

Extending COQ with Imperative Features and Its Application to SAT Verification^{*}

Michaël Armand¹, Benjamin Grégoire¹, Arnaud Spiwack², and Laurent Théry¹

¹ INRIA Sophia Antipolis - Méditerranée, France
{Michael.Armand,Benjamin.Gregoire,Laurent.Thery}@inria.fr
² LIX, École Polytechnique, France
Arnaud.Spiwack@lix.polytechnique.fr

Abstract. COQ has within its logic a programming language that can be used to replace many deduction steps into a single computation, this is the so-called reflection. In this paper, we present two extensions of the evaluation mechanism that preserve its correctness and make it possible to deal with cpu-intensive tasks such as proof checking of SAT traces.

1 Introduction

In the COQ proof assistant [3], functions are active objects. For example, let us consider the sum of two natural numbers. The sum function is defined recursively on its arguments using, say, Peano's definition. Then $1+2$ is an expression that can be *computed* to its expected value. In particular to prove $1+2=3$, we simply need to know that equality is reflexive, and the system takes care of checking that $1+2$ and 3 compute to the same value. Note that this computation (also called normalisation of λ -calculus) is not restricted to ground terms, like in our example: it can act as a symbolic evaluation on any term. Furthermore, COQ being based on the Curry-Howard isomorphism, writing a proof or a program is essentially the same. These remarks are the bases of proofs by reflection, which consist in replacing many deduction steps by a single computation. This technique has become popular in the COQ community since a few years. One of its most impressive application is the formal proof of the four-colour theorem [8].

Using reflection can greatly improve the checking time of proofs. However, as one pushes the limits of it, efficiency can become a concern. In that respect, a major improvement has already been achieved through the introduction of a dedicated virtual machine [9] allowing COQ programs to compare with (bytecode compiled) OCAML [12] ones. Still, there are strong restrictions left. First, there are no primitive data-structures. Every type is encoded using the constructs allowed by the system (primarily, inductive definitions). Also, there is no possibility to use destructive data-structures, which can be much more efficient than purely functional ones in some circumstances. To be able to go further on what can be efficiently programmed in COQ, we will add new data-structures such as

^{*} This work was supported in part by the french ANR DECERT initiative.

native integers and destructive arrays. The challenge is to achieve this, changing as little as possible, in order to preserve trust in the correctness of evaluation in COQ, and nevertheless to get an effective speed-up thank to the new features.

The paper is organised in the following way. In Section 2, we describe how it is possible in COQ to benefit from the arithmetic capabilities of the microprocessor. Section 3 is dedicated to arrays. We then propose two examples that illustrate the benefit of our extension. In Section 4, we present the toy example of the Mini-Rubik for which we use computation to prove that it is always solvable in less than 11 moves. Section 5 is dedicated to a more challenging example. In order to prove the unsatisfiability of large boolean propositions, we replay in a reflexive way proof traces generated by SAT solvers.

2 Extending Coq with Machine Integers

Arithmetic is currently defined in COQ as a standard inductive type. Thus, computations with numbers do not differ from other data-structures: it is a plain symbolic evaluation. What we aim at, here, is to rely on the arithmetic of the processor to speed-up computations within COQ. In order to add machine integers, a first possibility is to extend the theory underlying the COQ logic with:

- one primitive type `int`;
- the constructors $0, 1, 2, \dots, 2^n - 1$ of type `int`;
- the basic primitive functions over the type `int` such as $+, *, \dots$;
- the corresponding reduction rules for each primitive function.

It is also necessary to give it an equational theory, for instance, Peano theory together with a lemma stating that $(2^n - 1) + 1 = 0$. However, this approach has some drawbacks:

- It adds a large amount of new constructions to the theory. This goes against de Bruijn’s principle which states that keeping the theory and its implementation as small as possible highly contributes to the trust one has in a system. Furthermore, on a more practical side, it will have a deep impact in the implementation, since the terms will have to be extended with new syntactic categories (primitive types and primitive functions).
- It adds a lot of new reductions, not only for ground terms but also for theorems. For example, if we consider the theorem `n_plus_zero` that states that $\forall n: \text{int}. n + 0 = n$, it could be convenient to have `(n_plus_zero 7)` reduces to `(refl int 7)` where `refl` represents the reflexivity of equality. It is not clear that way that one captures all the necessary reductions.

For these reasons, we have taken an alternative approach. Efficient evaluation in COQ, as provided by the virtual machine, uses a compilation step. Before evaluating a term, it transforms it into another representation that is more suitable for performing reduction. The idea is to introduce the native machine integers not as part of the theory of COQ but only in this compilation phase. So, the type `int` is defined using the standard commands as a type with a single constructor that contains n digits:

Definition `digit := bool`.

Inductive `int : Type := ln(dn-1 ... d1 d0 : digit) : int`.

Note that, in the current development, the type `int` is not parametric in n . We use a specific n to get a direct mapping to machine words and their operations. Still, the integration is done in a generic way so integers for a different n could be derived easily.

The primitive functions are not defined directly. We relate the machine numbers `int` with the relative numbers \mathbb{Z} (the COQ representation of \mathbb{Z}) with the two functions $\overline{\bullet} : \text{int} \rightarrow \mathbb{Z}$ and its inverse $[\bullet] : \mathbb{Z} \rightarrow \text{int}$ and we prove that they satisfy the following two properties:

$$\begin{aligned} \forall i : \text{int}. \overline{[i]} &= i \\ \forall z : \mathbb{Z}. [z] &= z \pmod{2^n} \end{aligned}$$

Now, it is straightforward to define the primitive functions of `int` as the image of the corresponding function of \mathbb{Z} . For example, addition for `int` is defined as follows:

Definition `i1 +int i2 := [i1 +Z i2]`

It is also straightforward to derive the basic properties of these functions from the properties of the corresponding functions on \mathbb{Z} . This set of definitions and properties will let the user manipulate the type `int` in COQ as any other type. So we preserve the property that everything is defined from base principle.

