, P.-L. 1986. *categorical combinators, sequential algorithms, and functional programming*. Pitman, London, and John Wiley and Sons, New York.

Harper, R., Milner, R., and Tofte, M. 1991. *The definition of Standard ML*. The MIT Press.

Hudak, P. and Peyton Jones, S. L. and Wadler, P. 1992. Report on the programming language haskell. *SIGPLAN-notices 27*, 5 (March), section R.

Huet, G. and Lévy, J.-J. 1991. Call by need computations in non-ambiguous linear term rewriting systems. In J.-L. Lassez and G. D. Plotkin (Eds.), *Computational Logic, Essays in Honor of Alan Robinson*. The MIT Press.

Johnsson, T. 1984. Efficient compilation for lazy evaluation. *SIGPLAN Notices 19,* 6 (June), 58–69.

Kahn, G. and Plotkin, G. D. 1978. Domaines concrets. Technical Report 336, IRIA Laboria. (In French).

Kennaway, R. 1990. The specificity rule for lazy pattern matching in ambiguous term rewriting systems. In *European Symposium On Programming*. Springer-Verlag. Lecture Notes in Computer Science 432.

Laville, A. 1991. Comparison of priority rules in pattern matching and term rewriting. *Journal of Symbolic Computations 11,* 4 (April), 321–348.

Leroy, X. 1990. The ZINC experiment: an economical implementation of the ML languag e. Technical Report 117 (February), INRIA.

Leroy, X. *et al.*. 1995. The Caml Light system, version 0.7. Software and documentation available electronically at `http://pauillac.inria.fr/caml/distrib-caml-light.html`.

Maranget, L. 1991. GAML: A parallel implementation of lazy ML. In *Functional Programming and Computer Architecture*. Springer-Verlag. Lecture Notes in Computer Science 523.

Maranget, L. 1992a. Compiling lazy patern matching. In *Lisp and Functional Programming*. ACM Press.

Maranget, L. 1992b. La stratégie paresseuse. Ph. D. thesis, Université de Paris VII. (In French).

Plotkin, G. D. 1977. LCF considered as a programming language. *Theoretical Computer Science 5*, 223–255.

Puel, L. and Suárez, A. 1993. Compiling pattern matching by term decomposition. *Journal of Symbolic Computation 15,* 1, 1–26.

Sekar, R. C., Ramesh, R., and Ramakrishnan, I. V. 1992. Adaptive patter matching. In *International Colloquium on Automata Languages and Programming*. Springer-Verlag. Lecture Notes in Computer Science 623.

Wadler, P. 1987. Efficient compilation of pattern matching. In *Peyton Jones S. L. , The Implementation of Functional Programming Languages*, Chapter 6. Prentice-Hall.

Weis, P. 1990. The CAML reference manual, version 2.6.1. Technical Report 121, INRIA.