Now, the trick is to modify the compiler in such a way that it treats the type `int` as real machine integers. The main difficulty is that COQ requires strong reduction. This is not the case of traditional functional languages where only weak reduction is needed (no reduction under binders). Before explaining our modification to the compiler, we first give an overview of what symbolic weak and strong reductions are and then explain how the compiler works.

2.1 Strong Reduction by Symbolic Weak Reduction

In order to compute the strong normal form of a term t or to test the convertibility between two terms t_1 and t_2 , the COQ system uses a compiled implementation of the symbolic calculus [9,2]. We briefly recall what symbolic computation is starting from the pure λ -calculus extended with inductive types.

Each inductive type is defined by a name I and a fixed number of constructors $|I|$. In this context the constructor `ln` is represented by $C_{\text{int},1}$. The syntax of the λ -calculus is extended with constructors and case analysis:

$$a ::= x \mid \lambda x.a \mid a_1 a_2 \mid C_{I,i}(\mathbf{a}) \mid \text{case } a \text{ of } (\mathbf{x}_i \Rightarrow a_i)_{1 \leq i \leq |I|}$$

the reduction rules are

$$\begin{aligned} (\lambda x.a_1)a_2 &\Rightarrow a_1\{x \leftarrow a_2\} && (\beta) \\ \text{case } C_{I,j}(\mathbf{a}) \text{ of } (\mathbf{x}_i \Rightarrow a_i)_{1 \leq i \leq |I|} &\Rightarrow a_j\{\mathbf{x}_j \leftarrow \mathbf{a}\} && (\iota) \\ \Gamma(\mathbf{a}) &\Rightarrow \Gamma(\mathbf{a}') \quad \text{if } a \Rightarrow a' && (\text{context}) \end{aligned}$$

where there is no restriction on the context Γ . Reduction can happen anywhere, in particular under binders or inside a branch of a case. We are interested in

computing the strong $\beta\iota$ -normal form of the λ -term a . This is done by iterating weak symbolic reduction and *readback*.

We first introduce the symbolic calculus:

Extended terms	$b ::= x \mid \lambda x.b \mid b_1 b_2 \mid C_{I,i}(\mathbf{b}) \mid \text{case } b \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq I } \mid [k]$
Accumulators	$k ::= h \mid k v$
Atoms	$h ::= \tilde{x} \mid \text{case } k \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq I }$
Values	$v ::= \lambda x.b \mid C_{I,i}(\mathbf{v}) \mid [k]$

The value $[h v_1 \dots v_n]$ is called an accumulator. It represents h applied to arguments $v_1 \dots v_n$. The atom \tilde{x} is a symbolic variable. It represents the free variable x . Finally, $\text{case } k \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|}$ represents a suspended cases which cannot reduce since its argument does not reduce to a constructor.

The rules for weak reduction are defined as follows:

$$\begin{array}{ll}
 (\lambda x.b) v \rightarrow b\{x \leftarrow v\} & (\beta_v) \\
 [k] v \rightarrow [k v] & (\beta_s) \\
 \text{case } C_{I,j}(\mathbf{v}) \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|} \rightarrow b_j\{\mathbf{x}_j \leftarrow \mathbf{v}\} & (\iota_v) \\
 \text{case } [k] \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|} \rightarrow [\text{case } k \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|}] & (\iota_s) \\
 \Gamma_v(b) \rightarrow \Gamma_v(b') \text{ if } b \rightarrow b' & (\text{context}_v)
 \end{array}$$

where $\Gamma_v(\bullet) ::= \bullet v \mid b \bullet \mid C_{I,i}(\mathbf{b} \bullet \mathbf{v}) \mid \text{case } \bullet \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|}$

The rules (β_v) and (ι_v) are the standard rules for call-by-value function application and case reduction. The rule (β_s) (“symbolic” β -reduction) handles the case where the function part of an application is not a function: a free variable $[\tilde{x}]$ or an application of a free variable $[\tilde{x} v_1 \dots v_n]$ or a suspended case. In that case, the accumulator simply absorbs its argument. The rule (ι_s) explains what to do when an accumulator is argument of a case: we simply *remember* that the case construct cannot reduce by producing a new accumulator. The rule (context_v) enforces weak reduction (no reduction under binder) and a right to left evaluation order (the argument being evaluated before the functional part)¹.

In order to compute the normal form of a λ -term a , we first inject a into the symbolic calculus. This is done by replacing each free variable x of a by its corresponding symbolic value $[\tilde{x}]$. We obtain a closed symbolic term: the variable \tilde{x} is symbolic and not subject to substitution. In order to compute the normal form of a closed symbolic term b , we first compute its symbolic head normal form $\mathcal{V}(b)$ (also called value); then we *read back* the resulting value:

$$\begin{array}{ll}
 \mathcal{N}(b) & = \mathcal{R}(\mathcal{V}(b)) & (1) \\
 \mathcal{R}(\lambda x.b) & = \lambda y.\mathcal{N}((\lambda x.b) [\tilde{y}]) \quad y \text{ fresh} & (2) \\
 \mathcal{R}(C_{I,i}(\mathbf{v})) & = C_{I,i}(\mathcal{R}(\mathbf{v})) & (3) \\
 \mathcal{R}([k]) & = \mathcal{R}'(k) & (4) \\
 \mathcal{R}'(k v) & = \mathcal{R}'(k) \mathcal{R}(v) & (5) \\
 \mathcal{R}'(\tilde{x}) & = x & (6) \\
 \mathcal{R}'(\text{case } k \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|}) & = \text{case } \mathcal{R}'(k) \text{ of } (\mathbf{x}_i \Rightarrow \mathcal{N}(b C_i([\tilde{y}_i])))_{1 \leq i \leq |I|} & (7) \\
 & \text{where } b = \lambda x.\text{case } x \text{ of } (\mathbf{x}_i \Rightarrow b_i)_{1 \leq i \leq |I|} \\
 & \text{and } \mathbf{y}_i \text{ are sequences of fresh variables with } |\mathbf{y}_i| = |\mathbf{x}_i|
 \end{array}$$

¹ The evaluation order is important when using a virtual machine like the ZAM [11] with n -ary applications to execute the symbolic calculus.

The readback function \mathcal{R} is defined recursively. It transforms a value v into a normalised source term. Reading back an atom \tilde{x} (equation 6) simply consists in extracting the variable x . Reading back an accumulator $k v$ (equations 5) consists in applying the readback of the functional part to the readback of the argument. The interesting case is for function $\lambda x.b$ (equation 2). It consists in applying the functional value to a value $[\tilde{y}]$ representing a fresh variable. Here, “fresh” means that y is not a free variable of b . Then, we compute the value of the application, which reduces in one step to $b\{x \leftarrow [\tilde{y}]\}$, and reads it back as a normalised term a . The normal form of $\lambda x.b$ is $\lambda y.a$, which is correct up to α -conversion. The same idea is used to obtain the normal form of the branches of a case.

In [9], the authors prove the following theorem in the case of the λ -calculus:

Theorem 1. *If a is a closed, strongly normalizing λ -term, then $\mathcal{N}(a)$ is defined and is the normal form of a .*

The normal form of a term can be obtained by recursively computing its symbolic weak normal form and reading back the resulting value. The efficiency of the process clearly depends on the efficiency of the weak evaluation.

2.2 Compiling the Symbolic Calculus

Weak symbolic reduction can be implemented using a compiler and an abstract machine. The abstract machine we present here is a simplified version of the ZAM [11]. We write \hat{v} the values manipulated by the abstract machine. They are pointers to heap allocated blocks $[T : \hat{v}_1, \dots, \hat{v}_n]$, where T is a tag, and the \hat{v}_i are values belonging to the block.

A machine state (e, c, s) has three components: an environment e that contains a sequence of machine values $\hat{v}_1, \dots, \hat{v}_n$ (it associates to the variable of de Bruijn index i the value \hat{v}_i); a code pointer c that represents the term being executed; a stack frame s that contains function arguments, intermediate results and return context $\langle c, e \rangle$. The semantics of the instruction set and the compilation rules are given Figure 1.

The compilation scheme $\llbracket b \rrbracket c$ takes a term b that has to be compiled and a code c that corresponds to the continuation of b . If b is normalising, the execution of $(e, \llbracket b \rrbracket c, s)$ leads to $(e, c, \hat{v} :: s)$ where \hat{v} is the machine representation of the value v of b where the free variables have been substituted by their values in e .

Evaluating the code corresponding to a function $\lambda x.b$ builds a closure $[T_\lambda : c, e]$ where c is the code pointer corresponding to b and e the current environment. For application, a return context is pushed on the stack, then the argument and the function are evaluated, and finally, the APPLY instruction starts the evaluation of the closure. For constructors, the instruction MAKEBLOCK(n, T) builds a block² $[T : \hat{v}_1 \dots \hat{v}_n]$ which is the machine representation of constructors. The compilation of a case starts by a PUSHRA, which saves the return context

² The compilation of a block erases the inductive name in the constructor. For the correctness, we refer to [9].

$$\begin{aligned}
\llbracket x \rrbracket c &= \text{ACCESS}(i); c \quad \text{where } i \text{ is the de Bruijn index of } x \\
\llbracket \lambda x. b \rrbracket c &= \text{CLOSURE}(\text{GRAB}; \llbracket b \rrbracket \text{RETURN}); c \\
\llbracket b_1 b_2 \rrbracket c &= \text{PUSHRA}(c); \llbracket b_2 \rrbracket \llbracket b_1 \rrbracket \text{APPLY} \\
\llbracket C_{I,i}(b_1, \dots, b_n) \rrbracket c &= \llbracket b_n \rrbracket \dots \llbracket b_1 \rrbracket \text{MAKEBLOCK}(n, i); c \\
\llbracket \text{case } b \text{ of } (x_i \Rightarrow b)_{1 \leq i \leq |I|} \rrbracket c &= \\
&\quad \text{PUSHRA}(c); \llbracket b \rrbracket \text{SWITCH}(\llbracket b_1 \rrbracket \text{RETURN}, \dots, \llbracket b_{|I|} \rrbracket \text{RETURN}) \\
\\
(e, \text{ACCESS}(i); c, s) &\rightsquigarrow (e, c, e[i] :: s) \\
(e, \text{CLOSURE}(c_f); c, s) &\rightsquigarrow (e, c, [T_\lambda : c_f, e] :: s) \\
(e, \text{GRAB}; c, \hat{v} :: s) &\rightsquigarrow (\hat{v} :: e, c, s) \\
(e_f, \text{RETURN}, \hat{v} :: \langle c, e \rangle :: s) &\rightsquigarrow (e, c, \hat{v} :: s) \\
(e, \text{APPLY}, [T_\lambda : c_f, e] :: s) &\rightsquigarrow (e_f, c_f, s) \\
(e, \text{PUSHRA}(c_1); c_2, s) &\rightsquigarrow (e, c_2, \langle c_1, e \rangle :: s) \\
(e, \text{MAKEBLOCK}(n, T); c, \hat{v}_1 :: \dots :: \hat{v}_n :: s) &\rightsquigarrow (e, c, [T : \hat{v}_1 :: \dots :: \hat{v}_n] :: s) \\
(e, \text{SWITCH}(c_1, \dots, c_m), [T : \hat{v}_1 :: \dots :: \hat{v}_n] :: s) &\rightsquigarrow (\hat{v}_n :: \dots :: \hat{v}_1 :: e, c_T, s) \\
(e, \text{SWITCH}(c_1, \dots, c_m), [0 : \text{ACCU}, \hat{k}] :: s) &\rightsquigarrow (e, \text{RETURN}, \hat{v} :: s) \\
\text{where } \hat{v} = [0 : \text{ACCU}, [1 : \hat{k}, [T_\lambda : \text{GRAB}; \text{SWITCH}(c_1, \dots, c_m)]]] & \\
(e_f, \text{ACCU}, \hat{v}_a :: \langle c, e \rangle :: s) &\rightsquigarrow (e, c, [0 : \text{ACCU}, e_f :: \hat{v}_a] :: s)
\end{aligned}$$

Fig. 1. Compilation rules and semantics of the virtual machine

(used at the end of branches), then the argument is evaluated and the SWITCH instruction jumps to the corresponding branche.

What happens for the symbolic calculus? When an APPLY instruction is executed, the top stack value is not necessarily a closure, it can be the machine representation of an accumulator. An accumulator $[k]$ is represented like a closure: $[0 : \text{ACCU}, \hat{k}]$. Furthermore, k being of the form $h v_1 \dots v_n$, \hat{k} is represented as an environment: the sequence $\hat{h}, \hat{v}_1, \dots, \hat{v}_n$. The ACCU instruction takes the top value of the stack, pushes it at the end of the environment, rebuilds an accumulate block and returns. In that way, the APPLY instruction does not need to perform an extra test.

For the same reason, when a SWITCH instruction is executed, the top value is not necessarily a constructor, it can be an accumulator. If the tag is 0, the matched value is an accumulator, the SWITCH instruction builds an accumulate block representing the suspended case. In practice, 0 branches are automatically added to cases by the compiler, thus the SWITCH instruction of the ZAM can be used without extra test. Note that for atoms, \hat{x} is represented by the block $[0 : x]$ and case k of $(x_i \Rightarrow b)_{1 \leq i \leq |I|}$ is represented by $[1 : \hat{k}, [T_\lambda : c, e]]$ where \hat{k} is the machine representation of k , c and e are the code and the environment for the function $\lambda x. \text{case } x \text{ of } (x_i \Rightarrow b)_{1 \leq i \leq |I|}$.

In order to normalise a λ -term a using the virtual machine, we first compute $c = \llbracket a \rrbracket$ and start the evaluation with the abstract machine in the state (e, c, \emptyset) , where e is an environment associating to each free variable x of a its value $[\hat{x}]$ encoded by the heap block $[0 : \text{ACCU}, [0 : \hat{x}]]$. When the machine stops, we obtain a value v on the top of the stack. The readback function analyses which kind of value it is. It can either be a closure, a constructor, or an accumulator.

This can be done by a simple inspection of the tag. If the tag is T_λ , we have to normalise v $[\tilde{y}]$. This is done simply by restarting the machine in the state $(e, \text{APPLY}, v :: [\tilde{y}] :: \emptyset)$. The same technique is used to normalise the branches of a suspended case.

2.3 Adding Machine Integers

We are now ready to explain how we can take advantage of the compilation mechanism to boost the evaluation of a λ -calculus extended with inductive types using the machine-integer operations. Of course, the gain will only be effective for programs using the previously defined inductive type `int`.

We extend the λ -calculus with a global environment Δ associating global variables g to their definition (λ -term):

$$\begin{aligned} a &::= x \mid \lambda x. a \mid a_1 \ a_2 \mid C_{I,i}(\mathbf{a}) \mid \text{case } a \text{ of } (x_i \Rightarrow a_i)_{1 \leq i \leq |I|} \mid g \\ \Delta &::= \emptyset \mid \Delta :: (g, a). \end{aligned}$$

The reduction rules now depend on the global environment Δ and are extended with one rule for the reduction of global definitions: $g \rightarrow \Delta(g)$.

We assume that n has been chosen in such a way that the term $\text{ln}(d_{n-1}, \dots, d_0)$ is isomorphic to the machine word $d_{n-1} \dots d_0$ if the d_j are all constructors (true stands for the machine digit 1 and false for 0). In the following, the term $\text{ln}(d_{n-1}, \dots, d_0)$ is written p if all the d_j are constructors. We write \dot{p} for the machine representation of p . If m is a machine integer, we write $|m|$ its representation as a term of type `int`; we have $p = |\dot{p}|$ and $m = |\dot{m}|$.

In the following, we assume that we have a global definition $+$ performing the addition of two `int`. We denote by $+_a$ its associated definition and we write $+_M$ the processor addition. We assume that $+$ does what it is supposed to do, i.e.:

$$p_1 + p_2 \Rightarrow^* |\dot{p}_1 +_M \dot{p}_2|$$

This gives us a first way to boost the reduction of $+$ when the two arguments are of the form p_1 and p_2 ; instead of accessing to the global definition $+_a$ of $+$ and then reducing the application $+_a \ p_1 \ p_2$, one can directly compute $|\dot{p}_1 +_M \dot{p}_2|$. This solution does not work so well when additions are nested. For example, during the reduction of $(p_1 + p_2) + p_3$, the machine word $\dot{p}_1 +_M \dot{p}_2$ would be injected into its constructor representation at the end of the evaluation of $p_1 + p_2$ and then immediately re-injected into a machine word to perform the second addition. We have chosen a different solution that overcomes this problem.

We extend the symbolic calculus with machine integers and their primitive operations. The idea is to try to maintain as long as possible the terms of type `int` in their machine representation. The syntax of the new symbolic calculus is extended with machine integers:

$$\begin{aligned} b &::= x \mid \lambda x. b \mid b_1 \ b_2 \mid C_{I,i}(\mathbf{b}) \mid \text{case } b \text{ of } (x_i \Rightarrow b_i)_{1 \leq i \leq |I|} \mid [k] \mid m \\ v &::= \lambda x. b \mid [k] \mid C_{I,i}(v_1, \dots, v_n) \mid m \\ \Delta_b &::= \emptyset \mid \Delta_b :: (g, v) \end{aligned}$$

In the definition of values, we exclude the case $\text{ln}(v_1, \dots, v_n)$ where the v_i are all true or false. New reduction rules are added to the calculus. First, a special case is added for the constructor ln :

$$p \rightarrow \dot{p}$$

In other words, when a constructor of type int can be represented by a machine word, its value is the machine representation. Second, for each global definition representing a primitive operation over int , some special rules are added. Let us consider addition, we add two rules:

$$\begin{aligned} m_1 + m_2 &\rightarrow m_1 +_M m_2 \\ v_1 + v_2 &\rightarrow \Delta_b(+)\ v_1\ v_2 \end{aligned}$$

The first rule applies when the two arguments of $+$ are in machine representation. The result is given by the machine addition. The second one applies when one of the two arguments is not in machine representation. It can either be an accumulator or an ln constructor with one of its arguments being an accumulator. In that case, the usual rule for global variable is used: the variable is replaced by its associating value in Δ_b .

Pattern matching on terms of type int has also to be taken care of. The matched value can be a machine word whereas a constructor value or an accumulator is expected. For this reason, we add the rule:

$$\text{case } m \text{ of } (x_i \Rightarrow b)_{1 \leq i \leq |I|} \rightarrow \text{case } |m| \text{ of } (x_i \Rightarrow b)_{1 \leq i \leq |I|}$$

Finally the readback function only needs to be extended so to get rid of machine integers: $\mathcal{R}(m) = |m|$.

Theorem 1 can be extended to this new symbolic calculus:

Theorem 2. *If for all global definitions g with a special shortcut we have $\Delta_b(g)\ \mathbf{m} \rightarrow^* g_M\ \mathbf{m}$. For all closed term a , well typed and strongly normalising, then $\mathcal{N}(a)$ is defined and is then normal form of a .*

What remains to be modified is the virtual machine and the compilation scheme. Previously, the values of the virtual machine were only pointers to heap-allocated blocks. The values are now extended with machine integers. Two instructions TOINT and OFINT are added to the virtual machine. Their semantics is given by:

$$(e, \text{OFINT}; c, d_0 :: \dots :: d_{n-1} :: s) \rightsquigarrow (e, c, m :: s)\ m = d_0 \dots d_{n-1} \quad (1)$$

$$(e, \text{OFINT}; c, v_0 :: \dots :: v_{n-1} :: s) \rightsquigarrow (e, c, v :: s)\ \text{otherwise} \quad (2)$$

$$\text{where } v = [1 : v_0, \dots, v_{n-1}]$$

$$(e, \text{TOINT}; c, m :: s) \rightsquigarrow (e, c, v :: s) \quad (3)$$

$$\text{where } v = [1 : d_0, \dots, d_{n-1}]$$

$$(e, \text{TOINT}; c, v :: s) \rightsquigarrow (e, c, v :: s)\ \text{otherwise} \quad (4)$$

We also add one instruction for each primitive operations. For example, the instruction ADD corresponds to the addition:

$$(e, \text{ADD}; c, m_1 :: m_2 :: s) \rightsquigarrow (e, c, m_1 +_M m_2 :: s) \quad (5)$$

$$(e, \text{ADD}; c, v_1 :: v_2 :: s) \rightsquigarrow (e_+, c_+, v_1 :: v_2 :: \langle c, e \rangle :: s)\ \text{otherwise} \quad (6)$$

$$\text{where } \Delta_b(+)= [T_\lambda : c_+, e_+]$$

Finally, we modify the compiler with special cases for the compilation of the `In` constructor, for the primitive operations and for the pattern matching over elements of type `int`:

$$\begin{aligned}
 \llbracket \text{In}(b_0, \dots, b_{n-1}) \rrbracket c &= \llbracket b_{n-1} \rrbracket \dots \llbracket b_0 \rrbracket \text{OFINT}; c \\
 \llbracket b_1 + b_2 \rrbracket c &= \llbracket b_2 \rrbracket \llbracket b_1 \rrbracket \text{ADD}; c \\
 \llbracket \text{case } b \text{ of } (x_i \Rightarrow b)_{1 \leq i \leq |I|} \rrbracket c &= \\
 \text{PUSHRA}(c); \llbracket b \rrbracket \text{TOINT}; \text{SWITCH}(\llbracket b_1 \rrbracket \text{RETURN}, \dots, \llbracket b_n \rrbracket \text{RETURN}) \\
 \text{if type of } b = \text{int}
 \end{aligned}$$

The compilation of the `In` constructor generates an `OFINT` as last instruction and not a `MAKEBLOCK` as for the other constructors. The `OFINT` instruction checks if the first n arguments on the stack correspond to machine representation of digits (i.e. a block representing the constructors `true` or `false`). If all the arguments are constructors, the instruction builds the corresponding machine word (this corresponds to the reduction rule $p \rightarrow \dot{p}$ at the symbolic level). If one of the argument is not a constructor, the instruction is equivalent to `MAKEBLOCK`($n, 1$).

The compilation of a `+` first evaluates its arguments, then the `ADD` checks if they are machine words. If it is the case, the instruction simply performs the addition. If not, the instruction gets the value $[T_\lambda : c_+, e_+]$ of the `+`, inserts a return context $\langle c, e \rangle$ and performs an `APPLY`.

The compilation of a pattern matching on an object of type `int` is also modified. A `TOINT` instruction is inserted just before the `SWITCH`. If the top value of the stack is a machine word, the `TOINT` instruction replaces it by its corresponding block representation. If not, the `TOINT` instruction does nothing. The semantics of the `SWITCH` does not need to be modified.

The readback function should be able to analyse the value it gets. Before adding machine integer, this was done by matching the tag of the heap block. Remember that a heap block is a pointer, i.e. a machine integer. The problem is how to differentiate between pointers and integers. Note that a similar problem occurs for the implementation of the new instructions, which have to test if some values are blocks or integers. Fortunately, there is an easy solution. Since the implementation of the COQ virtual machine is based on the one of OCAML. The OCAML garbage collector makes the difference between a machine word representing a pointer and a machine word representing an integer using the following convention: a pointer is a machine word with least significant bit set to 0, an integer is a machine word with the bit set to 1. So, an integer p is encoded by the machine word $2p + 1$. This is why OCAML, and now COQ, have integers of only 31 bits on a 32-bit architecture.

2.4 Primitive Functions

As adding a new primitive function requires some expertise, we have developed a reasonable library of primitive functions for the `int` type. It contains the usual functions (addition, multiplication, square root, comparison, logical functions, shifts) but also some iterators. Functions like

```

Definition foldi (A:Type) (F:A->A) (a:A) (n_s n_e:int) :=
  if n_s <= n_e then
    (fix aux (i:int) (ai:A) {
      if i = n_s then F i ai else aux (i-1) (F i ai)
    }) n_e a
  else a.

```

that computes $F\ n_s\ (F\ n_{s+1}\ (\dots\ (F\ n_e\ a)\ \dots))$ cannot be defined on top of our library. Because of the definition of `int`, this is not structurally recursive so COQ cannot establish that it always terminates. So, we add them as primitive functions.

3 Extending Coq with Persistent Arrays

Arrays are among the most important data-structures. Unfortunately, logics like the one of COQ are stateless. So, it is impossible to deal directly with destructive arrays as the ones we find in mainstream programming languages. The work-around is usually to use some flavour of purely functional arrays [13]. This works pretty well when arrays are rather small. For larger ones, not having an $O(1)$ access to elements of the array quickly becomes unmanageable.

In order to introduce destructive arrays, one way to go is to add states to the logic. Monads [18] are the standard way to do this. Unfortunately, monads are quite difficult to manage in a prover without developing some infrastructure (see [5,15] for example). An alternative approach is to develop some kind of program analysis that is capable of discovering (automatically or semi-automatically) that it can safely use destructive arrays instead of functional ones (see [14] for example). If one wants this technique to be applicable to a large set of programs, such analysis is usually rather complex. Here, we are going to follow a third approach and use destructive arrays but with a functional interface. These arrays are called persistent arrays [1]. For COQ, persistent arrays are implemented in a very naive way. An array is simply composed of the list of elements of the array and a default value.

```

Inductive array (A : Type) : Type := mkArray(elems : list A)(default : A).

```

As there is no exception in COQ, the default value is mainly used as a return value when accessing outside the range of the array. Instead of a default value, all functions manipulating arrays could have been parametrised by a proof that the access is valid. This last solution has two drawbacks. First, a proof has to be provided each time a function is used. Second in a call by value strategy, it adds extra costs since the proof has to be reduced before the actual function is evaluated. With a default value, the two basic operations on arrays are defined in a straightforward manner:

```

Definition get (A : Type)(t : array A)(n : int) := get_elem (default t) (elems t) n.
Definition set (A : Type)(t : array A)(n : int)(a : A) :=
  mkArray (upd_elem (elems t) n a) (default t).

```

where `(get_elem d l n)` returns the $n + 1$ -th element of the list l if n is less than the size of the list, d otherwise; and `(upd_elem l n a)` returns l where the

$n + 1$ -th element has been replaced by a if n is less than the size of the list, l otherwise. Both definitions are very inefficient. The access is linear in the number of elements and the update is also linear and furthermore reallocates a large part of the list.

Now, the virtual machine is going to conform to this functional behaviour but using destructive arrays. The idea is quite simple. Among all the versions of the array that may co-exist during execution, the last one (the newest one) is privileged and is represented by a destructive array. Look-ups and updates applied to this last version are then very efficient. Older versions of the array are not destructive arrays but point to the last version through a list of indirections. These indirections explain which modifications have to be applied in order to retrieve the values of the old array from the last version. Look-ups of old versions are possible (this is a requirement of the functional interface) but rather slow (linear to the number of updates). Updates just add a level on indirection. For the implementation, we have directly adapted the OCAML code proposed by J.C. Filliâtre in his paper [6]. Persistent arrays are defined as follows:

```
type 'a parray_kind = Array of 'a array | Updated of int * 'a * 'a parray
and 'a parray = ('a parray_kind) ref
```

A persistent array is a reference on a `parray_kind` which is either a destructive array (`Array`) or an indirection `Updated(i,v,t)` indicating that the persistent array is t except that at position i the value is v . The OCAML implementation does not contain explicitly the default value. It is stored in the last position of the array.

For the `get` function, we look directly in the array or follow the indirections:

```
let rec get p n =
  match !p with
  | Array t ->
    let l = Array.length t in
    if 0 <= n && n < l then Array.get t n else Array.get t (l-1)
  | Updated (k,e,p) -> if n = k then e else get p n
```

Note that in a path to the destructive array, there could be several occurrences of n (this location could have been updated several times) but we stop at the first one. For the `set` function, an indirection is added at the right position:

```
let set p n e =
  let kind = !p in
  match kind with
  | Array t ->
    if 0 <= n && n < Array.length t - 1 then
      let res = ref kind in
      p := Updated (n, Array.get t n, res); Array.set t n e; res
    else p
  | Updated _ ->
    if 0 <= n && n < length p then ref (Updated(n, e, p)) else p
```

Note that if we are updating outside the array, everything is left unchanged. Two more functions complement the library of arrays:

Definition `copy` $(A : \text{Type})(t : \text{array } A) := t$.

Definition `reroot` $(A : \text{Type})(t : \text{array } A) := t$

If the functional behaviour is the identity for both, they implement two distinct operations. With the first one, we get a physical copy of the array (a new independent destructive array is allocated). With the second one, we get an array with fast access (it is a destructive array). This is done without copying by recursively reverting all the indirections that lead to the destructive array in the array that is passed as argument. A very nice application of this `reroot` operation can be found in [6].

The extension of the virtual machine of COQ follows the same methodology than for machine integers. Two translation functions are used to transform a COQ representation of array into its virtual machine representation and conversely. The only detail we had to take care of is that OCAML arrays are limited in size. If the size of the COQ array is greater than the maximum OCAML size then the virtual machine switches to the inefficient COQ representation. For the compilation, array primitives like `get` are compiled in a slightly different way. This is due to the implicit polymorphism of the virtual machine implementation. For the virtual machine, the `get` operation expects only two arguments, whereas the COQ version expects three arguments (the type A of the elements). So, the compilation scheme first evaluates the last two arguments, then the `get` checks if they are in machine representation. If not, the argument A is evaluated and the three arguments are applied to the COQ implementation of `get`.

4 First Application: The Mini-Rubik

The Mini-Rubik is the pocket version of the famous Rubik cube. It is composed of 8 small cubes only and has 3,674,160 configurations. It is then quite easy to explore them completely with computers. Here, we explain how the property that the Mini-Rubik is always solvable in less than 11 moves has been proved formally.

First, we need to give a model. For this, we use indexation and associate to each configure of the Mini-Rubik a unique number from 1 to 3,674,160. Second, we have to construct a reachability graph – a Cayley graph, using the terminology of group theory. As we are capable of indexing configuration, this is easy. In order to represent a set of configurations, we use an array of 3,674,160 booleans. We prove that it is solvable in 11 moves using an iterative process and two sets of configurations S_A and S_N . The first set S_A contains the configurations that have been reached so far. The second set S_N contains the new configurations that have been reached by the previous iteration. Initially these two set only contain the initial configuration. At each iteration, S_A and S_N are updated with all the configurations that can be reached in one move by one configuration in S_N . After 12 iterations, S_N should be empty.

This application is perfect for testing our extension. First, as $3,674,160 < 2^{31} - 1$, a configuration can be represented by a single native integer. Second, if it is not possible to allocate in COQ an array of 3,674,160 booleans (booleans

are defined as an inductive type with two constructors, so one boolean takes one word in memory), thanks to binary encoding we use an array of machine integer of length $118522 = 3,674,160/31 + 1$. Furthermore, from a given configuration, there are 9 configurations reachable in one move. So this means that look-ups in the array S_A will be 9 times more frequent than updates. This is perfect since our look-ups cost much less than our updates. In [16], we have already presented a formal proof of the Mini-Rubik. In this version, machine integers were available but not efficient arrays. The arrays were implemented using an *ad hoc* functional data-structure. Checking the proof that the Mini-Rubik is solvable in 11 moves took 4 minutes. Modified with our new arrays, not only did it reduce to 10 seconds, but it also greatly simplified the implementation.

5 Second Application: Verifying SAT Traces

The most efficient SAT solvers that are used to prove the unsatisfiability of booleans formulas are all based on the DPLL algorithm with learning (see [10] for a complete introduction). An interesting feature of this algorithm is that very little overhead is needed in order to generate a trace that explains why the formula is unsatisfiable. Formats for traces may slightly vary from one solver to the other but they are all based on the simple resolution rule:

$$\frac{\neg x \vee C \quad x \vee C'}{C \vee C'}$$

The variable x is called the resolution variable, C and C' are clauses, i.e. disjunctions of literals. The reflexive method to prove the unsatisfiability of boolean formulas in COQ works as follows:

- The initial problem is turned by some CNF transformation into a list of clauses. These clauses are called the roots
- The sat solver is called and returns the trace. This trace is composed of a list of resolution chains. Each chain corresponds to a clause that has been learned by the algorithm, so it is a logical consequence of the roots.
- A program written in COQ checks that the trace is correct: it builds the clauses that correspond to the resolution chains and finally checks that the last clause is the empty clause, i.e. \perp is a consequence of the roots.

From the implementation point of view, roots are represented by a list of lists of natural numbers. Each boolean variable x has a unique number n . The literal x is represented as $2n$, $\neg x$ as $2n+1$. The trace is also represented as a list of lists of natural numbers. Each number represents an index of a clause. The indexes are computed with the following convention: roots come first then the clauses built by resolutions. A resolution chain is a list of natural numbers $\{n_1, n_2, \dots, n_k\}$. In order to build the resulting clause, it is traversed from left to right: C_{n_1} is resolved with C_{n_2} , the result is then resolved with C_{n_3} and so on.

For the implementation of the checker, we have directly translated the C code of zVERIFY, the checker of zCHAFF [17], into our functional setting. The checker represents 363 lines of code and its correctness proof is 1621 line long. We use

Problem	Vars	Clauses	zCHAFF	ISABELLE	Coq	Cert	Typing	Check	Array	Parray	zVERIFY
dubois50	150	400	0.00	0.04	0.04	0.00	0.02	0.02	0.00	0.00	0.01
barrel5	1407	5383	0.50	1.10	0.47	0.00	0.32	0.15	0.00	0.00	0.07
barrel6	2306	8931	1.74	10.38	1.15	0.08	0.62	0.45	0.06	0.14	0.14
barrel7	3523	13765	5.20	5.63	1.45	0.17	0.80	0.48	0.07	0.16	0.26
6pipe	15800	394739	42.21	–	24.73	0.98	13.92	9.83	2.05	4.74	2.86
longmult14	7176	22390	408.55	–	73.63	7.72	27.07	38.84	9.10	16.92	7.34
hole11	132	738	14.82	9.36	9.51	0.41	2.96	6.14	1.39	2.89	0.90
hole12	156	949	144.49	61.10	58.28	2.44	18.47	37.38	13.12	16.88	4.85
hole13	182	1197	5048.23	–	1068.30	88.15	387.44	592.72	183.47	275.14	–

Fig. 2. Benchmarking the checker

native integers to represent literals and indexes. Arrays are used for the set of clauses and for a temporary cache to compute the result of a resolution chain. In order to tackle large examples, we had to take a special care in memory usage. For this reason, traces are preprocessed for garbage collecting: we track when a clause is not used anymore, so its index can be reallocated and we share common prefixes in resolution chains. Traces processed by the checker are then list of tagged lists of natural numbers. The tag indicates if the new resulting clauses has to be appended or reallocated in that case it contains the index of the substituted clause.

In order to evaluate what we have done, Figure 2 presents some benchmarks. The machine used for these benchmarks is a Linux Intel Xeon 2.33GHz with 6144 KB of cache and 3 Gigabytes of memory. For each problem, we give:

- the number of variables, the number of clauses and the time for zCHAFF to generate the trace;
- the time of a very similar effort done in ISABELLE/HOL by proof reconstruction [19];
- the time for the reflexive method in COQ (we first give the total time, and then split it in three: the time to parse and preprocess the trace, the time for COQ to typecheck the trace³, and the actual time of the checker);
- the time of the extracted version of the checker running in OCAML with OCAML int and native compilation, first with destructive arrays (Array) then with persistent arrays (Parray);
- the time of zVERIFY.

Times are given in seconds. The symbol – indicates that the verification fails by out-of-memory. We can draw some observations. First, checking the trace is always faster than generating it. While we are slower than zVERIFY (roughly 10 times slower but with a better memory management most probably due to our preprocessing), we are competitive with the proof reconstruction of ISABELLE/HOL. We have a better memory management. We are also faster for all benchmarks except the pigeon-hole problems, where proof reconstruction is as fast as our reflected approach. Second, a fair amount of time in COQ is spent for simply typechecking the generated trace. This clearly indicates that the type-checker of COQ has not yet been tuned to handle very large terms. Finally, the

³ COQ has to typecheck the trace because it is an argument of the call to the checker, so it appears explicitly in the proof term.

checking part in COQ behaves quite well with respect to its extracted version with native code compilation. Remember that the evaluation in COQ is usually comparable to the bytecode compilation of OCAML and between native and bytecode compilations there is usually a factor of 5 to 10. One reason for this good behaviour is that our array primitives in COQ are actually running in their native version. The two versions of the extracted version indicate that, for this kind of application, the cost of using persistent arrays instead of destructive ones is about a factor of 2.

6 Conclusion

In this paper, we have presented how COQ can be extended with some imperative features. This extension increases the trusted computing base of the system but we believe that what we have proposed here is a very good compromise between the impact the extension has on the architecture of the prover and the benefit in term of speed-up in proof checking. Our changes are localised to the abstract machine and its compiler. We didn't have to change any other part of the prover. In particular, we didn't change the logic. We have also developed a systematic and simple methodology to add efficient data-structures with a functional interface to the abstract machine and its compiler. This contributes to the trust one can put in this extension. The methodology has been used to integrate machine integers and persistent arrays.

Some kind of destructive arrays are available in provers like ACL2 [4], ISABELLE/HOL [5] or PVS [14], but some of these techniques are difficult to apply directly to a prover with a rich logic such as COQ and anyway all of them would require a major modification in the architecture of the prover. To our knowledge, the idea of using persistent arrays inside a prover is new. If it does not provide the full power of destructive arrays as in the other provers, for large applications, it gives a clear speed-up with respect to functional arrays. The loss in efficiency with respect to destructive arrays is largely compensated by the fact that we remain in the comfortable setting of functional behaviour.

Our overall goal is not to have an evaluation inside COQ that competes with mainstream programming languages. It is more to have a reasonable computing power within the prover. For example, being able to check the property of the Mini-Rubik in 4 minutes was sufficient enough. The SAT example is more interesting. We manage to get within COQ what was done by extraction in [7]. Without our extension, it would have been impossible to handle large examples. The fact that we could very quickly be competitive with what was achieved by finely-tuned proof reconstruction [19] in HOL and ISABELLE/HOL is clearly good news. It opens new perspectives for the use of reflexive methods inside COQ. Finally, if our initial motivation was efficiency, memory usage has revealed to be sometimes an even more crucial limiting factor. Our machine integers and persistent arrays are much more compact than their corresponding functional representations – or their traditional encoding. Unfortunately, and this is maybe the only drawback of having this light integration, these objects only

exist within the abstract machine. In particular, they cannot be stored in proof objects, therefore have no impact on their size. For this reason, we had to develop an *ad hoc* inductive type in order to store efficiently the traces generated by the SAT solver.

References

1. Baker, H.G.: Shallow Binding Makes Functional Arrays Fast. ACM SIGPLAN notices 26, 145–147 (1991)
2. Barras, B., Grégoire, B.: On the Role of Type Decorations in the Calculus of Inductive Constructions. In: Ong, L. (ed.) CSL 2005. LNCS, vol. 3634, pp. 151–166. Springer, Heidelberg (2005)
3. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science. Springer, Heidelberg (2004)
4. Boyer, R.S., Moore, J.S.: Single-Threaded Objects in ACL2. In: Krishnamurthi, S., Ramakrishnan, C.R. (eds.) PADL 2002. LNCS, vol. 2257, pp. 9–27. Springer, Heidelberg (2002)
5. Bulwahn, L., Krauss, A., Haftmann, F., Erkök, L., Matthews, J.: Imperative Functional Programming with Isabelle/HOL. In: Mohamed, O.A., Muñoz, C., Tahar, S. (eds.) TPHOLs 2008. LNCS, vol. 5170, pp. 134–149. Springer, Heidelberg (2008)
6. Conchon, S., Filliâtre, J.-C.: A persistent union-find data structure. In: ACM Workshop on ML, pp. 37–46 (2007)
7. Darbari, A., Fischer, B., Marques-Silva, J.: Formalizing a SAT Proof Checker in Coq. First Coq Workshop, published as technical report tum-i0919 of the Technical University of Munich (2009)
8. Gonthier, G.: Formal Proof – The Four-Color Theorem. Notices of the AMS 55(11) (2008)
9. Grégoire, B., Leroy, X.: A compiled implementation of strong reduction. In: ICFP, pp. 235–246 (2002)
10. Kroening, D., Strichman, O.: Decision Procedures, An Algorithmic Point of View. Texts in Theoretical Computer Science. Springer, Heidelberg (2008)
11. Leroy, X.: The ZINC experiment: an economical implementation of the ML language. Technical report 117, INRIA (1990)
12. Leroy, X.: Objective Caml (1997), <http://ocaml.inria.fr/>
13. Okasaki, C.: Purely Functional Data Structures. Cambridge University Press, Cambridge (1998)
14. Shankar, N.: Static Analysis for Safe Destructive Updates in a Functional Language. In: Pettorossi, A. (ed.) LOPSTR 2001. LNCS, vol. 2372, pp. 1–24. Springer, Heidelberg (2002)
15. Swierstra, W.: A Hoare Logic for the State Monad. In: Urban, C. (ed.) TPHOLs 2009. LNCS, vol. 5674, pp. 440–451. Springer, Heidelberg (2009)
16. Théry, L.: Proof Pearl: Revisiting the Mini-Rubik in Coq. In: Mohamed, O.A., Muñoz, C., Tahar, S. (eds.) TPHOLs 2008. LNCS, vol. 5170, pp. 310–319. Springer, Heidelberg (2008)
17. Princeton University. zChaff, <http://www.princeton.edu/~chaff/zchaff.html>
18. Wadler, P.: Monads for Functional Programming. In: Jeuring, J., Meijer, E. (eds.) AFP 1995. LNCS, vol. 925, pp. 24–52. Springer, Heidelberg (1995)
19. Weber, T., Amjad, H.: Efficiently checking propositional refutations in HOL theorem provers. J. Applied Logic 7(1), 26–40 (2009)